

## Towards Effective E-Participation of Citizens in the European Union: The Development of AskThePublic

Nils Messerschmidt  
SnT, University of Luxembourg  
[nils.messerschmidt@uni.lu](mailto:nils.messerschmidt@uni.lu)

Kilian Sprenkamp  
University of Zurich  
[kilian.sprenkamp@uzh.ch](mailto:kilian.sprenkamp@uzh.ch)

Amir Sartipi  
SnT, University of Luxembourg  
[amir.sartipi@uni.lu](mailto:amir.sartipi@uni.lu)

Xiaohui Wu  
University of Zurich  
[xiaohui.wu@uzh.ch](mailto:xiaohui.wu@uzh.ch)

Igor Tchappi  
SnT, University of Luxembourg  
[igor.tchappi@uni.lu](mailto:igor.tchappi@uni.lu)

Liudmila Zavolokina  
University of Lausanne  
[liudmila.zavolokina@unil.ch](mailto:liudmila.zavolokina@unil.ch)

Gilbert Fridgen  
SnT, University of Luxembourg  
[gilbert.fridgen@uni.lu](mailto:gilbert.fridgen@uni.lu)

### Abstract

*E-participation platforms are an important asset for governments in increasing trust and fostering democratic societies. By engaging public and private institutions and individuals, policymakers can make informed and inclusive decisions. However, current approaches of primarily static nature struggle to integrate citizen feedback effectively. Drawing on the Media Richness Theory and applying the Design Science Research method, we explore how a chatbot can address these shortcomings to improve the decision-making abilities for primary stakeholders of e-participation platforms. Leveraging the "Have Your Say" platform, which solicits feedback on initiatives and regulations by the European Commission, a Large Language Model-based chatbot, called AskThePublic is created, providing policymakers, journalists, researchers, and interested citizens with a convenient channel to explore and engage with citizen input. Evaluating AskThePublic in 11 semi-structured interviews with public sector-affiliated experts, we find that the interviewees value the interactive and structured responses as well as enhanced language capabilities.*

**Keywords:** E-Participation, Citizen Involvement, Large Language Models, Retrieval-Augmented Generation, Media-Richness-Theory.

### 1. Introduction

Fueled by a decade of crises since the global financial crisis of 2007-2008, including the Euro crisis, the refugee crisis, and the COVID-19 pandemic, polarization is on the rise across the European

continent (Matthijs, 2020). The increasing support for Euroscepticism and anti-political establishment parties is a particular concern for the bloc in major European countries such as France, Germany, and Italy, who are the largest net contributors to the European Union (EU) budget (Begg, 2017) and whose citizens have mostly enjoyed the benefits of European cooperation since the 1950s. Meanwhile, the Brexit movement has led the UK to withdraw from the EU entirely. A common theme for anti-establishment parties is populist, anti-system or protest tendencies, and, in fewer cases, even radical or extremist ones (Casal Bértoa & Rama, 2021).

One proposed way to increase trust in governments and thus decrease anti-systemic tendencies is citizen participation (Kim & Lee, 2012). More recently, governmental agencies have been promoting e-participation tools, such as the *Have Your Say*<sup>1</sup> platform, which invites its users to provide feedback to ongoing initiatives led by the European Commission. Despite successfully receiving citizen feedback, the platform lacks effective ways to structure, analyze and make meaningful conclusions out of this feedback, thus decreasing the ability to leverage the platform for primary stakeholders. These include EU policymakers, who aim to incorporate various perspectives for inclusive decision-making; researchers and journalists, who can leverage the provided feedback for scientific or popular research, and lastly, interested EU citizens, who can use the platform for information purposes.

Addressing the platform's shortcomings, the Media Richness Theory (MRT) offers a theoretical explanation for the various degrees of effective communication depending on the chosen medium. In accordance

<sup>1</sup><https://ec.europa.eu/info/law/better-regulation/have-your-say>

with the theory, a medium is more effective, the more interactive it is (Daft and Lengel, 1984; 1986). Assuming this theoretical lens, this would call for more interactivity on the *Have Your Say* platform to be of more effective use to its stakeholders, in line with the concept of "conversational governments" (Baldauf & Zimmermann, 2020). One method to increase interactivity while simultaneously avoiding the common issue of higher cost associated with higher media richness (Androusoy et al., 2019), e.g., face-to-face interactions, is the use of Large Language Models (LLMs), which have gained increasing public interest since the release of ChatGPT in 2022. This poses the following research question (RQ):

*How can an LLM-based chatbot leveraging e-participation data be designed to increase the informedness of primary stakeholders involved in EU policy-making processes?*

Our project resulted in the development of a proof-of-concept prototype based on the Design Science Research (DSR) methodology (Hevner et al., 2004; Peffers et al., 2007), i.e., an LLM-enabled chatbot named *AskThePublic*, designed to enhance e-participation mechanisms by facilitating the effective feedback analysis of its primary stakeholders.

