An AI-Enhanced Solution for Large-Scale Deliberation Mapping and Explainable Reasoning

Nikos Karacapilidis[1(✉)], Dimitris Tsakalidis[2], and George Domalis[2]

[1] University of Patras, 26504 Rio Patras, Greece
karacap@upatras.gr
[2] Novelcore, 26500 Patras, Greece

Abstract. This work aims to respond to the profound lack of dialogue between citizenship and policy making institutions by proposing a novel solution that enables the transition to inclusive, transparent, accountable and trustworthy deliberation practices. The proposed solution builds on cutting-edge AI tools and technologies to develop a sustainable digital platform, and bridges theories from the fields of argumentation and digital democracy. It may transform scattered islands of emerging knowledge and practices, as well as fragmented discussion threads, into an integrated and coherent dialogue, and provides mechanisms for expanding this dialogue and converting it into tangible actions. Much attention is paid to issues related to knowledge extraction, knowledge graph-based representation of large-scale deliberation, argument mining, aggregation and visualization, as well as to explanation and awareness services about the evolution and outcome of a deliberation.

Keywords: Digital Governance · Digital Democracy · Citizen-centric e-Governance · Artificial Intelligence in Government · conceptual framework

1 Introduction

Citizens worldwide are increasingly worried about the socio-political crises and conflicts emerging around the world and seek new ways to exercise their social responsibility. In such settings, they ask to be engaged in democratic and inclusive discussions about how these crises and conflicts can be prevented, mitigated and even resolved in new ways, such as through mobilization of grassroot collective actions. However, current methods to engage citizens in fruitful deliberations, gain trust, and harness their commitment to act are problematic; without the practical means to engage them into policy making, a meaningful and effective joining-up of bottom-up citizen-led initiatives and top-down policy making has not yet been realized.

At the same time, political science research highlights the need of balance as well as communication, interaction and exchange of knowledge between 'democracy' (democratic institutions, consultations with citizens) and 'technocracy' (specialized knowledge of experts and policy makers), as they are complementary, each of them needing inputs

M. Papadaki et al. (Eds.): EMCIS 2022, LNBIP 464, pp. 305–316, 2023.
https://doi.org/10.1007/978-3-031-30694-5_23

from the other, while both making significant but different contributions to the design of effective and socially acceptable public policies (Androutsopoulou *et al.*, 2018). In particular, participants in democratic processes need extensive knowledge and expertise on the social problems they are dealing with. On the other hand, experts dealing with important social problems tend to ignore important aspects of public policies, such as their impact on employment, social inequalities, and quality of life. To reduce these negative tendencies, experts need inputs from democratic political processes concerning the values of citizens and other stakeholder groups, as well as their diverse perspectives and approaches.

The approach proposed in this paper aspires to address the above issues by unleashing the power of democratic and participatory processes towards the aggregation of ideas and the co-creation of efficient and effective solutions to multi-dimensional societal problems. It aims to exploit and meaningfully integrate internal and external data, by considering all the operational stakeholders as key co-creators of value information-knowledge-action chains, thus sustaining and inspiring better-informed collaboration towards innovative actions. The proposed approach creates a novel deliberation solution to transform scattered islands of emerging knowledge and practices, and fragmented discussion threads into an integrated and coherent dialogue and provides mechanisms for expanding this dialogue and converting it into tangible actions.

The key contributions of the proposed solution are: (i) novel knowledge extraction algorithms to yield factual and affective knowledge; (ii) a knowledge graph-based representation of large-scale deliberation enriched with state-of-the-art natural language understanding, argument mining, aggregation and visualization mechanisms to turn unstructured user-generated content into knowledge and actions, and (iii) explanation and awareness services about the evolution and outcome of deliberation to enable better informed collaboration, augment sense making and increase transparency of the overall process.

The remainder of this paper is structured as follows: Sect. 2 reports on background issues concerning large-scale online deliberation methods and tools, argumentation and social knowledge mining, group decision making in large-scale deliberations, and knowledge graphs for deliberation mapping. Section 3 describes the proposed solution for the facilitation and enhancement of large-scale deliberations, along with its potential and expected impact. Finally, Sect. 4 outlines concluding remarks and future work directions.

