Investigating Trust in the incorporation of NLP applications in Digital Democracy Platforms

Nikolaos Giarelis [1], Nikos Karacapilidis [1], Georgios Kournetas [1] and Ilias Siachos [1]

[1] *IMIS Lab, University of Patras, 26504 Rio Patras, Greece*

Abstract

This study focuses on two basic natural language processing applications, namely clustering and summarization of the opinions expressed by participants in a digital democracy platform, aiming to investigate the extent that users trust them in terms of reliability, transparency, ethics, inclusiveness, trustworthiness, and accuracy. Results demonstrate a positive attitude in most cases, with the highest rank observed when referring to the reliability of opinion clustering. However, participants' confidence is less strong when evaluating the ethical implications of opinion clustering and the inclusivity of the opinion summarization process. Conducting non-parametric Kruskal-Wallis statistical tests, this study also reveals that English language proficiency plays a key role in shaping respondents' beliefs about the ethicality and accuracy of opinion clustering. Additionally, it highlights a positive correlation between familiarity with web applications and participants' perception of the accuracy of opinion clustering. To the best of our knowledge, this is the first attempt to gain such insights, which may reveal useful information about the utilization and deployment of these applications in digital democracy solutions.

Keywords

Digital Democracy, Natural Language Processing, Trust, Software Platforms.

1. Introduction

Digital democracy has been defined as *"the pursuit and the practice of democracy in whatever view using digital media in online and offline political communication"* [1]. Generally speaking, current digital democracy platforms are rudimental in the way they structure data, scarcely support evidence-based reasoning, lack features to enhance personal understanding, and fail to support effective deliberation and decision-making [2]. Solutions building on social media technologies are inapt to promote public discussion and cannot enable the realization of constructive, informative and rational dialogue. On the other hand, while participatory democracy solutions (such as Consul - see *https://joinup.ec.europa.eu/collection/joinup/solution/joinup-archive/release/100-beta*) have provided a much more constructive and inclusive environment to promote citizens engagement in collective decision making, they hardly support evidence-based thinking since deliberation data is neither presented nor collected in a way that makes it easy for people (or machines) to make sense of (or extract) the knowledge embedded in a democratic dialogue.

Such requirements can be only fulfilled by Natural Language Processing (NLP) technologies, which aim to make machines capable of understanding and reasoning with human language, and thus automatically processing the associated information. These technologies become increasingly intelligent, in that they lead to developments that may mimic or even outperform humans in diverse digital democracy acts, while also enabling new approaches to collaboration and knowledge co-creation. Functionalities that are usually facilitated and significantly augmented by NLP technologies include information extraction, text classification, sentiment analysis and semantic text matching. In the context of this study, we focus on two basic functionalities associated with digital democracy

Proceedings EGOV-CeDEM-ePart 2023, September 04–07, 2023, Budapest, Hungary
EMAIL: giarelis@ceid.upatras.gr (A.1); karacap@upatras.gr (A.2); kournetag@upnet.gr (A.3); ilias.siachos@upnet.gr (A.4)
ORCID: 0000-0003-2611-3129 (A.1); 0000-0002-6581-6831 (A.2); 0000-0001-8668-296X (A.3); 0000-0001-5489-3500 (A.4)

platforms, namely those of clustering and summarization of the opinions expressed by participants, aiming to investigate the extent that users trust these functionalities in terms of reliability, transparency, ethics, inclusiveness, trustworthiness, and accuracy when incorporated in a digital democracy platform. As argued in [3], trust in such emerging technologies has been considered to play a significant role in human-AI partnership, in that it does not only enable the adoption of the associated software platforms but also impacts users' behavior and interaction, enabling the long-term usage and the continuous improvement of these platforms. To the best of our knowledge, this is the first attempt to gain such insights, which may reveal useful information about the utilization and deployment of these NLP functionalities in digital democracy solutions addressing large-scale deliberation.

The remainder of this paper is structured as follows: Section 2 reports briefly on the two NLP functionalities elaborated in our study. Section 3 describes our methodology, research approach and data collection. Section 4 presents the re-search findings in the form of descriptive and inductive statistics, as well as through qualitative data analysis. Finally, Section 5 outlines concluding remarks, comments on the limitations of this study, and sketches future work directions.