## 2. Related Work

### 2.1. Citizen Involvement in the EU

Since the 1960s, governments have increasingly recognized public participation as a cornerstone in improving the legitimacy, accountability, and transparency of decision-making (Lourenço & Costa, 2007). Since then, an increasing number of legislative systems have solidified the right to public participation in law, recognizing it as a fundamental right to contribute meaningfully to decision-making processes (Armenia & Irina, 2022). E-participation is a more recent form of citizen involvement using information and communication technologies for electronic government-citizen consultations, thus promoting engagement and collaboration in decision-making (Alarabiat et al., 2021; Macintosh, 2004; Sæbø et al., 2008). It strengthens transparency, advances social inclusion, and fosters democratic societies (Alarabiat et al., 2021). However, implementing e-participation is challenging due to low citizen adoption (Hassan & Hamari, 2020; Venkatesh et al., 2000). Digital literacy gaps, restricted internet access, and poor advertisement can hinder its effectiveness and impact (Alarabiat et al., 2021; Toots, 2019).

EU member states have developed a tradition of

digital citizen dialogue to reinforce public engagement and active participation in democratic governance (Mariani et al., 2024). Policymakers at the local, regional, national, and European levels further increased e-participation efforts in line with the 2020 Digital Agenda as part of the Europe 2020 strategy: Panopoulou et al. (2009) identified 255 different e-participation initiatives within the EU with varying scopes. E-participation employs a variety of tools and platforms to engage citizens, including social media channels like Facebook, YouTube, and X (ex-Twitter). Additionally, dedicated online portals, including bespoke tools such as *Have Your Say* were created by governmental agencies to drive citizen involvement. Given that legitimate e-participation data is received, there is fragmented knowledge on how to analyze it, particularly for large datasets. For example, Aitamurto et al. (2016) used concept extraction and sentiment analysis on crowd-sourced policy data to evaluate how citizen input shaped a city's transportation plan, while Hagen et al. (2015) applied textual analysis and topic modeling to classify e-petition texts, thereby exploring policy impacts of emerging themes. Similarly, Arana-Catania et al. (2021) combined topic modeling and text summarization in a direct democracy platform to address information overload and support participation. However, all of these systems have in common that they require prior knowledge of Machine Learning (ML) tools: A gap that LLMs can bridge merely with natural language.

### 2.2. Media Richness Theory

The MRT, initially developed by Daft and Lengel (1984; 1986), provides a theoretical framework for understanding how entities, including public and private organizations and individuals, select communication media to process information effectively. It posits that communication channels vary in capacity to reduce equivocality and manage uncertainty. These early works show that successful entities carefully match the richness of their communication methods with the complexity and uncertainty of the messages they need to share. Critically, they highlighted that organizational success depends on balancing the richness of the media used with the complexity and ambiguity of the tasks, emphasizing that a core challenge often stems from a lack of clarity rather than insufficient data (Daft & Lengel, 1984). The MRT has been applied to the communication channels governments offer, such as face-to-face, telephone, messenger services, websites, or social media (Sanina et al., 2017). Although face-to-face interaction allows for a high richness,

it requires a high cost for governments to operate (Androusofoulou et al., 2019).

Building on these existing communication channels, the concept of "conversational government" (Baldauf & Zimmermann, 2020) offers a solution. It combines the high level of richness found in face-to-face interactions with the cost-effectiveness of automated systems. Utilizing chatbots, governments can thus provide real-time, personalized responses to citizens while lowering costs (Androusofoulou et al., 2019). Further, there are several examples of how chatbots aligned with MRT principles can advance government-citizen interaction beyond simply providing static information, enabling more dynamic, context-sensitive, and user-focused exchanges. For instance, building on the foundations of automated conversation and enhanced data accessibility, Cortés-Cediel et al. (2023) developed a holistic framework for employing chatbots in public administration, demonstrating how chatbots can enrich information retrieval, promote greater transparency, and broaden the scope of citizen engagement with governmental services. Similarly, Segura-Tinoco et al. (2022) illustrate how chatbots facilitate meaningful deliberation by extracting and summarizing complex policy debates, thereby assisting citizens in navigating large-scale civic discussions. Last, LLMs have been utilized to analyze community data, offering insight into refugee needs and supporting data-driven decision-making, thus voicing the needs of refugees and including them in citizen participation (Sprenkamp et al., 2025). Collectively, these works highlight the potential of chatbots to deliver rich, context-aware communication channels that directly reflect MRT's emphasis on managing complexity and ambiguity through media choices, ultimately contributing to more responsive, inclusive, and effective governance.

### 2.3. Large Language Models

Artificial Intelligence (AI) has seen significant advancements. In Natural Language Processing (NLP), early static embeddings like GloVe (Pennington et al., 2014) were replaced by context-aware methods following the introduction of the Transformer architecture (Vaswani et al., 2017), which employs self-attention for parallel sequence processing. Generative models further extended NLP capabilities through pre-training on large corpora followed by fine-tuning (Radford, 2018), culminating in LLMs like GPT-4 (Achiam et al., 2024), which integrate multi-modal processing and reinforcement learning for alignment with human preferences.