2 Background Issues

2.1 Large-Scale Online Deliberation Methods and Tools

Current deliberation platforms are rudimental in the way they structure data, scarcely support evidence-based reasoning, lack features to enhance personal understanding and situational awareness, and hardly support effective deliberation and decision-making. If we look at *social media solutions*, a wide research literature demonstrates how online dialogue on these platforms is prone to toxic behaviors such as biased and un-supported information, rumors, misinformation, hate speech and echo chambers effects. These technologies are therefore inapt to promote public discussion and fail to enable the realization of constructive attitude, informative and rational dialogue, civility and equality.

On the other hand, *participatory democracy solutions* such as *Consul, Democracy OS, Loomio* and *Decidim* have demonstrated large adoption in supporting a variety of democratic processes, such as solicitation of ideas on public issues, community voting, and participatory budgeting. While this second category of solutions is able to promote active change in specific policy making contexts, and provides a much more constructive and inclusive environment to promote citizens engagement in collective decision making, it shares some of the weaknesses of social media; it provides simple discussion features and hardly supports evidence-based thinking since discussion data is neither presented nor collected in a way that makes it easy for people or machines to make sense of the knowledge embedded in the dialogue. Moreover, when the discussion scales, it is hard for participants to grasp the status and progress of the deliberation.

To address these shortcomings, *issue-centric solutions* such as *Kialo, Deliberatorium, Cohere, DebateGraph* and *The Evidence Hub* enable people to interact by building deliberation maps that are made up of interlinked questions, answers and arguments. Such tools help communities be much more systematic and complete in their deliberations about complex topics, enhance evidence-based dialogue, build common ground, support the development of shared understanding of complex problems and improve the quality of online argumentation (De Liddo *et al.*, 2012). However, the uptake and impact of these solutions is hindered by a lack of usable and intuitive interfaces for online dialogue.

2.2 Argumentation and Social Knowledge Mining

Argumentation mining lies between natural language processing, argumentation theory and information retrieval, aiming to automatically detect the arguments expressed in a deliberation process, their individual or local structure and the interactions between them. The main goal of argumentation mining is to automatically extract arguments from generic textual corpora, in order to provide structured data for computational models of argument and reasoning engines. Recent advances in Computational Argumentation and Natural Language Processing (NLP) enable the development of novel methods that may capture arguments and inform stakeholders about the evolution of a deliberation through contextualization, representation and aggregation of argumentation in diverse contexts (Cabrio and Villata, 2018).

Argumentation mining systems developed so far adopt a pipeline architecture through which they process unstructured textual documents and produce as output a structured document, where the detected arguments and their relations are annotated so as to form an argument graph (Lippi and Torroni, 2016). Such a pipeline consists of three basic subtasks, namely argumentative sentence detection, argument component boundary detection, and argument structure prediction. There are many similarities between these subtasks that are typically addressed by prominent Machine Learning (ML) and NLP techniques. Approaches to argumentation mining adopt either a *discourse-level perspective*, aiming to analyze local argumentation structures, or an *information-seeking perspective,* aiming to detect arguments that are relevant to a predefined topic. Consequently, such approaches call for a subsequent argumentation aggregation step, which can aggregate similar arguments for the same topic.

As far as argument aggregation is concerned, a variety of models have been already proposed based on *argument-wise* and *framework-wise* methods (Bodanza *et al.*, 2017).

In the former, individually supported arguments are aggregated by a voting mechanism, while in the latter the aggregation comes from merging the individually supported criteria or different argumentation frameworks through a collectively decided method, depending on the specific argumentation context under consideration. The framework-wise approach is considered more efficient in the context of deliberative democracy, while the argument-wise approach could be the most efficient one in the context of a debate among experts. In a similar research line, contextualized word embeddings that classify and cluster topic-dependent arguments have been recently proposed in the literature. Two of the most popular approaches are *Embeddings from Language Models* (Peters *et al.*, 2018) and *Bidirectional Encoder Representations from Transformers (BERT)* (Devlin *et al.*, 2018). Contrary to traditional word embeddings, these approaches calculate the embeddings for a sentence dynamically, by considering the context of a target word. This generates word representations that better match the specific sense of the word in a sentence.

2.3 Group Decision Making in Large-Scale Deliberations

New technological paradigms such as social networks, e-participation, e-democracy and e-marketplaces enable the participation of big numbers of stakeholders in the decision-making process. Consequently, these paradigms make it possible to obtain more and more subjective and objective data. At the same time, the group decision making process is characterized by the following: (i) the scale of groups participating in the process has become much larger than before, varying from dozens to thousands; (ii) people involved in the process come from different organizations and in most cases have different backgrounds, interests and constraints; (iii) individuals can express opinions at different times or places, while the final solution is no longer attributable to a single decision maker, but rather to a large-scale group making decisions jointly (Karacapilidis, 2014).

