
Enhancing structured democratic dialogues: evaluating the impact of AI assistants on participant contributions

Yiannis Laouris*

Futures Design Unit,
Future Worlds Centre,
Thrakis 2, 2313 Nicosia, Cyprus
Email: laouris@cnti.org.cy
*Corresponding author

Lazar Kovacevic

Inverudio Inc,
Chicago, IL USA
Email: lakinekaki@gmail.com

Marcus Hallside

Future Worlds Centre,
Thrakis 2, 2313 Nicosia, Cyprus
Email: marcus.hallside@innovativecompliance.com

Abstract: The structured democratic dialogue (SDD) process involves distinct stages designed to enhance participants' understanding and refinement of each other's contributions. Initially, participants generate titles for their responses to a predetermined triggering question, followed by brief clarifications. As the dialogue progresses, further exploration of ideas necessitates additional elucidation, deepening their meaning and concreteness. This paper investigates the integration of an AI-enriched model into the process, through a Python-based program utilising ChatGPT-4, named AI-SDD-Assistant©. The study evaluates the AI-assistant's effectiveness in enhancing participants' original contributions by making them specific, measurable, assignable, relevant, and time-bounded (SMARTer). Participants utilised the AI-assistant to refine their ideas based on clarification questions from others in subsequent stages. Results indicate that the majority of participants approved the AI-assistant's enhancements, resulting in more specific, measurable, assignable, and relevant responses, with time-bound estimates. While one-third of the participants remained skeptical, they still endorsed the use of AI-assisted revisions by organisers.

Keywords: structured democratic dialogue; SDD; virtual assistant; AI-enhancement; collective intelligence; dialogue facilitation; stakeholders.

Reference to this paper should be made as follows: Laouris, Y., Kovacevic, L. and Hallside, M. (2026) 'Enhancing structured democratic dialogues: evaluating the impact of AI assistants on participant contributions', *Int. J. Applied Systemic Studies*, Vol. 13, No. 3, pp.267–289.

Biographical notes: Yiannis Laouris is the 2025–2026 President of the International Society for the Systems Sciences, President of the Institute for 21st Century Agoras, and Lead Scientist of Future Worlds Center. He is a peace and democracy activist honoured with 13 awards (and ‘Honorary Fellow’ of the School of Sciences of the University of Lancashire, UK) for his pioneering work in utilising digital technologies and structured democratic dialogue to harnesses collective intelligence and wisdom to reimagine and redesign humanity’s future based on principles of freedom, equality, and justice, as described in his book: *Democratic [R]evolution: A Call for Collective Imagination and Systemic Action*.

Lazar Kovacevic is an independent researcher specialising in the application of IT to education, creativity, collaboration, and social action. He has done many projects in areas of (web) information retrieval systems, text analysis and natural language processing, machine learning, data mining, collaboration, etc. He enjoys participating in multidisciplinary environments, working on interdisciplinary solutions to real-world problems. He co-authored several papers discussing creative features in time series, ranging from physical and biological to physiological and psychological processes (i.e., a healthy heart exhibits more innovative features than an unhealthy one). He developed algorithms to increase the diversity of perspectives in search results.

Marcus Hallside, a retired Israeli IT entrepreneur, is an active member of the International Society for Systems Sciences, Member of the Board of Institute for 21st Century Agoras, and associate researcher at Future Worlds Center. In 1999, he received a \$2 million FAA-grant to lead a multi-year air cargo security R&D program, resulting in the first operational Known Shipper/Trusted Trader initiative for the new US Transportation Security Administration. Marcus applies structured dialogic design to complex socio-technical challenges, utilising structured democratic dialogue, to engage people from diverse backgrounds and perspectives in meaningful conversations, thereby building consensual, normative visions of the future.

This paper is a revised and expanded version of a paper entitled ‘Bridging AI governance and collective intelligence: a systems science approach to global challenges’ presented at 21th On-Line HSSS National & International Conference, 20th National & International Conference, Systemic Human Resources Management and New Technologies, Tripoli, Greece, 8–11 October 2025.

1 Introduction

In the era of hyper-connectivity and amid contemporary grand societal challenges concerning the governance of humanity, imperative questions arise: how do we identify and involve the appropriate stakeholders for each unique situation; how can we develop and execute systems ensuring equitable and wise outcomes for all parties involved; how do we safeguard the genuineness of individuals’ perspectives; how do we attain genuine rather than fallacious equality among citizens. These concerns have become paramount, especially in an era where people spend significant time online (Laouris, 2015). The structured democratic dialogue (SDD) methodology has emerged as a practical and powerful approach for addressing these challenges. It facilitates small groups of authentic stakeholders to engage in face-to-face collaborative democratic deliberations (referred to

as co-laboratories of democracy: Christakis and Baush, 2006; Laouris et al., 2014), allowing them to exchange viewpoints, organise complex socio-technical issues in manageable formats (Laouris and Romm, 2022a), reach a common understanding of the challenges they face, and importantly, establish consensus on pathways for progress (Christakis and Baush, 2006). The term ‘authentic’ emphasises the importance of engaging the true stakeholders, i.e., those whose lives will be affected by any decisions or changes in the system in which they are embedded [Laouris et al., (2008a), pp.332–334]. SDD ‘enables people from all walks of life to act as systems scientists and to harness their collective wisdom without needing them to understand all the complexities and jargon of systems science to address and resolve complex socio-technical challenges’ [Laouris and Michaelides, (2018), p.919]. The structured process empowers stakeholders to address challenges collaboratively and reduce the situational complexity index (Flood, 1988).