2. NLP applications

2.1. Clustering

Clustering is the task of grouping a set of objects, in such a way that objects in the same cluster are more similar (based on various metrics) to each other than to those in other groups. Nowadays, clustering approaches are divided into several categories, based on the techniques employed, with the main categories being partition-based, hierarchical and density-based. Partition-based clustering approaches assign datapoints to clusters by extracting the center point of each cluster. K-Means [4] and K-Medoids [5] are the two most prominent. K-Means calculates the center of data points by an iterative procedure until some criteria for convergence are reached. K-Medoids follows a similar philosophy with K-Means, with the differentiating factor being that it can process discrete data. The data point closest to the center of data points, is rendered as the medoid of the corresponding cluster. The advantages of these approaches include relatively low time complexity and high computing efficiency. On the other hand, they do not efficiently handle non-convex data (i.e., relatively sensitive to the outliers). Additionally, the number of clusters must be predefined, which may impact the clustering result.

Hierarchical clustering approaches extract the hierarchical relationships among data. These approaches initially correspond each data point to an individual cluster. At each step, two clusters are merged into a new cluster, based on their proximity, until there is only one cluster left. BIRCH [6], ROCK [7] and Chameleon [8] are some typical approaches of this kind. Hierarchical approaches are preferred when handling datasets of arbitrary shape and types. The hierarchical relationships among clusters are more easily extracted, offering a relatively high scalability. However, these approaches have high computational complexity, and the number of clusters must be predefined.

The basic principle behind density-based approaches is that data points belonging to the same cluster must form a high-density region in the data space [9]. The typical ones include DBSCAN [10], OPTICS [11] and Mean-shift [12]. Density-based approaches have the advantages of high efficiency clustering while handling arbitrary-shaped data. Some drawbacks include low quality clustering results when the density of data space is not even, and increased memory requirements.

As suggested in the literature [13], word embeddings models can be used to infer a representative vector representation of the document (i.e., the document embedding). These embeddings can be then used by a clustering algorithm, to discover clusters of similar texts. Many word embedding models have been introduced in the literature after the introduction of the pioneering Word2Vec model [14]. The aim of these models is to introduce semantic information for textual terms, thus increasing the accuracy of various Natural Language Processing (NLP) tasks.

2.2. Summarization

Summarization is an NLP task, which deals with the creation of a short summary that represents the most important information from a single document or from multiple ones. Regarding this task, many approaches exist, which are classified into three major categories, namely the extractive, abstractive, or hybrid approaches. The extractive approaches split the input document into sentences, which are then ranked according to their importance and relevance to the overall document; these sentences are then concatenated to produce an output summary of the top-n most important sentences. Unlike the former approaches, abstractive ones utilize various techniques as to generate summaries comprising different text than the original document(s). Recent advancements in deep learning led to the development of abstractive approaches that create an internal representation of the input document(s), using pre-trained language models. By utilizing this representation, they are able to generate an abstractive summary. Finally, the hybrid approaches utilize the techniques employed by both the extractive and abstractive ones.

Extractive approaches can be further classified into various subcategories depending on their employed underlying techniques that they utilize to rank and extract the top-n sentences. These include: (i) statistical-based approaches ([15], [16]), which utilize statistical metrics such word or sentence frequency; (ii) graph-based approaches such as TextRank [17] and LexRank [18], which model the document into a graph of sentences, and then employ various graph-based measures (e.g., PageRank) for the sentence ranking and extraction step and (iii) semantic-based approaches such as the one presented in [19], which utilize the technique of Latent Semantic Analysis. This technique models the sentences and phrases into a co-occurrence matrix, and then ranks and extracts the top-n sentences.

Recent advancements in deep learning and transformers led to the creation of abstractive models based on the transformer architecture, including Unified Language Model [20], BART [21], and Text-To-Text Transfer Transformer model (T5) [22]. A detailed analysis of these models falls out of the scope of this paper.

3. Our Study

3.1. Conceptualizing and measuring trust

Trust is a concept of paramount importance when technological artifacts are used (or about to be used) by individuals and teams. It has been studied through different but complementary perspectives, including the traditional cognitive one where trust is being formed gradually over time and/or is based on categorization, disposition and third-party recommendations, the social / relational one where emphasis is given on social relations rather than on purely instrumental motives, and the emotional one where trust is not calculative and emotions reflect concerns whose underlying value is very strong, despite explicit belief to the contrary (for details, we refer to [3]).

To investigate and measure trust in the NLP applications described in Section 2, when these are incorporated in digital democracy platforms, we adopt in this study the *AttrakDiff Semantic Differential Scale*, which has been proven to be a valid and reliable instrument for assessing the attractiveness, pragmatic quality, hedonic quality, and overall appeal of interactive products [23]. This instrument has been extensively used in the field of human-computer interaction, making it suitable for evaluating trust in our case. Notable studies that have also built their research approach and questionnaires based on the same instrument include those of [24] and [25].