Retrieval-Augmented Generation (RAG) methods

enhance LLMs by integrating external knowledge. In a typical RAG workflow, the input query is first converted into a vector embedding using an encoder, and all data in the external database are similarly embedded and stored for efficient retrieval. When a query is made, its embedding is used to search the database for the most relevant embeddings, and the retrieved information is then combined with the LLM to generate a more accurate and context-rich response (Lewis et al., 2020).

LLMs and NLP techniques have been applied in various e-participation contexts to enhance public engagement, streamline service delivery, and support decision-making processes. In citizen input analysis, AI-based systems have automated feedback processing, reducing manual workload while ensuring consistent and scalable results (Borchers et al., 2024). In public administration, NLP tools have facilitated policy-making by bridging communication gaps between policymakers and citizens (Guridi et al., 2024). Similarly, LLM-powered chatbots have been employed in public services to improve service accessibility, though challenges linked to privacy, transparency, and trust remain critical concerns (Dreyling et al., 2024). Moreover, citizen complaint management systems using contextual feedback mechanisms driven by LLMs have enhanced user trust and engagement by providing meaningful, real-time responses (Karren et al., 2024). These applications demonstrate the practical potential of LLMs in enabling more interactive, transparent, and efficient e-participation systems over previous ML-based implementations. However, this novel approach is yet to be applied to a real-world context.

### 3. Methodology

This study is part of a larger project focused on improving EU citizen participation through an LLM-enabled chatbot. As such, enabling policymakers to effectively analyze e-participation data is a necessary first step to drive more active citizen participation. To create a solution for this real-world problem (Nunamaker Jr et al., 2015), we employ the DSR methodology (Hevner et al., 2004; Peffers et al., 2007), namely, 1) problem identification, 2) definition of solution objectives, 3) design and development, 4) demonstration, 5) evaluation, and 6) communication.

To (1) identify the problems related to EU citizen participation and (2) derive solution objectives, we observed the *Have Your Say* platform and compared it to insights from academic literature. These findings informed the definition of our problem statements and solution objectives, which aimed to address these specific challenges through the development

of an advanced chatbot. For the (3) design and development phase, we iteratively created our prototype—a web-based LLM-powered chatbot, called *AskThePublic*<sup>2</sup>, the source code can be found on GitHub<sup>3</sup>. It is built with React for the frontend and FastAPI for the backend, uses LangChain for NLP tasks, interacts with MongoDB and OpenAI, and is containerized with Docker to support reproducibility and streamlined deployment.

Next, it is essential to (4) demonstrate and (5) evaluate the artifacts created to determine if the research objectives have been met. For the demonstration and evaluation, we conducted 11 semi-structured interviews (ID1-ID11). In order to be eligible, participants needed to be EU citizen who were affiliated with the local government or EU institutions through research projects. We gathered participants from five countries (France, Germany, Greece, Italy, and Luxembourg) and varying backgrounds, including law (8,10), finance (ID7,9), security and privacy (ID3-6), and energy (ID1-2,11). The participants had no previous work experience related to citizen participation or the *Have Your Say* platform. Prior to the session, a short survey was administered to assess the participants' perceptions of the EU. During the sessions, participants interacted with the chatbot, performing tasks such as querying content on specific EU initiatives, focusing on the usability and effectiveness of the solution within their field of expertise, as well as informing directions for future development based on their professional assessment. The sessions, ranging between 44 and 74 minutes, have been recorded, transcribed, and analyzed using a qualitative approach (Saldaña, 2021) consistent with Gioia's methodology for inductive research (Gioia et al., 2013). The session guideline is available upon request. Finally, the (6) communication of this research is achieved through the present article.

## 4. Results

### 4.1. Problem Statement & Solution Objectives

The *Have Your Say* platform provides several opportunities for broad citizen participation. Yet, it faces several limitations that can hinder its effectiveness. We found three major problems in analyzing data within the *Have Your Say* platform, presented in the following. Table 1 matches those with the solution objectives, and offers an overview of which participants provided a favorable evaluation, with the remaining ones provided either neutral, or critical feedback of the solution.

**Lack of Interaction** constitutes a major issue, as the platform primarily facilitates one-way communication. Citizens can submit feedback, but the lack of real-time responses or interactive features limits their involvement, aligning with the lower levels of the public participation spectrum, namely informing, consulting, and involving rather than collaborating or empowering (Shipley & Utz, 2012). In particular, when users interact with the platform, they must first navigate to an initiative of interest, which can be challenging since most EU citizens are unclear about the objectives and context of various initiatives. After selecting an initiative, users are expected to navigate the timeline and supporting documentation, often consisting of text-heavy external files. The feedback submitted on the platform is ordered by time of submission, while lacking filtering or clustering options. The proposed solution should address this by providing interactive functions for more effective communication (Daft and Lengel 1984; 1986).