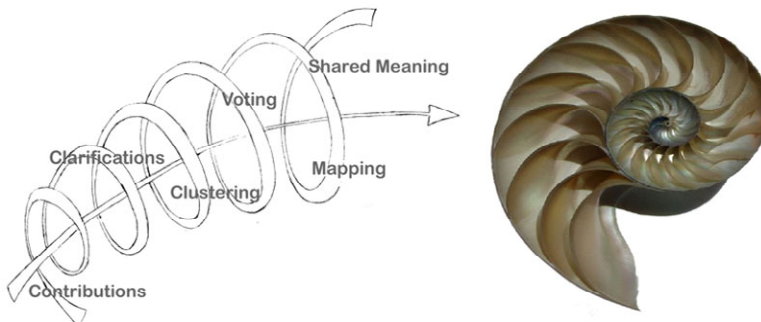
Two significant shortcomings of SDD are that it works best in small groups and two in face-to-face synchronous interactions. The third shortcoming, sometimes brought up by opponents of SDD, is that lay people (SDD engages stakeholders from all walks of life) might lack the skills, knowledge and expertise to support their arguments. In the following paragraphs, we elaborate on how SDD scientists attempt to alleviate these shortcomings. This paper examines whether inserting AI assistance within the SDD process can alleviate some of these shortcomings.

The increasing rate of change and the rapidly growing complexity make the need for effective large-scale reforms a pressing emergency. To this end, the authors have been experimenting with 3rd, 4th and 5th generation tools (Laouris, 2023a, 2003b; Laouris and Dye, 2024), utilising emerging communication technologies to render SDD wholly or partly virtual and partly or fully asynchronous. Various hybrid models that combine synchronous vs. asynchronous and virtual vs physical interactions have been tried (Laouris and Metcalf, 2024), and at least three models of scaling up have been developed and tested (Laouris, 2022, 2023a; Laouris and Midgley, 2025).

An instrumental characteristic of a successful SDD (or any other type of constructive dialogue) is that it creates an environment in which the ideas or concepts being discussed can ‘grow’ both in terms of becoming more concise, i.e., well-bounded concerning other similar concepts. Moreover, also their meanings are becoming more ‘spherical’ and ‘deeper’ (i.e., SMART: specific, measurable, assignable, relevant, time-bounded, and others). In the case of SDD, the sequence of stages in the process promotes evolutionary learning (Figure 1). The SDD views stakeholders as a learning community aspiring to reach a state where all participants achieve a deeper but also shared understanding of the issue they are discussing (law of requisite meaning, Turrisi, 1997; see also Christakis and Bausch, 2006; Laouris, 2012). During the first two stages (i.e., idea generation and idea clarification) of the process (see Figure 2, upper part), participants share their existing knowledge and different points of view. Peers are encouraged to request clarifications, but they are not allowed to make any value statements regarding the statements or contributions of others. The authenticity and autonomy of each participant are ‘protected’ through compliance with SDD guidelines that call for respect and tolerance (law of requisite autonomy in decision, Tsivacou, 1997). Upon completing these two initial stages, participants are supported and urged to expand their explicit knowledge of the issue at hand by encouraging them to explore together relations of similarity and influence between their ideas. The interaction between individuals within an SDD process is critical in further developing their ideas, which aligns with Vygotsky’s (1978)

theory of constructivism. In other words, ‘communities of interaction’ contribute to the amplification and development of new knowledge. Learners discover new aspects and perspectives as they explore and compare ideas for similarity or influence, hence expanding their explicit knowledge about the issue. Their ideas develop in-depth and quality but not necessarily in numbers. Learners achieve a much deeper understanding, as illustrated in the nautilus spiral, with the nautilus separators increasing in surface (Figure 1).

Figure 1 Evolutionary learning. knowledge is acquired in incremental phases. during each subsequent stage of the SDD process, participants acquire meaning and wisdom in an evolutionary manner (see online version for colours)



Adhering to the laws of structured dialogic design cultivates autonomy, facilitates evolutionary learning (Dye and Conway, 1999), and assists participants in achieving meaning and wisdom. Out of these primarily cognitive processes, action emerges as a natural consequence (law of requisite action; Laouris et al., 2008a), translating to commitment, endurance, and intentionality of learning. The SDD process facilitates better understanding and renders learners self-driven and enthusiastic, therefore serving the learner-centred principle. In line with philosophers such as Mikhail Bakhtin (in: Todorov, 1984; see also Bohm, 1987, and Pietroski, 2021 discussing Chomsky’s views on meaning), SDD posits that meaning is not pre-formed but instead emerges iteratively through dialogue. This means an original author’s meaning is not yet discovered until that author engages in dialogue. Each iterative statement moves toward clarification of meaning, while each attempt to make an understandable statement moves toward clarification of communication.

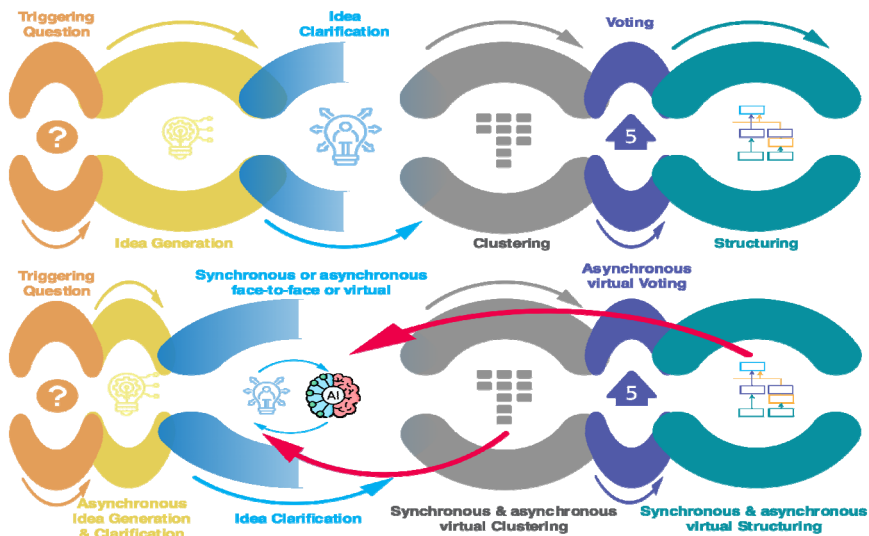
To summarise, we claim that the SDD process empowers people from all walks of life to act as experts and harness their collective wisdom without needing them to understand all the complexities and jargon of systems and complexity science [Laouris and Michaelides, (2018), p.919]. Nevertheless, some critics argue that participants may be disadvantaged in addressing complex issues or designing ideal solutions if they lack the necessary skills, professional knowledge or expertise. When participants do not fully understand important technical aspects of an issue, the quality of their proposed solutions can be compromised. This concern is based on the idea that while SDD is designed to democratise problem-solving by including diverse perspectives, the absence of specialised knowledge might lead to oversimplified or suboptimal outcomes and limit the acceptance or diffusion of their proposals (Kalliola and Nakari, 1999). This criticism does not apply to cases where the anticipated stakeholder participants are ‘experts’ with rich