State-of-the-art approaches attempt to address the following major challenges (Tang and Liao, 2019): *(i) reduction of the decision makers' dimension:* clustering analysis is the most widely used method so far, aiming to reduce the complexity and cost of the associated problems, as well as to identify common opinion patterns (e.g. clusters with similar opinions and a spokesman who represents each cluster); *(ii) weighting and aggregating decision information:* the development of a reasonable method that considers the diverse characteristics of individuals and subgroups to determine weights is very crucial, making simple aggregation strategies such as arithmetic average or weighted average not appropriate; *(iii) management of participants' behavior:* existing studies often adopt a social network analysis perspective to investigate the consensus reaching process and the associated detection and elimination of conflicts among decision makers (Liu *et al.*, 2019); *(iv) cost management:* diverse consensus models with minimum cost have been already proposed to address this challenge, which is associated to the feedback mechanism of the whole process; *(v) knowledge distribution and information increase:* this concerns the diverse social relationships that may exist among decision makers as well as the consideration of additional information such as trust and reputation of them.

Most prominent tools and technologies build on concepts and techniques from Artificial Intelligence and Operational Research to enable a sophisticated data analysis, while also discovering patterns of data and inferring data content relationships and rules from

them (Karacapilidis *et al.*, 2014). Such tools and technologies certainly facilitate diverse aspects of decision making. Although there exist certain limitations in their suitability, they may aid users to make better and faster decisions. However, there is still room for further developing the conceptual, methodological and application-oriented aspects of the problem. One critical point that is still missing is a holistic perspective on the issue of large-scale group decision making. This originates out of the growing need to develop applications by following *a more human-centric (not problem-centric) view*, in order to appropriately address the requirements of contemporary knowledge-intensive settings. Such requirements stem from the fact that decision making has also to be considered as *a social process* that principally involves human interaction. The structuring and management of this interaction requires the appropriate technological support.

2.4 Knowledge Graphs for Deliberation Mapping

Knowledge Graphs (KGs) facilitate the storage and representation of knowledge in a direct and expandable way (Wang *et al.*, 2014). Recent KG-based approaches can represent knowledge extracted from either structured (e.g., tabular and matrix data) or unstructured data (e.g., media and textual data). The advantage of KGs against the classical knowledge bases is that they generally perform better in data-intensive environments, since they allow for: (i) easier data schema expansion and alternation, (ii) better knowledge extraction and representation, (iii) masking of the underlying data complexity, (iv) integration of knowledge from external sources (e.g., Wikipedia), and (v) exploitation of graph algorithms. In addition, KGs can effectively represent both information related to the relations between entities and information that concerns each individual entity (Lin *et al.*, 2017).

With respect to the representation of the knowledge existing in such graphs, *KG embeddings* have been recently adopted. KG embeddings provide low-dimensional dense vectors, which incorporate important information related to the entities (i.e. nodes) and relations (i.e. edges) of a KG. Most important, KG embeddings assist traditional ML models in performing a list of tasks more accurately. Such tasks may concern entity classification, inference of relations, network analysis and prediction of links between the entities of a KG.

KGs have already been applied to several practical domains including question-answering, language models, entity matching, chatbots, dialog systems, recommendation engines, fraud detection, and prediction of future research collaborations (Wang *et al.*, 2017). Furthermore, they excel in real world applications, where complex data from multiple sources can only be processed together, aiming to gain important insights. As far as the implementation and the utilization of a KG are concerned, several well-tried and mature programming libraries and tools exist in the literature. For instance, the *Neo4j* database provides the user with already implemented graph and ML algorithms, thus enabling the construction of robust and production-ready KGs. Graph databases can be seamlessly used along with widely used ML frameworks such as *TensorFlow, PyTorch* and *scikit-learn* to build meaningful ML models and pipelines.