Based on the concepts of the abovementioned instrument, we developed a questionnaire (shown in the Appendix of this paper) aiming to quantify the users' responses with respect to the extent that they trust clustering and summarization functionalities when these are integrated in digital democracy platforms (on a 5-point Likert scale). Trust was investigated through six dimensions, namely reliability, transparency, ethics, inclusiveness, trustworthiness, and accuracy. In addition to the twelve basic closed-ended questions, participants were also asked to respond to two open-ended ones regarding whether 'they believe that the processes of clustering and summarizing opinions in a digital democracy platform may augment collaboration and knowledge co-creation among participants' and what are 'their main concerns when they are aware that the digital democracy platform they use employs AI technologies'.

The incorporation of the *AttrakDiff Semantic Differential Scale* offers several advantages for this research. First, it has been proven to have high reliability and validity, ensuring that our measurement of trust accurately reflects users' perceptions. Second, the scale is sensitive to different aspects of user experience, allowing us to distinguish between the *pragmatic* (e.g. usefulness and usability) and *hedonic* (e.g. emotional satisfaction) dimensions of trust. This distinction is particularly relevant in the context of AI-enhanced digital democracy platforms, as it enables us to capture the citizens' trust in using AI algorithms and the overall emotional appeal of the platform.

3.2. Research approach and data collection

The data reported in this study was collected through a questionnaire answered by 122 individuals. This questionnaire was developed through and hosted in an online survey platform (Google Forms); the corresponding link was disseminated via email and social media channels. We adopted the convenience sampling method due to the associated ease of access to the target participants, their availability at the time this study was carried out, and their willingness to participate in it. The target participants were within our acquaintanceship network, something that enabled a high response rate and assured the veracity of the responses collected.

Before answering the questionnaire, participants were briefly informed about the processes of clustering and summarization of opinions through AI algorithms and the potential of the incorporation of these processes in a digital democracy platform. This information was given by the researchers involved in this study, aiming to ensure that participants had a clear understanding of the study's concepts and objectives. It is also noted that participants were provided with a consent form explaining the purpose of the study, the voluntary nature of participation, and the anonymity of their responses. The data collection period lasted for three weeks, after which the responses were collated and analyzed.

3.3. Demographics

The majority of the participants were in the age group '25-34 years old' (36.2%). The sample displayed a balanced gender distribution (45.4% women and 53.1% men). Educational levels were relatively uniform across categories, except than a high representation of participants holding a M.Sc. degree (36.9%). In terms of occupational status, the majority of respondents were full-time employees (63.8%), followed by students (16.9%) and self-employed individuals (11.5%). The majority of participants were classified as intermediate (25.4%) or fluent (68.5%) English speakers. A considerable portion of the sample indicated themselves as being 'very comfortable' (68.5%) or 'somewhat comfortable' (21.5%) when using web applications. Finally, the frequency of e-government service usage among respondents was distributed as follows: 'monthly' (34.6%), 'rarely' (23.8%), 'weekly' (20.8%), and 'never' (6.9%).

4. Data Analysis

4.1. Descriptive Statistics

This section presents the descriptive statistics of our case study in the form of summary graphs. Specifically, it reports on the frequency distribution of the participants' responses regarding their trust in the processes of clustering and summarization of the opinions of individuals in the context under consideration. As mentioned in the previous section, we investigated trust through the dimensions of reliability, transparency, ethics, inclusiveness, trustworthiness, and accuracy.

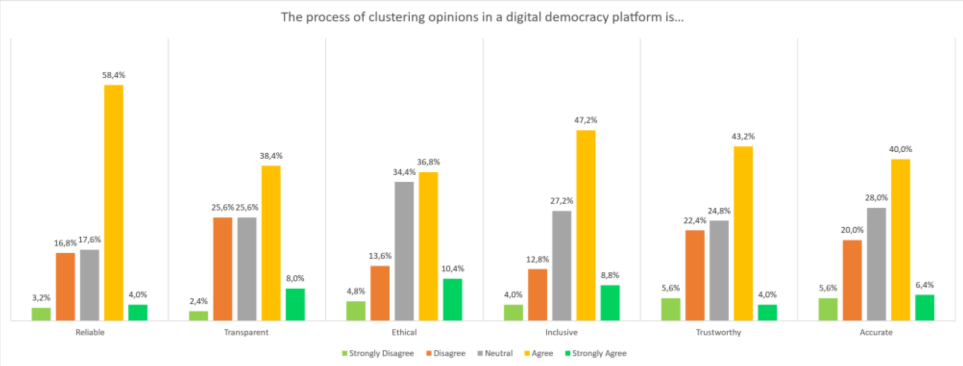

Figure 1: Frequency distributions of responses to questions concerning clustering of opinions.