knowledge and sufficient experience to evaluate and analyse the current situation (see examples in Laouris and Michaelides, 2018 and Michaelides and Laouris, 2024). Other critics state that while public participation is essential for legitimacy and democratic governance, it might not always produce the best technical solutions. The latter is especially true in highly specialised fields where professional expertise is crucial. For instance, the literature discusses the challenges of achieving meaningful engagement when participants lack the requisite expertise, which can sometimes lead to decisions that are not technically sound or are based on incomplete understandings of the issues (Pruitt and Thomas, 2007, p. 241; Althaus et al., 2022, p.66). The SDD can and has been utilised by groups of experts, such as the case of IT experts across Europe convening to evaluate the obstacles that prevent the broader public from benefiting from and participating in the broadband society (Laouris et al., 2008b) or to develop technology-transfer road maps to assist the development of assistive and ICT products (Roe et al., 2011).

2 Method

The experiments described in this paper were conducted in the context of a virtual, primarily synchronously implemented SDD with experts from across the globe interested in exploring and discussing the possible negative repercussions and ways of addressing the emerging powerful AIs for society (Laouris et al., in preparation). In response to the triggering question (TQ), ‘What actions -or options, initiatives, policies, regulations, and collaborative efforts - can humanity undertake to protect itself from the potential negative repercussions of AI development and deployment?’ eighteen participants generated 64 responses. The participants included AI developers, and experts in dialogue (e.g., SDD/ISM), ICT, cyber security, research, policy, ethics, sustainable development, human rights, global education, and other societal issues.

Figure 2 Typical (upper part) vs. AI-enriched models of the SDD process (see online version for colours)



The SDD process has been widely described (e.g., Christakis and Bausch, 2006; Flanagan, 2006; Flanagan and Christakis, 2021; Laouris and Michaelides, 2018; Laouris and Romm, 2022a). A revised model that allows for hybrid synchronous vs. asynchronous virtual interactions was applied (for the evolution of the process, see Laouris and Dye, 2024; for hybrid models, see Laouris and Metcalf, 2024) for this project.

Figure 3 Python code of the AI-SDD-Assistant© with ChatGPT-4 prompt (see online version for colours)

```
import re, openai, requests, time, json

openai.api_key = HIDDEN

def gpt(text, prmpt):
    time.sleep(1)
    response = openai.chat.completions.create(
        model="gpt-4",
        messages=[
            {"role": "user", "content": text+"\n\n"+prmpt},
        ]
    )
    response = json.loads(response.json())
    try:
        if response['choices'][0]['finish_reason'] != 'stop':
            return response['choices'][0]['message']['content']
    except:
        print(response)
    return ""

with open("Planetary_AI_SDD_20231119.XML", 'r', encoding='utf-8') as file:
    xml_data = file.read()

trigger = re.search(r'<Name>(.*?)</Name>', xml_data, re.DOTALL).group(1)

pattern = r'<idea>.*?<Num>(ld+)</Num>ls*<IdeaText>(.*?)</IdeaText>.??<clarification>(.*?)</clarification>.??</idea>'

matches = re.findall(pattern, xml_data, re.DOTALL)
ideas_dict = {}

for match in matches:
    num = match[0]
    title = match[1]
    clarification = match[2]
    ideas_dict[num] = (title, clarification)
    text = ""
    - To the Question "{trigger}" a workshop participant provided the following response:
    {title}:
    [clarification]
    - I would like you to improve this response making it SMARTer, at the same time more Specific, Measurable, Assignable, more Relevant to the Question"" format(trigger=trigger, title=title, clarification=clarification)
    ai = gpt(text, prmpt="Please assume the role of an AI expert who is knowledgeable about the threats and possible negative repercussions of AI. ")
    print('~~~~~')
    print(num, title)
    print(clarification)
    print('AI: ', ai)
```

The experiment explicitly aimed to explore whether AI could support participating stakeholders to improve the clarifications of their contributions they provide during the clarification stage of the SDD (see blue sub-diagram in Figure 2), thus enhancing their own and their peers learning and understanding regarding the meaning of their contribution. The experiments were conducted using an AI-enriched model of the SDD process, which utilised a Python-based software named AI-SDD-Assistant© programmed to interact with ChatGPT version 4. Each participant’s clarification was submitted to the AI assistant, asking it to make it SMARTer, i.e., more specific, measurable, assignable, relevant, and placed within a reasonable time horizon.

Specifically, the prompt’s structure comprised three parts:

- 1 ‘Please assume the role of an AI expert who is knowledgeable about the threats and possible negative repercussions of AI’
- 2 ‘To the question {trigger question} a workshop participant provided the following response: {title} {clarification}’
- 3 ‘I would like you to improve this response, making it SMARTer, i.e., specific, measurable, assignable, and more relevant to the {trigger question}.’