3 The Proposed Solution

3.1 Research Methodology

The development of the proposed solution follows the *Design Science* paradigm, which seeks to create innovations that define the ideas, practices, technical capabilities, and products through which the analysis, design, implementation, and use of Information Systems can be effectively and efficiently accomplished (Hevner *et al.*, 2004). This paradigm has been extensively adopted in the development of Information Systems in order to address what are considered to be *wicked problems*, i.e., problems characterized by unstable requirements and constraints based on ill-defined contexts, complex interactions among issues of the problem, inherent flexibility to change design processes and artifacts, and a critical dependence upon human cognitive and social abilities to produce effective solutions. At the same time, our approach is in line with the *Action Research* paradigm, which aims to contribute both to the practical concerns of people in a problematic situation and to the goals of social science by joint collaboration within a mutually acceptable ethical framework (Rapoport, 1970). As such, it concerns the improvement of practices and strategies in the particular cognitively complex environment under consideration, as well as the acquisition of additional knowledge to improve the way stakeholders address issues and solve problems (Checkland and Holwell, 1998).

3.2 Conceptual Architecture

The proposed platform offers a holistic and modular solution that securely hosts and effectively supports large-scale deliberation processes. All the individual modules are designed to be built on top of a cloud service system, configured to be aligned with the needs of all types of stakeholders. This modular approach constitutes the backbone of our solution, which is capable of thoroughly addressing the complexity of deliberative processes, while also enhancing trust, transparency and legitimacy of policy making.

Guided by advancements in (Explainable) AI, ML, NLP, Graph Theory and Argumentation, our human-centric approach will produce an efficient and scalable platform that can support deliberation processes of different models and at all levels, from local to global. The proposed technical solution, whose *three-layer architecture* is illustrated in Fig. 1, ensures the seamless integration (at both a conceptual and a technical level) and interoperability of diverse components and services. It enables a *synergy of human and machine reasoning* towards facilitating and augmenting the participation and deliberation of diverse types of stakeholders in structured discursive interactions. In addition, it exploits *rich semantics at machine level* to enable the meaningful incorporation and orchestration of interoperable services, aiming to reduce the inherent data-intensiveness of the context under consideration. In particular, the proposed solution seamlessly integrates:

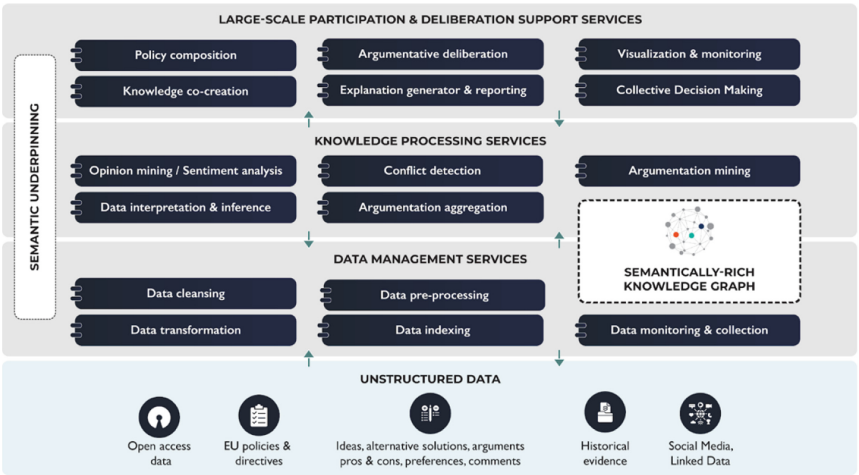

Fig. 1. The architecture of the proposed large-scale deliberation platform.

- ***Participation and deliberation support services*** (Fig. 1, top layer) that (i) support incremental formalization of argumentative deliberation, i.e., a stepwise and controlled evolution from a mere collection of individual ideas and resources to the production of highly contextualized and interrelated knowledge artifacts; (ii) augment sense-making through advanced visualization and monitoring dashboards that offer an informative and user-friendly overview of a deliberation process in terms of participants' engagement and contributed knowledge, while also providing insights about the structure, evolution and dynamics of a deliberation process; (iii) provide advanced knowledge exchange and co-creation functionalities by offering a deliberation environment that supports interpretation of diverse knowledge items and their interrelationships; (iv) are geared towards facilitating collective decision making and consensus building through innovative virtual workspaces that enable participants rank alternative solutions; (v) aid policy makers discover and elucidate key points relevant to the deliberation and accordingly compose evidence-informed policies and practices, and (vi) aid citizens and other stakeholders to get explanations and reporting about the inferential process of the underlying AI algorithms and decision making mechanisms, thus promoting trust in the deliberation outputs.
- ***Knowledge processing services*** (Fig. 1, middle layer) that enable a sophisticated analysis on the associated textual content of deliberative processes. Building on prominent AI techniques, this set of services facilitates the comprehension of the structure and meaning of argumentative deliberation by breaking down the input received into a machine-readable format. By understanding a set of linguistic and structural cues, our solution enables a precise interpretation of the corresponding texts and their transformation into actionable, measurable and easily accessible knowledge, thus augmenting the quality of human-computer interaction. Items identified populate and are meaningfully linked with the platform's knowledge graph to support a sophisticated representation of deliberation entities and their dynamics. This set of services builds on and extends state-of-the-art ML frameworks and neural architectures for