Figure 1 demonstrates clearly that the majority of respondents accept and have trust in the process of clustering the opinions of citizens within a digital democracy platform. As revealed from the answers received, the response 'agree' consistently prevails over the alternative options, often demonstrating a significant majority. This observation is amplified in the answers referring to the perceived reliability of opinion clustering, where the prevalence of affirmative responses reaches its apex. Conversely, when evaluating the ethical implications of opinion clustering within digital democracy platforms, the assurance reported by participants is comparatively tempered; as illustrated in the corresponding bar chart, the percentage of individuals selecting 'agree' (36.8%) only slightly surpasses the percentage of individuals being 'neutral' (34.4%).

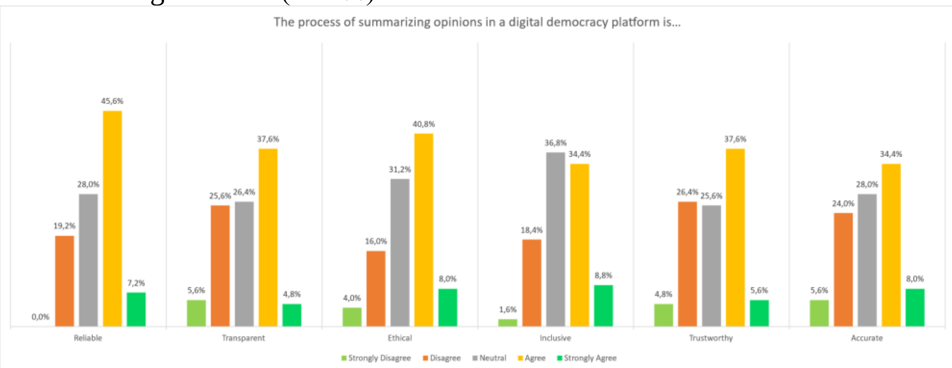

Figure 2: Frequency distributions of responses to questions concerning summarization of opinions.

Figure 2 illustrates a prevailing positive disposition among survey participants concerning the process of summarizing opinions in the context of a digital democracy platform. Nevertheless, the sentiment expressed for this process is more cautious when juxtaposed with the opinions registered for the process of opinion clustering; instances of 'disagree' responses emerge with greater frequency, while, interestingly enough, in the query referring to the inclusiveness of the opinion summarization process, the 'neutral' response (36.8%) marginally outweighs the affirmative 'agree' (34.4%).

4.2. Inductive Statistics

To explore potential associations between individuals' trust in the clustering and summarization applications incorporated in digital democracy platforms and their respective demographic attributes, the non-parametric Kruskal-Wallis H statistical tests were employed. This method was deemed appropriate due to the non-normal distribution of data gathered from the questionnaire, which is a consequence of the utilization of a five-point Likert scale. Our tests revealed the following three statistically significant correlations.

(i) 'Level of proficiency in English' vs 'The process of clustering opinions in a digital democracy platform is ethical'

Table 1
Kruskal-Wallis H test results

Test Statistics	
The process of clustering opinions in a digital democracy platform is ethical	
Kruskal-Wallis H	8.787
Df	3
Asymp. Sig	0.032

The p-value of this test (see Asymp. Sig.) is 0.032, which strongly indicates that there is a statistically significant difference between the levels of proficiency in English and the answers given to the statement 'The process of clustering opinions in a digital democracy platform is ethical'.

Table 2
Mean Ranks for each category of 'Level of Proficiency in English'

Ranks			
	What is your level of proficiency in English?	*N*	*Mean Ranks*
The process of clustering opinions in a digital democracy platform is ethical.	Beginner	4	64.50
	Intermediate	33	51.85
	Fluent	81	66.93
	Native Speaker	4	28.25
	Total	122	

Table 2 indicates a higher mean rank for fluent English speakers compared to intermediate English speakers. This observation suggests that intermediate speakers perceive the ethicality of opinion clustering within digital democracy platforms to be comparatively lower than the corresponding perception held by fluent speakers. Regarding the remaining categories of English proficiency, definitive conclusions cannot be drawn due to insufficient observations (comprising four instances for 'Beginners' and 'Native Speakers', respectively).

(ii) 'Level of proficiency in English' vs 'The process of clustering opinions in a digital democracy platform is accurate'.