Table 1 Questions used to evaluate the AI-SDD-Assistant©’s contributions

<i>Evaluate the degree to which, in your judgment, the AI has improved your original clarification.</i>				
<i>Not at all</i>	<i>Somewhat</i> (20–50%)	<i>A lot</i> (50–100%)	<i>Significantly</i> (200 + %)	<i>Enormously</i> (300 + %)
To what degree has AI contributed in making your clarification (Use same scales as above)				
<i>Not at all</i>	<i>Somewhat</i> (20–50%)	<i>A lot</i> (50–100%)	<i>Significantly</i> (200 + %)	<i>Enormously</i> (300 + %)
More specific				
More measurable				
Better assignable				
More relevant to the TQ				
Time bounded (provided estimates for time urgency)				
Based on this experience, how would you consider using ChatGPT when participating in an SDD?				
I would prefer to use it myself to edit and improve my clarification(s)				
I find it helpful if the organisers run the ideas through Chat GPT and provide me with an edited-improved version of my own contribution to review, correct and approve				

It is important to note that all participants' clarifications were submitted to the software in the exact same format. The AI assistant's responses were fed back to the participants, who were free to decide whether or how to use them to consider improving their original clarification based on the AI Assistant's response. At the end of the experiment, the participants were invited to score the usability of the AI Assistant's responses using three evaluation questions documented in Table 1:

- a the degree to which the AI has improved their original clarification
- b the degree to which the AI contributed in making their clarification SMARTer
- c whether they found it helpful to have the AI assistant provide them with an edited-improved version of their contribution to review, correct and approve.

Figure 2 presents the classical (upper part) and the hybrid SDD model (lower part). The embedding of the AI assistant has been tried in the hybrid model. More specifically, after the participants had offered their clarifications to the group (blue sub-diagram in the figure), the AI-SDD-Assistant© was run, and its responses were fed back to the participants. The additional feedback loops, depicted with the backward-pointing red arrows in the figure, exemplify that participants could utilise the AI assistant's responses to revisit or improve their clarification also later in the SDD process (e.g., during the stages of clustering or mapping).

3 Results

The data stems from an SDD conducted in the fall of 2023, engaging eighteen societal and scientific actors from across the globe. Twelve actors who agreed to participate in this study were given an AI-SDD-Assistant© enhanced version of their clarification, as they provided it in a shared Google Doc. If they had contributed more than one idea to the dialogue, their first one was used. They were requested to consider possible improvements to their clarification based on the AI's responses. Later, they were asked to evaluate their experience using the abovementioned questions.

3.1 *AI enriching original contributions*

The following sections present three examples demonstrating how the participants utilised the AI's responses to enrich their original clarifications. The original text submitted by an author as clarification is at the top, followed by the (typically extended) clarification offered by the AI in response to the Prompt as described in the methods section. Edits conducted subsequently by the author are highlighted in red, while additions of text present in the author's original clarification but omitted by AI are highlighted in light blue. Where an author responded to the AI assistant's suggestions, those comments are shown in brown.

3.2 *Example 1: transparent AI research and development*

The first example (Figure 4) concerns the 'transparent AI research and development' proposal.

Figure 4 Example 1: transparent AI research and development (see online version for colours)

Encouraging transparency in AI research and development can help prevent potential negative repercussions of AI deployment. This involves promoting open-source AI technologies, sharing research findings, and ensuring clear communication about the capabilities and limitations of AI systems. By fostering a culture of transparency and accountability, stakeholders can collaboratively address ethical, safety, and societal concerns related to AI technologies.

AI Response [Transparent AI Research & Development] including human edits:

Encouraging transparency in AI research and development can help prevent potential negative repercussions of AI deployment. This involves promoting open-source AI technologies, sharing research findings, and ensuring clear communication about the capabilities and limitations of AI systems. By fostering a culture of transparency and accountability, stakeholders can collaboratively address ethical, safety, and societal concerns related to AI technologies.

- Specific:** In **relation** to the cons and pros of AI, one crucial factor to consider is the transparency of AI conceptualization, research, and development. Concretely, the risks associated with the misuse of AI technology require us to ensure full transparency in every phase of its development. **Measurable:** Specifically, a transparency protocol must be implemented to detail all AI-related projects' objectives and methods, thereby avoiding surreal or misguided expectations.
- Measurable:** **Ideally**, this should include indicators that quantifiably measure the level of transparency, such as a transparency index that ranges from 0 to 100.
- Assignable:** **This** protocol should be assignable, meaning it could be delegated to a dedicated team or individual, preferably an AI Ethics Officer or similar role. Their function would be to oversee compliance with the protocol maintaining the standards while keeping abreast of new AI development trends and their potential threats.
- Relevant:** **Finally**, to stay relevant to the pros and cons of AI, this protocol should serve to mitigate potential risks, magnify benefits, **This** and ensure that ethical considerations are always paramount in AI research and development.
- Timer-bounded:** **The above** will ensure that AI develops at a pace that society can manage and guide, rather than the other way around.

~~Therefore, a transparent AI research and development approach stands as an essential factor in controlling the cons while maximizing the pros of AI technology.~~

The author made three types of corrections in Example 1. First, the AI failed to include the original clarification provided by the author, who had to add it, as depicted with the green-coloured text. It should be noted, however, that our Prompt has not provided instructions on including the original or AI-modified text in the response. Second, while the AI provided ideas on how to make the clarification measurable, more relevant to the TQ, and time-bounded, those were not presented as separate sections, and the author had to move sentences to new paragraphs to organise the presentation better. Note the additions (marked in red) of subsection headings for specific, measurable, relevant, and time-bounded. Third, ChatGPT often concludes a response with a summary sentence, which the author did not consider necessary and deleted it (red strikethrough over text).

3.3 Example 2: Regulations that require responsible and ethical AI development

The next example (Figure 5) concerns the idea of 'regulations that require responsible and ethical AI development'.

The AI response had similar shortcomings as the previous example. It did not properly include the original text (added by the participant; in light blue) and did not structure the response as requested for the SMART criteria. More importantly, the responses for Measurable and Relevant were too short. The original author had to make more corrections to the syntax to convert the response into a logically structured, flowing text.