NLP to harness the complexity and address the uncertainties stemming from the associated data and knowledge. It includes services that (i) deploy novel argumentation mining mechanisms paying particular attention to argumentation quality; (ii) trigger context-dedicated argumentation aggregation algorithms to serve the clustering of similar deliberation items, as well as the consolidation and homogenization of participants' input by leveraging prominent text summarisation techniques; (iii) enable opinion mining techniques to identify, extract and process opinions from text by using a lexical approach in combination with prominent ML/NLP techniques, and (iv) handle conflict detection by analyzing the results of argument mining and offering recommendations for conflict resolution.

- ***Data management services*** (Fig. 1, bottom layer) that enable the purposeful monitoring and collection and efficient pre-processing of tractable information that exists in our solution's data sources. Transformation of different kinds of textual documents into a canonical form, structuring of these documents from layout information (e.g., detection of comments and supplementary material), data cleansing (e.g., removing noise from inputs, discarding useless parts of the documents), as well as linguistic annotations that facilitate data indexing are some of the functionalities foreseen in this category of services.

3.3 Advancements of the Proposed Solution

With respect to *large-scale online deliberation*, the proposed solution provides two main technological advancements to the state of the art:

- ***Improved Sense Making.*** Large scale deliberations are hard to monitor and make sense of. The proposed platform will develop powerful deliberation analytics and visual interfaces to make sense and assess the state, progress and quality of a deliberation process, as well as alerts that guide users to the parts of the discussion where they can offer most.
- ***Improved Evidence-Based Thinking.*** Large scale discussions often produce shallow content and low-quality debate. The proposed solution takes a knowledge-based approach to improve deliberation quality. The foreseen deliberation platform will provide a series of features and services to recommend scientific literature to participants during the debate. In this way, the proposed solution fosters evidence-based thinking and more informed discussion, which improve the overall quality of the deliberation.

As far as *argumentation mining and aggregation* are concerned, the proposed solution will shape novel ways of supporting and facilitating online deliberations: (i) the foreseen framework will employ automatically extracted arguments in order to improve decision and policy making and support strategic actions; (ii) it will adopt a joint discourse-level and information-seeking perspective paying much attention to argumentation clustering and argumentation aggregation procedures for the context under consideration, and (iii) it will employ novel argumentation mining pipelines paying particular attention to argument quality, while also facilitating the creation of explanations that disclose how the information on which the machine relies to make its own decisions is retrieved and interpreted (Karacapilidis *et al.*, 2017).

The proposed solution will also advance *large-scale decision-making support technologies*, by adopting a knowledge-based decision-making view, enabled by the meaningful accommodation of the results of the argumentation mining and aggregation processes. According to this view, decisions will be considered as pieces of descriptive or procedural knowledge referring to an action commitment. In such a way, the decision-making process is able to produce new knowledge, such as evidence justifying or challenging an alternative or practices to be followed or avoided after the evaluation of a decision, thus providing a refined understanding of the problem under consideration. On the other hand, in a decision-making context the knowledge base of facts and routines alters, since it has to reflect the ever-changing external environment and internal structures of citizen assemblies. Knowledge management activities such as knowledge elicitation, representation and distribution influence the creation of the decision models to be adopted, thus enhancing the decision-making process.

The abovementioned synergy of decision-making and knowledge management will be further strengthened in the proposed platform by the incorporation of features enabling decision makers to perform argumentation on the issues raised. Many collaborative decision-making problems have to be solved through dialoguing and argumentation among a group of people. In such contexts, conflicts of interest are unavoidable and support for achieving consensus and compromise is required. Independently of the model used for decision making, argumentation is valuable in shaping a common understanding of the problem. It can provide the means to decide which parts of the information brought up by the decision makers will finally be the input to the model used. Moreover, argumentation may stimulate the participation of citizens and decision makers and encourage constructive criticism. To address the above category of requirements, a user-friendly argumentative deliberation-based decision-making support environment will be developed (Christodoulou *et al.*, 2016).