**Table 1. Overview of Selected DSR Steps**

Problem Statements	Solution Objectives	Positive Evaluation
<i>Lack of Interaction</i> (Daft and Lengel, 1984; 1986; Shipley and Utz, 2012)	Real-Time Interaction with Citizen Feedback	ID 1-3, 6-11
<i>Lack of Structure</i> (Arana-Catania et al., 2021; Armenia & Irina, 2022; Lourenço & Costa, 2007)	1) Structured Feedback 2) Stakeholder Clustering	1) ID 1, 3-4, 9-11 2) ID 1-5, 7-8, 10-11
<i>Lack of Understanding</i> (Alarabiat et al., 2021; Toots, 2019; Wright, 2009)	Automated Multilingual Synthesis	ID 1, 10-11

**Lack of Structure** is another major issue in feedback collection due to the use of natural language. Responses vary in length, focus, and clarity, making it difficult for stakeholders to extract meaningful insights without automated tools to categorize and filter input (Arana-Catania et al., 2021). In particular, to polarizing topics, over 1000 feedback statements are submitted, requiring the user to distinguish qualified feedback from unqualified one, particularly from anonymous platform users. Frequently, feedback does not relate to the topic of the initiative or merely scratches the surface, making it difficult for EU citizens, but also for policymakers to fully comprehend the feedback and to leverage it as part of the implementation process. The proposed solution is expected to offer tools that make the structuring of the feedback by stakeholder groups more efficient (Armania & Irina, 2022; Lourenço & Costa, 2007).

**Lack of Understanding** is a third issue stemming from the diverse linguistic and cultural landscape of the EU. As feedback may be written in any of the 24 official languages of the EU, without providing tools for integrated translation, barriers to understanding and

<sup>2</sup><http://askthepublic.eu/>

<sup>3</sup><https://github.com/HICSS26/AskThePublic>

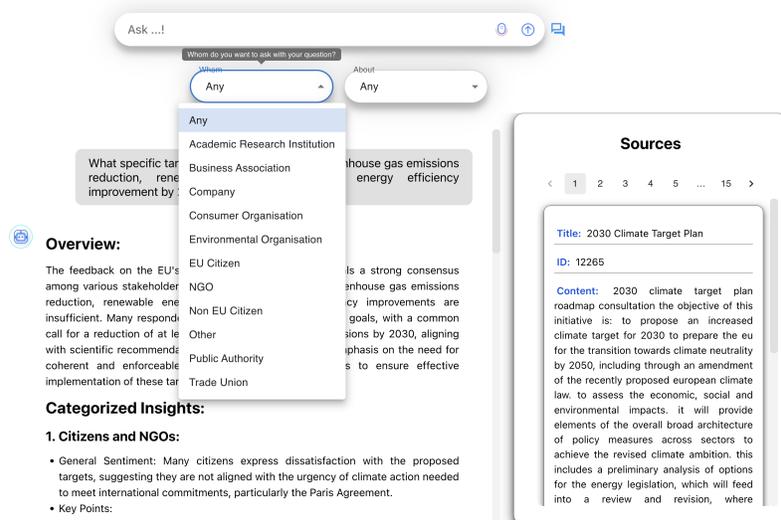


Figure 1. Analysis page displaying a detailed response to the user's query

accessibility arise for both citizens and policymakers (Wright, 2009). In particular, certain initiatives may receive mainstream media attention in selected countries. In those circumstances, it is common to find that citizens of those countries submit most of the available feedback statements. Thus, the proposed solution is expected to offer output comprehensible by all EU citizens (Alarabiat et al., 2021; Toots, 2019).

#### 4.2. Artifact Design

To address the problems mentioned in Section 4.1, we developed *AskThePublic* (Figure 1). It provides an intuitive interface for user interaction, and at its core, it utilizes RAG to answer user queries effectively.

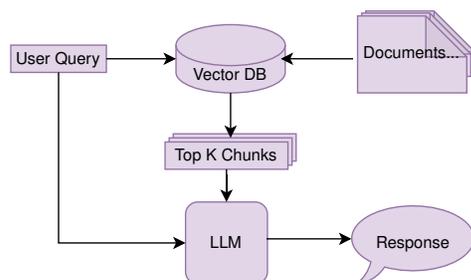


Figure 2. The RAG pipeline

The interface is designed to be accessible and straightforward. It contains a search bar that allows users to input queries in natural language. This central element is complemented by two drop-down filters: “Whom” and “About”. These filters enable users to

narrow their queries by selecting specific groups they want to inquire about and the topics they are interested in exploring further. By default, the system searches across all defined groups and all stored comments within various topics to provide comprehensive responses.

RAG is the core framework driving the system, combining advanced information retrieval techniques with the generative capabilities of LLMs. We rely on RAG because even models like GPT-4 can produce inaccurate information, and RAG has proven effective in mitigating this issue. To support this, the system starts by scraping feedback data from the *Have Your Say* platform. This feedback, often unstructured and multilingual, is preprocessed to remove noise and inconsistencies. The data is then transformed into dense vector embeddings using OpenAI’s text-embedding-3-small, which encode the semantic meaning of the text into a high-dimensional numerical representation. These embeddings are stored in a MongoDB Atlas vector database, enabling efficient semantic search during retrieval.

When a user submits a query, the system first processes it by generating a vector embedding for the input, matching it against the database of stored feedback using cosine similarity. The system then returns the  $K$  entries from the database with the highest cosine similarity as the context for LLM. At this stage, the system leverages GPT-4o-mini by feeding the retrieved context along with the user’s query into the LLM. This integration delivers reliable, insightful summaries and actionable recommendations in an intuitive, human-like format. The RAG pipeline

process to generate final response is shown in Figure 2.