Table 3
Kruskal-Wallis H test results

Test Statistics	
The process of clustering opinions in a digital democracy platform is accurate.	
Kruskal-Wallis H	9.429
Df	3
Asymp. Sig	0.024

The p-value of this test (see Asymp. Sig.) is 0.024, which strongly indicates that there is a statistically significant difference between the levels of proficiency in English and the answers given to the statement 'The process of clustering opinions in a digital democracy platform is accurate'.

Table 4
Mean Ranks for each category of 'Level of Proficiency in English'

Ranks

	What is your level of proficiency in English?	*N*	*Mean Ranks*
The process of clustering opinions in a digital democracy platform is ethical.	Beginner	4	89.50
	Intermediate	33	50.18
	Fluent	81	65.78
	Native Speaker	4	40.25
	Total	122	

Table 4 indicates a higher mean rank for fluent English speakers in comparison to their intermediate counterparts. This finding implies that intermediate speakers hold a relatively lower perception of the accuracy of opinion clustering within digital democracy platforms when contrasted with the perspective expressed by fluent speakers. The above remarks are in accordance with discussions claiming that proficiency in English language is nowadays an increasing prerequisite to enter technology driven spaces and markets, and that a lack thereof is deterring people to take advantage of the Web and its emerging applications.

(iii) 'How comfortable are you when using web applications (e.g., online forms, document editors, and social media)' vs 'The process of summarizing opinions in a digital democracy platform is transparent'

Table 5
Kruskal-Wallis H test results

Test Statistics	
The process of summarizing opinions in a digital democracy platform is transparent.	
Kruskal-Wallis H	11.382
Df	4
Asymp. Sig	0.023

The p-value of this test (see Asymp. Sig.) is 0.023, which strongly indicates that there is a statistically significant difference between the levels of how comfortable participants are when using web applications (e.g., online forms, document editors, and social media), and the answers given to the statement 'The process of summarizing opinions in a digital democracy platform is transparent'.

Table 6
Mean Ranks for each category of 'Being comfortable with web applications'

Ranks			
	How comfortable are you when using web applications (e.g., online forms, document editors, and social media)?	*N*	*Mean Ranks*
The process of summarizing opinions in a digital democracy platform is transparent.	Very uncomfortable	2	57.50
	Somewhat uncomfortable	1	22.00
	Neutral	10	54.80
	Somewhat comfortable	25	44.32
	Very comfortable	84	67.98
	Total	122	

Table 6 presents a higher mean rank for individuals who exhibit a high degree of comfort when utilizing web applications as opposed to those who are only somewhat comfortable. This finding implies that familiarity with web applications may serve as a contributing factor in fostering a more favorable perception with respect to the transparency of opinion clustering within a digital democracy platform.

4.3. Analysis of open-ended questions

As mentioned in Section 3.1, our questionnaire included two open-ended questions. The analysis of related responses revealed several key themes and sub-themes, which are summarized below (direct quotes are given in italics).

Open-ended question #1: *Do you believe that the processes of clustering and summarizing opinions in a digital democracy platform may augment collaboration and knowledge co-creation among participants? Please justify your answer.*

Positive impact

- Facilitate the identification of common themes and areas of agreement *("… clustering groups similar opinions to identify common themes and issues …", "clustering and summarizing can help to identify points of agreement/disagreement, facilitating more focused and productive discussions", "AI could create easy to consume text that will facilitate achieving common understanding on complex topics").*
- Enhance engagement and understanding of diverse perspectives *("… these techniques make it easier to engage in informed discussions and decisions", "these processes make the digital democracy platform more inclusive, thus more engaging", "summarizing opinions can make it easier for participants to understand and engage with each other's perspectives, leading to a more productive discussion").*
- Prioritize issues for focused discussion *("… clustering and summarizing opinions can help prioritize issues, making it easier to focus on the most important topics").*

Concerns and limitations

- Potential exclusion of important details or nuances ("… clustering/summarizing may exclude important details, on the subject, where the participants do not agree on").
- Misrepresentation of opinions ("I think that summarizing is misleading, as key points might be similar, but hide different ethics and actions", "… some people with conflicting opinions may disagree about the wording of the summary, or the categorization").