The proposed solution integrates novel mechanisms to aggregate citizens' and subgroups' opinions into collective positions. Different aggregation functions will be tested to assess the robustness of the results obtained. These mechanisms will comply with the foreseen transition model for the scaling of an ongoing deliberation. In addition, the proposed solution can further elaborate the consensus reaching process through the development of new models in which consensus is measured through aggregated collective opinions at both the intra-subgroup and the inter-subgroup levels. Aiming to augment the explainability and interpretability of the models and data involved in the overall large-scale group decision making process (Samek *et al.*, 2019), the proposed solution will also develop and integrate in the foreseen deliberation framework a dedicated explanation mechanism that will benefit the user in terms of *justification* (exposing the reasoning behind a decision may help the user decide how much credence to give in it), *user involvement* (allowing the user to add her knowledge and inference skills to the overall decision process), and *system acceptance* (in that the system's functionality is fully visible and its suggestions are adequately justified).

Finally, the proposed solution will employ a novel *knowledge graph* to model stakeholders' knowledge and interactions jointly. We plan to advance current knowledge aggregation methods that are based on neural architectures such as attention mechanisms and Graph Neural Networks (GNNs). The foreseen advancement will build on

large-scale pre-training via transformers and variants like BERT models. Large-scale pre-training models will aid the acquisition (and injection in the KG) of factual knowledge. Considering its overall objectives, the proposed solution will rethink the way of knowledge aggregation in an efficient and interpretable manner. Specifically, it enables the construction of dynamic knowledge graphs, together with novel mechanisms to capture the dynamics of a deliberation, thus addressing limitations of traditional knowledge representation and reasoning by meaningfully monitoring and analyzing the temporal dimension. It will advance current approaches to thoroughly address the scalability issue, which is certainly crucial in large-scale knowledge graphs. The ubiquitous trade-off between computational efficiency and model expressiveness will be addressed through transformer-based models to encode graph entities, relations and path sequences, as well as GNNs to aid the learning of connectivity structure under an encoder-decoder framework. For the analysis of deliberation data, the foreseen KG builds on a graph-based text representation, namely *graph-of-docs* (Giarelis *et al.*, 2020). The proposed knowledge graph structure and related advancements also serves explainability purposes, aiming to aid stakeholders build a complete and informed mental model of the inferential process of the underlying machine learning algorithms and the knowledge-based decision-making support system and promote trust for its outputs. The proposed explanations generator engine adopts a human-in-the-loop approach towards the development of interactive interfaces to support model interpretability and inference explainability.

4 Conclusions

This paper has described a novel solution that adopts a pluralist and bottom-up approach to increase the quality of deliberation and its ability to influence public policy. The proposed solution aims to facilitate and augment the scaling of this approach through an AI-enhanced digital democracy platform that builds on prominent ML/NLP technologies to enable lay and expert stakeholders exchange and reform their opinions, co-create actionable solutions, and collectively reach decisions in a highly transparent and trustful way. We argue that this solution will contribute to the improvement of the quality of democracy nowadays, which demands the active involvement and effective participation of citizens in policy making from the design to the implementation phase. It can be viewed as a digital transformation tool that is an essential enabler of a socially cohesive society, where all individuals and groups have a sense of belonging, participation, inclusion, recognition and legitimacy.

The proposed solution has been shaped through long and fruitful collaboration among diverse types of stakeholders (representing academia, citizens and civil society, government and public authorities, and ICT-focused SMEs), through which a series of rich application scenarios have been sketched and analyzed. The main limitation of our study is that though the proposed approach has undergone a first level assessment and validation by experienced practitioners, which has been highly positive, its application has to be carefully planned by taking into account the capacity and available resources of diverse organizations. Future work directions include the full implementation and integration of the proposed solution's modules and services, as well as the collection of feedback through its assessment in diverse deliberation settings, ranging from a local

to an international level. Its application will be evaluated through a set of dedicated Key Performance Indicators, focusing on the usefulness and ease of use of the proposed approach.