The output of the chatbot is structured into three parts. The first part provides an overview, highlighting key themes and sentiments from public opinions. The second part offers categorized insights, showcasing feedback segmented by different user groups and summarizing key points for each group. Finally, the third part delivers actionable insights, presenting two to three specific, evidence-based recommendations or highlighting conflicts that require resolution. To increase user confidence and transparency, the system includes a Source section, which highlights the resources used to generate the answer to the given query. A new chat icon positioned beside the search bar enables users to refresh the page and initiate a new query.

### 4.3. Demonstration and Evaluation

A total of 357 in-vivo codes were derived during the analysis of the 11 semi-structured interviews. These were then clustered into the solution objective categories, namely, real-time interaction (85 in-vivo codes), structured feedback management (69 in-vivo codes), and multilingual synthesis (63 in-vivo codes). A fourth category, prospective stakeholder usage, was introduced to inform future research directions, covering the remaining 140 in-vivo codes. The analysis of the pre-study shows that the participants shared a rather positive perception of the EU, with an average of 8.0 on a scale from 1 (very negative) to 10 (very positive). In the following, the results of the qualitative assessment according to the solution objectives are presented.

**Real-Time Interaction** Although some participants (ID4-6) were sceptical whether the RAG capabilities provide enough benefit over established chatbots (“Everything sounds reasonable. I probably would need to ask the same question to ChatGPT, because, based on what I read, though it is neither right nor wrong, I do not see why we need another chatbot.”, ID5), the majority of the participants (ID 1-2,7-9) indicated they are more likely to use *AskThePublic* over the *Have Your Say* platform due to the enhanced interaction: “I would be much more likely to use it than to browse through every single initiative or directive and to read the comments separately, this is obviously made much easier. And it also reduces the entry barrier, so to say, to inform myself. (ID8)” One participant narrowed the benefit of the solution down to a single statement: “Basically, the more advanced search engine for public opinions. (ID4)” Participants have particularly praised its intuitiveness: “It’s easy. In the beginning, you are doing a test. If you are not sure, then you see how it works. (ID9)” The tool was further noted

to be well suited for information purposes (ID 2-3, 6-10), particularly in cases where one does not have pre-existing knowledge of a domain: “So, if you do not have a very precise question in mind, but you rather want to assess opinion in certain areas of topics, then, it is being summarized really well. [...] I really like how it works because you can receive various opinions in a fairly well-structured and simple way. (ID3)” Moreover, several participants (ID 1-2,7,10-11) noted the various features of the chatbot positively: “This is something I really enjoy, that I have my text on the left side, and on the right side the sources for validation. (ID10)”

**Structured Feedback Management** Participants generally view the structure favorably (ID 1,3-4,9-11), as exemplified by ID10: “I actually like that it gives a broad answer at the top, then going into more insights and also having an implication section. So it has this tunneling effect.” The structure allows its users to conduct further research: “It is not too long, not too short, so it provides you an overview and from there you can take it further and research yourself. (ID9)” Participants particularly highlighted the ability to engage with specific stakeholder groups directly (ID1-5, 7-8, 10-11): “If you ask specific groups of people questions, it seems that you can engage with these people. You can engage with the answers that they have already provided on the platform. (ID7)” At the same time, through RAG-based computation, *AskThePublic* only retrieves stakeholder feedback that relates directly to the query, in stark contrast to the *Have Your Say* platform. While this was noted positively, all participants questioned how the sources are being selected (ID1-11): “If you say, 1000 people have responded, but now I get four random universities as output, I question the integrity of the process, what makes them stand out (ID6)?” Participants noted an apparent lack of credibility when it comes to certain stakeholders that are being displayed: “Perhaps the German AI Association I have heard before, but I do not know any other, so I cannot judge if that is a source where perhaps three people are sitting who simply do not have any relevance, or if it is some sort of NGO which is truly important. (ID3)” One participant suspected a potential bias towards more polarizing opinions that are selected to be displayed: “So for example, not everything about dynamic tariffs is bad, but it feels like this kind of overview right now focuses pretty heavily on the bad side. (ID11)” When being asked very specific questions, the chatbot was found to create answers artificially. This is recognized as a major threat to the validity: “So if you go and search yourself, and you find out that there is nothing, this is very bad for the reputation of the system. (ID9)”

**Multilingual Synthesis** The solution translates stakeholder feedback given in any official EU language to the language the question was asked in, making it easier for EU citizens to understand the feedback (ID1, 10-11), as exemplified by ID 10: *“The generation of my grandparents, or even my parents, they do not speak English so well. So for them, this is valuable because if they are struggling to understand the text, they will go to some tabloid newspaper, even if they do not provide accurate information.”* However, it was noted that this only applies to the immediate text output from the chatbot and does not apply to the sources panel. Further, finding the right sourcing balance is commonly criticized (ID4-5, 10-11). On the one hand, participants found larger countries with bigger impact most relevant, as exemplified by ID 4 (*“In this question, I only got answers from Germany and Slovakia. Germany, I understand, but as for Slovakia, frankly, I do not understand why it is being displayed”*), on the other hand, participants noted that they would like to have a bigger variety, as exemplified by ID 10 (*“I would prefer to have three bullet points from three different countries instead of having the majority from one country.”*)

## 5. Discussion

Utilizing DSR, we are able to answer our RQ. We do so by discussing the solution objectives and design features of the designed artifact *AskThePublic* (Section 5.1). Further, we give insights on how different stakeholders can utilize the given design (Section 5.2).