Open-ended question #2: *What are your main concerns when you are aware that the digital democracy platform you use employs AI technologies?*

Bias and fairness

- Biased algorithms *("AI technologies can be biased if they are trained on biased data or if their algorithms have built-in biases", "That they are designed by technical engineers without experts on demography, anthropologists and social scientists").*
- Handling of less popular opinions *("Less popular opinions may be fading into the background", "Some important ideas/opinions may not be included in a category or the summary, so they won't be heard by everyone and that could affect the decisions taken", "If the AI algorithms are not designed and tested properly, they may unintentionally discriminate against certain groups or unfairly amplify certain opinions over others").*

Transparency and accountability

- Lack of transparency in AI decision-making and accountability of AI-generated outcomes *("The reliability and transparency of the underlying algorithms", "... another concern is the transparency and accountability of the AI systems used; participants may be worried about how their data is being collected, stored, and used by the AI system, and may want to ensure that the system is transparent about its processes and accountable for its decisions", "One of the main concerns about using AI in digital democracy platforms is the lack of transparency; participants may not understand how AI is being used to make decisions, and may not be able to access or understand the data used to train the AI algorithms").*

Data privacy and protection

- Concerns about data collection, storage and use *("Main concerns are privacy, data protection, transparency and accountability", "... I'm suspicious that my personal data are not secure", "... users might be concerned about how their data is being used and whether it is being kept secure").*
- Unauthorized data sharing *("My main concern is the possibility of slightly altering my opinion as well as the possibility of transferring personal data without my knowledge", "I'm not concerned about technology, just about its usage and how we guarantee anonymity")*.

Reliability and accuracy

- AI misunderstanding or misinterpreting opinions *("I am concerned about the results and about the possibility of my answers being partially misunderstood by the AI technologies", "Depending on the way it is implemented it can greatly affect the outcome, for better or for worse", "AI is not able to process phrases and metaphors in the same way as a person does, so it could misunderstand the point of someone's words").*
- Trust in AI technologies *("even if the summarization or clustering is 99% percent accurate, which is very wishful thinking, what do we do with the rest 1%?", "... machines are quite useful and important but they will never be (like) humans").*

Inclusiveness and representation

- Ensuring participation from all demographic groups *("The main concern would be that not all opinions especially those of people who don't have access to digital means are taken into consideration")*.

5. Discussion and Conclusions

This paper reports on the results of a survey aiming to investigate trust in the incorporation of two NLP applications, namely opinion clustering and opinion summarization, in digital democracy platforms. Trust has been investigated through six dimensions, namely reliability, transparency, ethics, inclusiveness, trustworthiness, and accuracy. Results indicate that the response 'agree' prevails in most cases, with the highest difference observed when referring to the reliability of opinion clustering. However, participants' confidence is less strong when evaluating the ethical implications of opinion clustering and the inclusivity of the opinion summarization process. Conducting non-parametric Kruskal-Wallis statistical tests, this study has also revealed that English language proficiency plays a key role in shaping respondents' beliefs about the ethicality and accuracy of opinion clustering. Additionally, the research highlights a positive correlation between familiarity with web applications and participants' perception of the accuracy of opinion clustering. Finally, qualitative data analysis on responses to two open-ended questions has formulated a series of themes and sub-themes to enable a better understanding of the main issue investigated in this study.

Our findings reveal that clustering and summarizing opinions in a digital democracy platform may have both positive and negative effects. On one hand, these processes can facilitate the identification of common themes and areas of agreement, leading to more focused and productive discussions. They can also enhance engagement and understanding of diverse perspectives, making it easier for participants to appreciate each other's views and contribute to the issue under consideration. On the other hand, there are concerns regarding the potential exclusion of important details, suppression of individuality, and misrepresentation of opinions. Ensuring transparency and impartiality in these processes can help build trust among participants and increase the likelihood of meaningful collaboration and knowledge co-creation.

Moreover, it was revealed that designers of digital democracy platforms should consider the incorporation of mechanisms that allow participants to elaborate both aggregated and individual opinions; this can address concerns related to the suppression of individuality and the potential exclusion of important details or nuances. Inclusivity and representation of diverse groups should also receive much attention to ensure that the clustering and summarization of opinions do not marginalize or exclude certain voices.

Our study also indicated that users of digital democracy platforms have various concerns about the use of AI technologies, which are related to issues including bias, fairness, transparency, accountability, privacy, data security, reliability, and accuracy. Addressing these concerns is crucial in fostering trust in AI-driven digital democracy platforms and ensuring that they effectively support collaboration and knowledge co-creation among participants. To mitigate concerns related to bias and fairness, designers of digital democracy technologies should strive to use a diversity of training data and involve experts from fields such as demography, anthropology, and political and social science. Increasing transparency and accountability in AI decision-making can be accomplished by making algorithms better understandable and explainable to users [26].

Privacy and data security concerns can be addressed by implementing robust data protection mechanisms, transparent data handling practices, and ensuring that users have control over their personal information. Users should be informed about how their data is collected, stored, and used, as well as their rights regarding data access, correction, and deletion. To address concerns related to reliability and accuracy, AI technologies should be rigorously tested and validated to ensure that they can effectively process and analyze user-generated opinions. This includes evaluating AI algorithms for their ability to understand and interpret complex language, metaphors, and diverse perspectives.