References

Androutsopoulou, A., Karacapilidis, N., Loukis, E., Charalabidis, Y.: Combining technocrats' expertise with public opinion through an innovative e-participation platform.IEEE Trans. Emerg. Top. Comput. **2018** (2018). https://doi.org/10.1109/TETC.2018.2824022

Bodanza, G., Tohmé, F., Auday, M.: Collective argumentation: a survey of aggregation issues around argumentation frameworks. Argument Comput. **8**(1), 1–34 (2017)

Cabrio, E., Villata, S.: Five years of argument mining: a data-driven analysis. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-18), pp. 5427–5433. AAAI Press (2018)

Checkland, P., Holwell, S.: Action research: Its nature and validity. Syst. Pract. Action Res. **11**(1), 9–21 (1998)

Christodoulou, S., Karacapilidis, N., Tzagarakis, M.: Exploiting alternative knowledge visualizations and reasoning mechanisms to enhance collaborative decision making. In: Tweedale, J.W., Neves-Silva, R., Jain, L.C., Phillips-Wren, G., Watada, J., Howlett, R.J. (eds.) Intelligent Decision Technology Support in Practice. SIST, vol. 42, pp. 89–106. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-21209-8_6

Devlin, J., Chang, M.-W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)

Giarelis, N., Kanakaris, N., Karacapilidis, N.: On a novel representation of multiple textual documents in a single graph. In: Czarnowski, I., Howlett, R.J., Jain, L.C. (eds.) IDT 2020. SIST, vol. 193, pp. 105–115. Springer, Singapore (2020). https://doi.org/10.1007/978-981-15-5925-9_9

Hevner, A.R., March, S.T., Park, J., Ram, S.: Design science in information systems research. MIS Q. **28**(1), 75–105 (2004)

Karacapilidis, N., Malefaki, S., Charissiadis, A.: A novel framework for augmenting the quality of explanations in recommender systems. Intell. Decis. Technol. J. **11**(2), 187–197 (2017)

Karacapilidis, N., Christodoulou, S., Tzagarakis, M., Tsiliki, G., Pappis, C.: Strengthening collaborative data analysis and decision making in web communities. In: Proceedings of the 23rd International World Wide Web Conference (WWW2014), Companion Volume - Workshop on Web Intelligence and Communities, Seoul, Korea, 7–11 April 2014, pp. 1005–1010 (2014)

Karacapilidis N. (ed.): Mastering Data-Intensive Collaboration and Decision Making: Cutting-Edge Research and Practical Applications in the Dicode Project. Studies in Big Data Series, vol. 5, Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-319-02612-1

De Liddo, A., Sándor, Á., Shum, S.B.: Contested collective intelligence: rationale, technologies, and a human-machine annotation study. Comput. Support. Coop. Work **21**(4–5), 417–448 (2012)

Lin, H., Liu, Y., Wang, W., Yue, Y., Lin, Z.: Learning entity and relation embeddings for knowledge resolution. Proc. Comput. Sci. **108**, 345–354 (2017)

Lippi, M., Torroni, P.: Argumentation mining: state of the art and emerging trends. ACM Trans. Internet Technol. **16**, 2 (2016). Article 10. https://doi.org/10.1145/2850417

Liu, B.S., Zhou, Q., Ding, R.X., Palomares, I., Herrera, F.: Large-scale group decision making model based on social network analysis: trust relationship-based conflict detection and elimination. Eur. J. Oper. Res. **275**(2), 737–754 (2019)

Peters, M., et al.: Deep contextualized word representations. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics, vol. 1, pp. 2227–2237 (2018)

Rapoport, R.N.: Three dilemmas in action research. Hum. Relat. **23**(6), 499–513 (1970)
Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Müller, K.-R. (eds.): Explainable AI: Interpreting, Explaining and Visualizing Deep Learning. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-030-28954-6
Tang, M., Liao, H.: From conventional group decision making to large-scale group decision making: what are the challenges and how to meet them in big data era? A state-of-the-art survey. Omega 102141 (2019). https://doi.org/10.1016/j.omega.2019.102141
Wang, Q., Mao, Z., Wang, B., Guo, L.: Knowledge graph embedding: a survey of approaches and applications. IEEE Trans. Knowl. Data Eng. **29**(12), 2724–2743 (2017)
Wang, Z., Zhang, J., Feng, J., Chen, Z.: Knowledge graph embedding by translating on hyperplanes. In: Proceedings of AAAI 2014, pp. 1112–1119 (2014)