### 5.1. Designed Artifact and Future Research

The evaluation of the designed artifact demonstrates that each of our three solution objectives holds significant promise for enhancing e-participation.

Participants appreciated the chatbot’s intuitiveness and ease of use, noting that its real-time interaction capabilities significantly enhanced user engagement compared to traditional platforms. This finding aligns with the MRT, which posits that richer media facilitate more effective communication ((Daft & Lengel, 1984). The chatbot’s ability to provide immediate, interactive responses mirrors the high richness of face-to-face communication, thereby increasing its effectiveness in engaging users as a practical example of conversational government (Baldauf & Zimmermann, 2020).

*AskThePublic’s* ability to organize unstructured feedback into coherent and actionable insights was another key strength noted by participants. However, in the absence of relevant information to display, the chatbot might hallucinate, thus leading some participants to question the added benefits over LLMs

without RAG, underlining that the system is only as good as the available data. Provided there is relevant data, the system effectively filters and structures feedback, ensuring that only relevant information is presented (Lewis et al., 2020), thus focusing on transparency, a quality needed in e-participation systems (Alarabiat et al., 2021). While this concept was rooted in Sprenkamp et al. (2025), who analyzed social media data to make it available for policy-making, *AskThePublic* allows for generating insights from e-participation data in natural language, a large contribution over prior ML-based tools (Armenia & Irina, 2022; Hassan & Hamari, 2020; Mariani et al., 2024). As participants found this structured format facilitated easier navigation and deeper understanding of the e-participation data, we foresee that systems like *AskThePublic* will attract higher adaptation rates, thus addressing a common problem among modern e-participation tools (Hassan & Hamari, 2020; Venkatesh et al., 2000).

The multilingual capabilities of *AskThePublic* helped make feedback accessible across the EU’s 24 official languages. This feature directly addresses language barriers identified as significant determinants affecting e-participation (Alarabiat et al., 2021). By utilizing LLMs for translation and synthesis, the chatbot fosters a more inclusive environment, aligning with the principles of conversational government that advocate for language-inclusive communication channels (Baldauf & Zimmermann, 2020). Further, while the chatbot answers in any of the official EU languages, it further synthesizes context from various source documents in the RAG-system (Sprenkamp et al., 2025). Concerns regarding the balance and diversity of sources were raised, indicating that the chatbot should implement strategies to ensure a more equitable distribution of feedback from various countries. Notably, the current embedding fetching mechanism prioritizes data points in the language of the query, which may inadvertently limit the diversity of retrieved multilingual feedback, highlighting the need to balance language prioritization within the retrieval process. These limitations highlight the need for comprehensive multilingual support to further inclusion.

In addition, users proposed new features for *AskThePublic*. To mitigate the bias towards polarizing opinions and ensure a more balanced representation of feedback, the system could incorporate algorithms that promote diversity in source selection, aligning with research on unbiased data representation in public administration (Alarabiat et al., 2021). Another option is to utilize custom models based on open-source projects. Currently, we utilize the GPT API, a

closed-source model. Thus, the training data and potential biases cannot be identified; here, open-source LLMs like EURO-LLM (Martins et al., 2024) could be leveraged to align more closely with EU values of transparency and data sovereignty. While open-source models often lack the accuracy of closed-source models, this has been recently challenged by the development of DeepSeek (Guo et al., 2025), which exceeds the efficiency of state-of-the-art LLMs while being on par in accuracy. Thus, we expect governments to be empowered in the future to utilize open-source solutions. Several users flagged that *AskThePublic* could integrate unique functionalities not offered by existing chatbots like ChatGPT, such as specialized data visualization tools or tailored stakeholder engagement features, enhancing its distinct value proposition (Borchers et al., 2024). Furthermore, expanding the chatbot's integration capabilities with other e-participation platforms and social media channels can broaden its accessibility and usability, as highlighted in related works on digital engagement strategies (Macintosh, 2004). By implementing these recommendations, future iterations of *AskThePublic* can enhance its effectiveness, inclusivity, and trust, thereby further supporting democratic engagement.

Last, in line with MRT (Daft & Lengel, 1984, 1986), communication media that allow immediate feedback and multiple cues are more effective at reducing equivocality. *AskThePublic*'s chat interface, by offering real-time, two-way dialogue and transparent sourcing, functions as a richer medium compared to the static Have Your Say portal. This increased richness directly maps to participants' reports of greater engagement and clearer understanding.