Our research has some limitations that should be considered when interpreting the results. One of the primary limitations is the lack of observations within specific subgroups of demographic characteristics, which may have constrained the number of statistically significant correlations between participants' trust and their demographic attributes. First, regarding the level of English proficiency, our sample included only four beginner speakers and four native speakers, potentially limiting our ability to draw conclusions about the overall impact of language proficiency on the study outcomes. Second, in terms of current employment status, our sample had only six part-time employed individuals and three participants who selected the 'other' option, which may have restricted our understanding of the relationship between employment status and trust in the context under consideration. Third, our sample exhibited a low variability in participants' comfort levels with web applications, as only three individuals responded 'neutral' and one 'very uncomfortable'; this could affect the generalization of our findings, as it may not accurately represent the full spectrum of user experiences.

The above limitations, which are mainly due to the disadvantages of the convenience sampling method adopted in our case study, highlight the need for future research to include larger and more diverse samples, ensuring that various demographic subgroups are well-represented. Such studies would allow for a more comprehensive understanding of the relationship between demographic attributes and trust in digital democracy platforms and incorporated technologies, thereby contributing to the development of more inclusive, transparent and effective solutions.

6. Acknowledgements

The work presented in this paper is supported by the inPOINT project (*https://inpoint-project.eu/*) which is co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH—CREATE—INNOVATE (Project id: T2EDK-04389).

7. References

[1] van Dijk, J.A. (2012). Digital Democracy: Vision and Reality. In Public Administration in the Information Age: Revisited, pp. 49-62, IOS Press.

[2] Convertino, G., Westerski, A., De Liddo, A. and Díaz, P. (2015). Large-Scale Ideation & Deliberation: Tools and Studies in Organizations. Journal of Social Media for Organizations, 2(1).

[3] Gkinko, L., & Elbanna, A. (2023). Designing trust: The formation of employees' trust in conversational AI in the digital workplace. Journal of Business Research, 158, 113707, doi: 10.1016/j.jbusres.2023.113707.

[4] MacQueen, J. (1967). Classification and analysis of multivariate observations. 5th Berkeley Symp. Math. Statist. Probability, 281–297.

[5] Park, H.-S., & Jun, C.-H. (2009). A simple and fast algorithm for K-medoids clustering. Expert Systems with Applications, 36(2), 3336–3341. https://doi.org/10.1016/j.eswa.2008.01.039

[6] Zhang, T., Ramakrishnan, R., & Livny, M. (1996). BIRCH: an efficient data clustering method for very large databases. ACM Sigmod Record, 25(2), 103–114.

[7] Guha, S., Rastogi, R., & Shim, K. (2000). ROCK: A robust clustering algorithm for categorical attributes. Information Systems, 25(5), 345–366.

[8] Karypis, G., Han, E.-H., & Kumar, V. (1999). Chameleon: Hierarchical clustering using dynamic modeling. Computer, 32(8), 68–75.

[9] Kriegel, H.-P., Kröger, P., Sander, J., & Zimek, A. (2011). Density-based clustering. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 1(3), 231–240.

[10] Ester, M., Kriegel, H.-P., Sander, J., Xu, X., & others. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. Kdd, 96, 226–231.

[11] Ankerst, M., Breunig, M. M., Kriegel, H.-P., & Sander, J. (1999). OPTICS: Ordering points to identify the clustering structure. ACM Sigmod Record, 28(2), 49–60.

[12] Comaniciu, D., & Meer, P. (2002). Mean shift: A robust approach toward feature space analysis. IEEE Transactions on Pattern Analysis and Machine Intelligence, 24(5), 603–619.

[13] Mehta, V., Bawa, S., & Singh, J. (2021). WEClustering: Word embeddings based text clustering technique for large datasets. Complex & Intelligent Systems, 7, 3211–3224.

[14] Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. ArXiv Preprint ArXiv:1301.3781.

[15] Luhn, H. P. (1958). The automatic creation of literature abstracts. IBM Journal of Research and Development, 2(2), 159–165.

[16] Ko, Y., & Seo, J. (2008). An effective sentence-extraction technique using contextual information and statistical approaches for text summarization. Pattern Recognition Letters, 29(9), 1366–1371.

[17] Mihalcea, R., & Tarau, P. (2004). TextRank: Bringing order into text. Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing, 404–411.