Figure 5 Example 2: regulations that require responsible and ethical AI development (see online version for colours)

This is going to be particularly difficult given that AI companies have operated without much oversight and there is an enthusiasm of policymakers and business leaders that generative AI is going to allow companies and countries to be economically competitive. However, there are few policies that require a proactive approach to AI development that considers the ethics and implications. Unfortunately, engineers and tech entrepreneurs are not used to considering the wider society when they develop technologies. Therefore, it is going to be imperative that both education and regulations work together to force this wider consideration of societal responsibilities and ethics. I have been involved with attempts to create AI regulations and ethical AI standards. These are not easy to do. There is a desire by industry and experts to avoid these kinds of regulations. This makes it even more imperative to figure out how to create some kind of agreed-upon framework that considers the larger systemic issues.

AI Response [Regulations that require responsible and ethical AI] including human edits:

On-the-downside, While AI has the potential to greatly improve many aspects of human life, it also poses significant threats if not properly regulated. For example, AI systems can be programmed to engage in unethical behavior and can even be used as weapons in the wrong hands. Without proper regulation, we risk allowing the development and deployment of AI systems that could harm society or certain individuals in significant ways. **[original added manually]** This is going to be particularly difficult given that AI companies have operated without much oversight and there is an enthusiasm of policymakers and business leaders that generative AI is going to allow companies and countries to be economically competitive. However, there are few policies that require a proactive approach to AI development that considers the ethics and implications. Unfortunately, engineers and tech entrepreneurs are not used to considering the wider society when they develop technologies. **Therefore,**

Specific: It is going to be imperative that both education and regulations work together to force this wider consideration of societal responsibilities and ethics. **[added here from the original]** I have been involved with attempts to create AI regulations and ethical AI standards. These are not easy to do. There is a desire by industry and experts to avoid these kinds of regulations. This makes it even more imperative to figure out how to create some kind of agreed-upon framework that considers the larger systemic issues.

On the upside, there is a growing recognition of these threats, and there are already efforts underway to develop regulations to prevent them. Potential regulations could include guidelines for ethical AI development, requirements for transparency in AI decision-making, and controls on how AI can be used in certain risky areas like weapons development. For these regulations to be effective, they would need to be specific enough to prevent harmful use of AI.

Measurable: Regulations must be measurable in their impact.

Assignable: Regulations must be assignable to specific entities that are responsible for enforcing them. **[the following moved here from end]** As an AI community, we should strive to work with policymakers to ensure that these regulations are put in place and effective in mitigating the risks of AI.

Relevant: Must be relevant to the particular threats posed by AI. **and**

Time-bounded: Must be time-bounded in the sense that they are enacted before AI advances to the point where they would be difficult to control.

3.4 Example 3: human rights as the foundations for AI alignment

The third example (Figure 6) concerns the idea of ‘human rights as the foundations for AI alignment.’

The AI produced a well-structured response regarding the SMART criteria. However, just like in the previous cases, it failed to include the original (or improved) author’s contribution in its response.

Figure 6 Example 3: human rights as the foundations for AI alignment

A breakthrough step has been done when in 1948 the UN made official the Human Rights declaration. I believe we should use this as a foundation to build up the different layers that involve human goals, citizen contribution, society responsibility, and AI wisdom framework because at the end of the day the Human Rights are the most fundamental declaration for our species, all other element should serve, protect, and ease its realization.

Making the Universal Declaration of Human Rights the ethical standard for healthy living with AI will require a serious review of this visionary declaration for clear indicators and parameters to be developed. The transnational courts have progressed in this regard and they should be included in this process.

AI Response [Human Rights as the foundations for AI alignment] including human edits:

A breakthrough step has been done when in 1948 the UN made official the Human Rights declaration. I believe we should use this as a foundation to build up the different layers that involve human goals, citizen contribution, society responsibility, and AI wisdom framework because at the end of the day the Human Rights are the most fundamental declaration for our species, all other element should serve, protect, and ease its realization.

Making the Universal Declaration of Human Rights the ethical standard for healthy living with AI will require a serious review of this visionary declaration for clear indicators and parameters to be developed. The transnational courts have progressed in this regard and they should be included in this process.

- Specific:** Developing AI systems should be aligned with the principles of human rights to reduce potential threats and negative impacts on society. This includes respecting privacy, ensuring equality and non-discrimination, and enabling accountability.
- Measurable:** Success can be evaluated by how well AI systems mitigate and manage risks related to privacy infringement, discrimination, and accountability issues. Quantitative measurement can be done by tracking occurrences of these instances, while qualitative measurement can involve assessments from third-party human rights organizations.
- Assignable:** AI developers, policy makers, and human rights organizations need to collaborate in defining human-rights-centered guidelines for AI development. Technology companies need to assign dedicated teams to ensure these guidelines are followed and implemented in AI systems.
- Relevant:** Aligning AI with human rights principles is directly relevant to the potential cons of AI, which include privacy violation, discrimination, and lack of accountability. By addressing these issues, the benefits of AI can be harnessed more effectively and safely.
- Time-bound:** There should be initial alignment guidelines launched within the next two years, with regular revisions and implementation reviews. This creates urgency keeping pace with the rapid development of AI technology, ensuring human rights are always safeguarded.

3.5 Engaging in three-party conversations

A few participants have sent comments the authors responding to the AI assistant’s suggestions, which demonstrates that they engaged in a kind of ‘dialogue’ with the AI assistant’s ideas. In the fourth example (Figure 7), the text in brown corresponds to the responses of the human author to the AI assistant’s suggestions. It is particularly

interesting how this human participant engaged in a three-party conversation, which included the AI and the facilitators.