## 5.2. Possible User Groups

**EU policymakers** can leverage *AskThePublic* to gather and analyze feedback more effectively, thus informing more inclusive policy decisions, which are based on public data. By incorporating features that balance data distribution across different demographics, the chatbot can enhance both the quality and reliability of the feedback used in policy-making, addressing challenges related to sampling bias and data integrity (Alarabiat et al., 2021).

**For politically engaged EU citizens**, *AskThePublic* holds significant potential by providing an intuitive and interactive platform for accessing and analyzing policy feedback. Participants further valued the empowerment and active participation the tool could foster (ID8, ID6, ID10). By making feedback more accessible and structured, the chatbot can enhance citizen engagement

and address common barriers to e-participation, such as lack of information and perceived impact.

**For journalists and researchers**, *AskThePublic* provides a valuable tool for extracting feedback on legislation. Thus, they can identify the sentiments of citizens or organizations towards a given topic, providing researchers and journalists with a reliable and unique source that can be used for publications. However, future iterations of *AskThePublic* need to design clear links to original feedback that can support accurate reporting and data-driven research.

## 6. Conclusion

While e-participation has been in the focus of the academic discourse (Armenia & Irina, 2022; Hassan & Hamari, 2020; Mariani et al., 2024), methods to analyze the submitted feedback are sparse and focusing on the fragmented usage of ML methods to classify citizen needs (Aitamurto et al., 2016; Arana-Catania et al., 2021; Hagen et al., 2015). The advent of LLMs enables novel applications that utilize NLP to analyze citizen feedback, needing little to no understanding of the technology, as shown through *AskThePublic*.

Using the DSR methodology (Hevner et al., 2004; Peffers et al., 2007), the development of *AskThePublic* has made three contributions. First, we provide an initial tool for the analysis of citizen feedback using LLMs, which are open to the general public, enabling users to ask questions without the need for expertise in ML or computational tools. Second, we give design recommendations for tools enabling the structured analysis of e-participation. We hope that these recommendations are useful for further development in academia and the public sector. Third, we identify the first user groups for our tool.

However, our study is not without limitations. Our evaluation was solely done with a rather small user group with largely positive views of the EU, leaving room for broader evaluations. Moreover, we developed and tested *AskThePublic* solely using the GPT-4 API, creating a bias towards the training data of this closed-source model. We see these limitations as opportunities for future work. Later, iterations of *AskThePublic* or similar tools should further evaluate our design recommendations with the goal of generating long-lasting design principles (Gregor et al., 2020). Moreover, the effect of different closed and open-source LLMs should be tested to align with the values of the EU, e.g., EURO-LLM (Martins et al., 2024).

**Acknowledgements** This research was funded in part by the Luxembourg National Research Fund (FNR)

and PayPal, PEARL grant reference 13342933/Gilbert Fridgen. We also thank the University of Zurich and the Digital Society Initiative for (partially) financing the study through Liudmila Zavolokina's DIZH postdoc fellowship. For the purpose of open access, and in fulfillment of the obligations arising from the grant agreement, the author has applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.

## References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. (2024). Gpt-4 technical report.
- Aitamurto, T., Chen, K., Cherif, A., Galli, J. S., & Santana, L. (2016). Civic crowdanalytics: Making sense of crowdsourced civic input with big data tools. *Proc of the 20th Int Academic Mindtrek Conf*, 86–94.
- Alarabiat, A., Soares, D., & Estevez, E. (2021). Determinants of citizens' intention to engage in government-led electronic participation initiatives through facebook. *Government Information Quarterly*, 38(1), 101537.
- Androustoupoulou, A., Karacapilidis, N., Loukis, E., & Charalabidis, Y. (2019). Transforming the communication between citizens and government through ai-guided chatbots. *Government information quarterly*, 36(2), 358–367.
- Arana-Catania, M., Lier, F.-A. V., Procter, R., Tkachenko, N., He, Y., Zubiaga, A., & Liakata, M. (2021). Citizen participation and machine learning for a better democracy. *Digital Government: Research and Practice*, 2(3), 1–22.
- Armenia, A., & Irina, G. (2022). E-participation in europe, a comparative perspective. *Public Administration Issues*, 5(1), 7–29.
- Baldauf, M., & Zimmermann, H.-D. (2020). Towards conversational e-government: An experts' perspective on requirements and opportunities of voice-based citizen services. *HCI in Business, Government and Organizations: 7th Int. Conf., HCIBGO 2020*, 22, 3–14.
- Begg, I. (2017). The eu budget after 2020. *SIEPS European Policy Analysis*, 1–9.
- Borchers, M., Cao, T.-B., Tavanapour, N., & Bittner, E. A. (2024). Designing ai-based systems to support the analysis of citizens' inputs from e-participation.
- Casal Bértoa, F., & Rama, J. (2021). Polarization: What do we know and what can we do about it? *Frontiers in Political Science*, 3, 687695.
- Cortés-Cediel, M. E., Segura-Tinoco, A., Cantador, I., & Bolívar, M. P. R. (2023). Trends and challenges of e-government chatbots: Advances in exploring open government data and citizen participation content. *Government Information Quarterly*, 40(4), 101877.
- Daft, R. L., & Lengel, R. H. (1984). Information richness: A new approach to managerial behavior and organizational design. *Research in organizational behavior*.
- Daft, R. L., & Lengel, R. H. (1986). Organizational information requirements, media richness and structural design. *Management science*, 32(5), 554–571.
- Dreyling, R., Koppel, T., Tammet, T., & Pappel, I. (2024). Challenges of generative ai chatbots in public services-an integrative review.
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the gioia methodology. *Organizational research methods*, 16(1), 15–31.
- Gregor, S., Chandra Kruse, L., & Seidel, S. (2020). Research perspectives: The anatomy of a design principle. *Journal of the Association for Information Systems*, 21(6), 2.
- Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al. (2025). Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning.
- Guridi, J. A., Cheyre, C., & Yang, Q. (2024). Thoughtful adoption of nlp for civic participation: Understanding differences among policymakers.
- Hagen, L., Harrison, T. M., Uzuner, Ö., Fake, T., Lamanna, D., & Kotfila, C. (2015). Introducing textual analysis tools for policy informatics: A case study of e-petitions. *Proceedings of the 16th annual international conference on digital government research*, 10–19.
- Hassan, L., & Hamari, J. (2020). Gameful civic engagement: A review of the literature on gamification of e-participation. *Government Information Quarterly*, 37(3), 101461.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS quarterly*, 75–105.