[18] Erkan, G., & Radev, D. R. (2004). LexRank: Graph-based lexical centrality as salience in text summarization. Journal of Artificial Intelligence Research, 22, 457–479.

[19] Steinberger, J., Jezek, K. (2004). Using latent semantic analysis in text summarization and summary evaluation. Proc. ISIM, 4(93–100), 8.

[20] Dong, L., Yang, N., Wang, W., Wei, F., Liu, X., Wang, Y., Gao, J., Zhou, M., & Hon, H.-W. (2019). Unified language model pre-training for natural language understanding and generation. Advances in Neural Information Processing Systems, 32.

[21] Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., & Zettlemoyer, L. (2019). Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. ArXiv Preprint ArXiv:1910.13461

[22] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., & Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. The Journal of Machine Learning Research, 21(1), 5485–5551.

[23] Hassenzahl, M., Burmester, M., & Koller, F. (2003). AttrakDiff: A questionnaire to measure perceived hedonic and pragmatic quality. Mensch & Computer, 57, 187-196.

[24] Sousa, S., Martins, P., & Cravino, J. (2021). Measuring Trust in Technology: A Survey Tool to Assess Users' Trust Experiences

[25] Takahashi, L. & Nebe, K. (2019). Observed Differences Between Lab and Online Tests Using the AttrakDiff Semantic Differential Scale. 14. 65-75.

[26] Tiddi, I. & Schlobach, S. (2022). Knowledge graphs as tools for explainable machine learning: A survey. Artificial Intelligence, vol. 302, 103627, doi: 10.1016/j.artint.2021.103627.

Appendix – The questionnaire used in our study

Investigating Trust in Digital Democracy Platforms

Two basic functionalities integrated in Digital Democracy platforms are those of *clustering* and *summarization* of the opinions expressed by participants.

- ***Clustering*** is the task of grouping a set of objects in a way that objects in the same cluster are more similar to each other than those in other groups.
- ***Summarization*** is a Natural Language Processing task that deals with the creation of a short summary that represents the most important information from a single or multiple documents.

Part A.
The aim of this questionnaire is to assess the degree that citizens trust these functionalities when integrated in digital democracy platforms. Please tell us to what extent do you agree with the following statements:

Reliability
The process of ***clustering*** *opinions in a digital democracy platform is reliable.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

The process of ***summarizing*** *opinions in a digital democracy platform is reliable.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

Transparency
The process of ***clustering*** *opinions in a digital democracy platform is transparent.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

The process of ***summarizing*** *opinions in a digital democracy platform is transparent.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

Ethics
The process of ***clustering*** *opinions in a digital democracy platform is ethical.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

The process of ***summarizing*** *opinions in a digital democracy platform is ethical.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

Inclusiveness
The process of ***clustering*** *opinions in a digital democracy platform is inclusive.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

The process of ***summarizing*** *opinions in a digital democracy platform is inclusive.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

Trustworthiness
The process of ***clustering*** *opinions in a digital democracy platform is trustworthy.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

The process of ***summarizing*** *opinions in a digital democracy platform is trustworthy.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

Accuracy
The process of ***clustering*** *opinions in a digital democracy platform is accurate.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

The process of ***summarizing*** *opinions in a digital democracy platform is accurate.*
○ Strongly Disagree
○ Disagree
○ Neutral
○ Agree
○ Strongly Agree

Part B. Open-ended Questions
Do you believe that the processes of clustering and summarizing opinions in a digital democracy platform may augment collaboration and knowledge co-creation among participants? Please justify your answer (100 words max).

Your answer goes here …

What are your main concerns when you are aware that the digital democracy platform you use employs AI technologies? (100 words max)

Your answer goes here …

Part C. Demographics
What is your age?
○ 18-24
○ 25-34
○ 35-44
○ 45-54
○ 55-64
○ 65 or older

What is your gender?
○ Male
○ Female
○ Prefer not to say
○ Other

What is your highest level of completed education?
○ High school diploma or equivalent
○ College diploma or equivalent
○ Bachelor's degree
○ Master's degree
○ Doctoral degree

What is your current employment status?
○ Employed full-time
○ Employed part-time
○ Self-employed
○ Unemployed
○ Student
○ Retired
○ Other

What is your level of proficiency in English?
○ Native speaker
○ Fluent
○ Intermediate
○ Beginner

How comfortable are you when using web applications (e.g. online forms, document editors, and social media)?
○ Very Comfortable
○ Somewhat Comfortable
○ Neutral
○ Somewhat Uncomfortable
○ Very uncomfortable

How frequently do you use e-government services?
○ Multiple times per week
○ Weekly
○ Monthly
○ Rarely
○ Never