Figure 7 Example 4: AI: author's comments to revised clarification (see online version for colours)

AI: Refine the Development and Implementation of Regulatory Measures for AI Deployment and Impact Assessment

Specific:	Develop comprehensive regulations addressing the creation, deployment, and potential impacts of AI on various sectors. Regulations should focus on the ethical use, privacy concerns, potential biases, data security, and possible socioeconomic impacts of AI. In my clarification I tried to move away from the word regulation towards the broader concept of managing the impacts. The word regulation is in the title, so the AI carried the word through and adjusted its interpretation of the content of my response accordingly. My interest here was more in the everyday practice in workplaces of how, why and where AI is used and the importance of well-informed guidelines to maintain decision quality, veracity and overall value to society.
Measurable	Use key performance indicators such as the number of AI systems that pass the regulatory standards, number of ethical breaches or issues encountered, or changes in public sentiment towards AI systems as the regulatory measures get implemented. Guidelines for the use of AI in the workplace should follow the latest regulations available and 'measured' accordingly. They should be updated as often as necessary in lieu of new regulations and information about AI.
Assignable:	Multiple stakeholders need to be involved in this process, including AI developers, government entities, non-government organizations focused on human rights and digital privacy, and even representatives from the public to present a broad perspective. Multiple stakeholders are important in setting up guidelines for practice. Though the ultimate responsibility will be with 'managers' in workplaces to lead the processes.
Relevant:	Given the increasing influence of AI across various sectors, a strong regulatory framework is relevant and necessary to ensure that its development and use are in alignment with ethical standards and public interest. Potential negative impacts and repercussions of AI need to be studied, understood, and mitigated through these regulations. I can work with the response here, except for the reference to regulations at the end. I would prefer management. The management of AI technologies in the everyday workplace needs to be considered. It is a more everyday issue. All aspects of use management cannot be captured in regulations. Regulations will focus on the big overarching areas of concern, though they can inform management practice.
Time-Bound:	Aim to finalize and implement these regulatory measures within the next two years, with subsequent reviews and necessary updates every half-year or annually. Management is a continuous process and doesn't have a deadline as such. Though the need for workplaces to implement clear guidelines for practice can happen immediately and can be updated as new information about AI and regulations become available.

3.6 Evaluation of AI's contributions

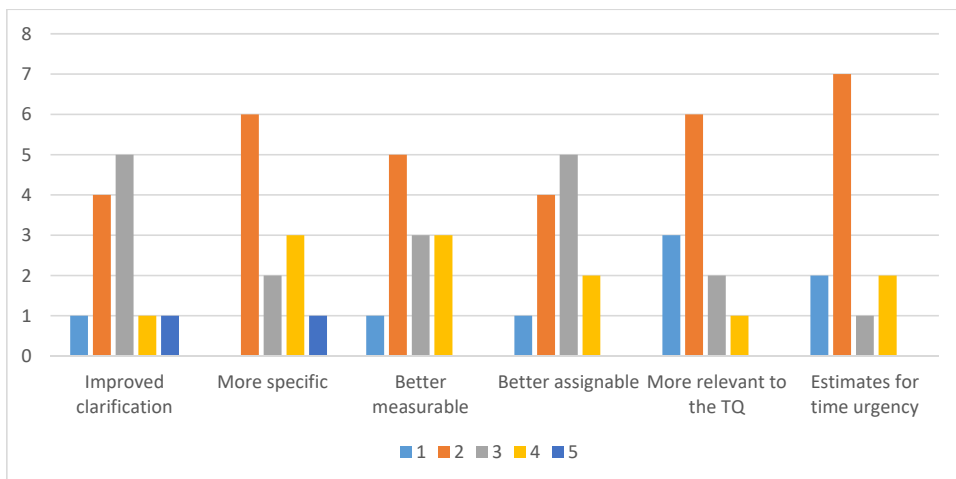
In Figure 8, we present the participants' evaluations regarding the utility of AI's contributions.

The distribution of responses suggests that opinions were split. Five out the 12 responders scored the AI assistant's contribution in improving their clarifications as either 'not at all' or 'somewhat' (blue and orange bars), another 5 scored it as 'a lot,' one as 'significantly' and one as 'enormously.' With regards to making it more *specific*, half responded 'not at all,' and five scored with as either 'significantly' or 'enormously.'

Regarding support for rendering it more *measurable*, eleven scored it high. Similarly, eleven out of twelve were satisfied with the assistant's help to make their clarification better assignable. The satisfaction scores were lower for the last two criteria, i.e., more relevant and more time-bounded.

In response to the last question, asking them to reflect on using AI assistants, seven responded that they found it helpful that the organisers run the ideas through the assistant and provided them with an edited-improved version of their contribution to review, correct, and approve. At the same time, five responded they would prefer to use AI assistance to edit and improve their clarification(s). One participant alerted that AI responses were easily recognisable by one of the industry's standard AI Detectors (<https://gptzero.me/>).

Figure 8 Evaluation results (see online version for colours)



4 Discussion

Our discussion is organised in sections. We begin with an introduction highlighting the key characteristics of deliberative democracy and how AI can assist in addressing the challenge of scalability. We then review some examples of AI agents that support human thinking, and continue with a critical comparison of AI-facilitated deliberation to traditional participatory models. We conclude with a discussion regarding our rationale for choosing to use SMART criteria in our experiments.

4.1 Deliberative democracy and the challenge of scale

Deliberative democracy focuses on active discussions among relevant stakeholders, grounding their positions on arguments, inclusivity, and public engagement as the foundation of collective decision-making (Dryzek, 2002; Parkinson and Mansbridge, 2012). Since the SDD methodology meets these criteria, and falls well into the repertoire of deliberative democracy tools, the subsequent discussion about deliberative democracy applies also to SDD. Unlike other models of democratic processes, which prioritise

voting and majoritarian decision-making, deliberative democracy focuses on dialogue and consensus-building. The legitimacy of any decisions reached is grounded on the quality of the arguments presented during the deliberation, instead of numerical superiority (Gutmann and Thompson, 2004). This model has been widely applied in various governance contexts, including citizens' assemblies, participatory budgeting, and public policy forums (Fung, 2006). Similarly, the SDD model has been applied in many contexts, including peace and conflict resolution (Laouris et al., 2009a, 2009b; Laouris et al., 2008a, 2008b); government and societal challenges (e.g., reforming the local authorities: Laouris and Michaelides, 2018, or the Wine Villages of Cyprus: Michaelides and Laouris, 2024); supporting pan-European groups of experts discover obstacles and design actions to improve access to broadband technologies for all but also for people facing accessibility challenges (Laouris and Christakis, 2007; Laouris et al., 2007; Laouris et al., 2008b; Laouris et al., 2017); discovering and collectively agreeing on research agenda priorities, thus influencing European Commission funding (Roe et al., 2011); reinventing democracy (Laouris and Romm, 2022b), and many others.