- Karren, K., Schmitz, M., & Schaffer, S. (2024). Improving conversational user interfaces for citizen complaint management through enhanced contextual feedback. *Proceedings of the 6th ACM Conference on Conversational User Interfaces*.
- Kim, S., & Lee, J. (2012). E-participation, transparency, and trust in local government. *Public administration review*, 72(6), 819–828.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., et al. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.
- Lourenço, R. P., & Costa, J. P. (2007). Incorporating citizens' views in local policy decision making processes. *Decision Support Systems*, 43(4), 1499–1511.
- Macintosh, A. (2004). Characterizing e-participation in policy-making. *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*, 1–10.
- Mariani, I., Mortati, M., Rizzo, F., & Deserti, A. (2024). An overview on e-participation. *Design Thinking as a Strategic Approach to E-Participation: From Current Barriers to Opportunities*, 31–56.
- Martins, P. H., Fernandes, P., Alves, J., Guerreiro, N. M., Rei, R., Alves, D. M., Pombal, J., Farajian, A., Faysse, M., Klimaszewski, M., et al. (2024). Eurollm: Multilingual language models for europe.
- Matthijs, M. (2020). Lessons and learnings from a decade of eu crises. *Journal of European Public Policy*, 27(8), 1127–1136.
- Nunamaker Jr, J. F., Briggs, R. O., Derrick, D. C., & Schwabe, G. (2015). The last research mile: Achieving both rigor and relevance in information systems research. *Journal of management information systems*, 32(3), 10–47.
- Panopoulou, E., Tambouris, E., & Tarabanis, K. (2009). Eparticipation initiatives: How is europe progressing. *European journal of ePractice*, 7(2009), 15–26.
- Peffer, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of management information systems*, 24(3), 45–77.
- Pennington, J., Socher, R., & Manning, C. (2014). GloVe: Global vectors for word representation. In A. Moschitti, B. Pang, & W. Daelemans (Eds.), *Proc. of the 2014 conf. on empirical methods in natural language processing (EMNLP)* (pp. 1532–1543).
- Radford, A. (2018). Improving language understanding by generative pre-training.
- Sæbø, Ø., Rose, J., & Flak, L. S. (2008). The shape of e-participation: Characterizing an emerging research area. *Government information quarterly*, 25(3), 400–428.
- Saldaña, J. (2021). *The coding manual for qualitative researchers*.
- Sanina, A., Balashov, A., Rubtcova, M., & Satinsky, D. M. (2017). The effectiveness of communication channels in government and business communication. *Information Polity*, 22(4), 251–266.
- Segura-Tinoco, A., Holgado-Sánchez, A., Cantador, I., Cortés-Cediel, M. E., & Rodríguez Bolívar, M. P. (2022). A conversational agent for argument-driven e-participation. *DG. O 2022: The 23rd Annual International Conference on Digital Government Research*, 191–205.
- Shipley, R., & Utz, S. (2012). Making it count: A review of the value and techniques for public consultation. *Journal of Planning Literature*, 27(1), 22–42.
- Sprenkamp, K., Dolata, M., Schwabe, G., & Zavolokina, L. (2025). Data-driven intelligence in crisis: The case of ukrainian refugee management. *Government Information Quarterly*, 42(1), 101978.
- Toots, M. (2019). Why e-participation systems fail: The case of estonia's osale. ee. *Government Information Quarterly*, 36(3), 546–559.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 6000–6010.
- Venkatesh, V., Morris, M. G., & Ackerman, P. L. (2000). A longitudinal field investigation of gender differences in individual technology adoption decision-making processes. *Organizational behavior and human decision processes*, 83(1), 33–60.
- Wright, S. (2009). The elephant in the room: Language issues in the european union. *European Journal of Language Policy*, 1(2), 93–119.