However, a persistent challenge in all models of deliberative democracy (including SDD; see Laouris and Midgley, 2025) is scalability. Ensuring that large and diverse populations can effectively engage in high-quality deliberation without compromising the depth of discussion or inclusiveness is a great challenge (Laouris et al., 2014, p.281; Parkinson and Mansbridge, 2012). In small-scale settings, deliberation fosters rational discourse and deep engagement, but when extended to larger-scales (e.g., national or global levels), issues such as cognitive overload, elite capture, and participant fatigue arise (Landemore, 2020). Traditional deliberative systems rely on face-to-face engagement, expert-driven facilitation, and structured dialogue formats, but these models struggle to accommodate large, dispersed, and digitally connected publics (Chambers, 2009).

Recent advances in digital and AI-driven facilitation offer new possibilities for addressing these limitations. By leveraging natural language processing, automated summarisation, and argument mapping, AI-enhanced deliberative platforms can expand the deliberative space, allowing a more representative public to contribute while maintaining deliberative quality (Noveck, 2018; Landemore, 2020). This shift aligns with the vision of open democracy, where governance structures are more responsive, decentralised, and participatory through technological augmentation (Landemore, 2020).

To summarise, the integration of AI into deliberative democracy represents a critical evolution in participatory governance. The following section explores how AI can enhance deliberative processes by addressing issues of scale, depth, and engagement.

4.2 Examples of AI agents supporting human thinking

Many authors have been experimenting with using an AI agent to support humans in considering more aspects of their thinking and thus making better decisions. For example, Kim et al. (2021) point out that deliberative discussion and consensus-reaching are challenging in online chats partly because of the lack of structure. They designed and developed an assistant called DebateBot, which structures discussions and encourages introverted participants to contribute. Their results suggest that a chatbot agent that partly structures the interactions improves the quality of discussions. In our experimental setup, the participants are supported to better structure their clarification highlighting how and why it is SMART. Lara (2021). experimented with a virtual assistant that, through

dialogue, neutrality and virtual reality technologies, could teach users how to make better moral decisions independently, thus increasing the individual's capacity to decide for themselves rather than directly influencing behaviour reflectively. Similarly, Ji et al. (2023) adopted a Socratic approach. They prompt users to critically evaluate and articulate their policy perspectives during the initial stages of their reflections. These authors managed to determine the AI agent's depth of reflection and potential persuasive abilities. Argyle et al. (2023) conducted large-scale experiments that demonstrated how online conversations about divisive topics can be improved with AI tools. They employed a large language model to make real-time, evidence-based recommendations to improve participants' perception of feeling understood. These interventions improved reported conversation quality and gave rise to democratic reciprocity. Even the tone of the conversations was improved without any systematic change in the content. Pinka (2021) disputes whether developing a truly ethical AI system is possible. He argues that, in most cases, it is presupposed that an artificially intelligent system can be ethical. The author explores the potential path to AI ethics by considering the role of imagination in the deliberative process via the work of John Dewey and his interpreters (Dewey, 1922; Fesmire, 2003; Johnson, 2020), showcasing one form of reinforcement learning that mimics imaginative deliberation.

The work of Zhang et al. (2023) has similarities with our line of work. While the usual approach is to use AI to support decision-making through machine learning or by automating or assisting human decisions, these authors used machine learning to help stakeholders surface ways to improve and make better judgments and decisions. Their web tool, 'deliberating with AI,' enables people to create and evaluate machine learning models to examine the strengths and shortcomings of past decision-making and deliberate on improving future decisions. They apply their tool in people selection, where decision makers (faculty) and decision subjects (students) use it to improve graduate school admission decisions. Another interesting approach was proposed by Chiang et al. (2024), who designed four different styles of devil's advocate bots powered by LLMs, varying their interactivity (i.e., interactive vs. non-interactive) and their target of objection (i.e., challenge the AI recommendation or the majority opinion within the group). They found evidence suggesting that LLM-powered devil's advocates that argue against the AI model's decision recommendation have the potential to promote groups' appropriate reliance on AI, hence leading to an increased level of decision accuracy.

4.3 AI-facilitated deliberation compared to traditional participatory models

Participatory governance has historically relied on structured deliberative processes to facilitate collective decision-making. The most widely used approaches include the Delphi, open method of coordination (OMC), and consensus conferences. Each of them utilises structured discussions to reach informed conclusions. Even though these methods are effective in small-to medium-scale deliberations, they face significant scalability, efficiency, and inclusivity challenges.

The Delphi method is an expert-driven process that refines opinions over multiple rounds of structured feedback (Linstone and Turoff, 1975; Rowe and Wright, 1999). It is effective for achieving consensus among experts. At the same time, it is constrained because participation is limited to only a few experts. AI-enhanced deliberation can not only extend but also democratise the Delphi process. For example, AI can facilitate large-scale, multi-stakeholder deliberations while keeping the structured feedback loops.

AI-driven sentiment and argument analysis can mitigate the dominance of influential voices, thus providing a more balanced synthesis of perspectives. Furthermore, AI-generated summaries and automated thematic clustering support stakeholders in reaching convergence on key issues faster, hence reducing the number of iterative cycles required. More importantly, an AI-enhanced Delphi process could support a globally distributed panel of experts in deliberating on specific policy proposals, incorporating diverse perspectives.

The OMC is used widely by the European Union. Instead of relying on legally binding directives, the OMC facilitates policy coordination among member states through iterative consultations, benchmarking, and peer review (Zeitlin, 2005; Sabel and Zeitlin, 2008). While OMC has been instrumental in fostering transnational policy learning, its effectiveness is often limited by bureaucratic procedures that lead to long response times, delaying decision-making. Also, the fact that consultations occur periodically provides limited adaptability to emerging issues. Finally, the level of engagement and the quality of the discussion may vary depending on priorities and capacities imposed by the government(s). AI could analyse policy deliberations across member states systematically and in real-time, identifying common trends, convergences, or divergences. Thus, AI could permit continuous rather than periodic consultation, therefore enabling dynamic policy refinement. As a result, AI-generated reports and real-time deliberation dashboards can ensure greater public visibility of ongoing discussions. For example, AI-assisted deliberation could facilitate cross-national comparisons of digital governance frameworks, enabling policymakers to refine regulations based on continuous, evidence-based feedback rather than periodic consultations.

Finally, consensus conferences are a participatory method in which a representative group of citizens engages in structured deliberations on scientific or technological issues before making policy recommendations (Joss and Durant, 1995; Fung, 2006). Consensus conferences boost democratic legitimacy. But they also suffer from various limitations. First, organising an in-person conference requires substantial financial and administrative resources. Second, the number of participants is often tiny and possibly self-selected, thus raising concerns about demographic inclusivity and bias. Third, securing engagement and a high-quality discussion is incredibly challenging as the number of participants rises. AI-powered deliberation, on the other hand, can alleviate some of these constraints. AI-facilitated platforms can integrate input from thousands of participants and maintain structured engagement. AI tools can extract, cluster, and synthesise common themes from participants' contributions and provide a comprehensive summary of perspectives. AI-driven systems can also tailor deliberative topics to participant expertise, hence ensuring more engagement. For instance, AI-supported citizen deliberations on AI ethics could enable large-scale, structured input from stakeholders across the globe, dynamically synthesised into actionable policy recommendations.

4.4 Rationale for using SMART criteria in AI-enhanced deliberation

Traditional deliberation processes, including the SDD methodology, often struggle with ambiguously formulated contributions, which can create interpretative challenges. Ambiguities in participant inputs may lead to misalignment between the intended meaning of an idea and how it is understood by others, potentially reducing the effectiveness of deliberative outcomes. An AI assistant can mitigate these challenges by

analysing semi-structured clarifications, detecting inconsistencies or vague formulations, and providing suggestions for refinement.

Research has shown that structured inputs improve consensus-building and decision quality by optimising the collective intelligence of deliberative participants (Sunstein, 2006; Rowe and Frewer, 2000). Compared to free-text discussions, where contributions can be imprecise, overly broad, or challenging to evaluate, the specific, measurable, assignable, relevant, and time-bounded (SMARTer) framing imposed through our AI-Assistant ensures that contributions are better defined and more actionable. The assistant provides feedback on the specificity and clarity of contributions, helping participants refine their clarifications by expressing their ideas in a way that allows for evaluation, prioritisation, and implementation. The structured format enables improved understanding, prevents deliberative drift (i.e., discussions becoming too abstract or vague), hence reducing the risk of misinterpretation.

The SDD process, when applied correctly in face-to-face implementations, enables a group of stakeholders to co-discover expert wisdom through the interrogative process of clarification. In face-to-face applications, inserting an AI agent could be technologically and conceptually challenging. Technologically, because the system must operate in real-time. In our experiment, the AI assistant was employed asynchronously to the idea clarification stage. However, even if the assistant can be used in real-time, one cannot ignore the cognitive overload introduced when participants start to divide their attention between what is being discussed in the room and what the assistant provides as feedback or recommendations. In our experiment, the AI is used asynchronously to support the idea clarification phase, allowing participants to refine their contributions at their own pace.

When asynchronous deliberating groups are asked to prompt each other with questions for clarification, the effort can become cognitively overwhelming both in terms of generating questions and in terms of responding to those questions. The AI assistant Prompts participants to draft ‘smarter’ clarifications, enhancing the quality of the dialogue. By improving the clarity and specificity of contributions upfront, AI assistance shortens the clarification stage.

We assert that AI-augmented deliberation represents a step forward in making participatory decision-making more scalable, structured, and impactful.

5 Limitations, conclusions, and further directions

While our study provides valuable insights, it might suffer several limitations, which are discussed below in connection with suggestions regarding future directions.

- Dependence on the Prompt’s syntax: the syntax of the Prompts submitted to an LLM can result in different responses. In our case, the Prompt was always the same and was embedded within the Python script. Nevertheless, further experimentation is needed to settle on optimal ways of interacting with LLM models. At least two weaknesses in our Prompts were identified: One, our Prompt did not include instruction to include the original (if it was deemed sufficiently clear) statement made by the author, or an improved version of it. That is why the AI included the original or edited clarification (e.g., in Example 2) and in others not (e.g., Examples 1, 3, 4). Second, our Prompt did not impose any constraints for meeting SMART

criteria. To avoid vagueness and generalities, the AI could be asked to reject or flag its own outputs before returning them to the participant if they are overly broad, ambiguous, or lack clear intent.

In future interactive versions, the assistant could engage in a Socratic-style conversation with the participant in ways analogous to Ji et al. (2023), but instead of the AI assistant posing the issue and asking questions, having the SDD participant submit (to the AI) her clarification and the AI engaging in a conversation by posing a clarifying question if the human input is too vague, or asking questions such as ‘To whom would you assign this ..?’, or ‘What would be a reasonable time horizon ..?’, or ‘How would you measure ...?’. By doing so, the AI assistant encourages the human to think more about her idea, but also receives useful feedback which will help it generate better and more specific recommendations.

- Subsequently, the AI assistant can encourage the participant to come up with specific measurable components such as percentages, frequency, or benchmarks. Regarding assignability, the AI could prompt users to define better who is responsible for implementing or who could prevent the proposed action (e.g., governments, the private sector, international organisations, etc.). In such an interactive implementation, the assistant could also prompt participants to specify a timeline for their proposed actions, offering some examples, e.g., ‘by 2030,’ ‘within 5 years,’ Or simply asking, ‘By when should this be implemented?’
- Usability may depend on subject matter: the subject matter (i.e., AI) upon which the research findings in this report are based may take better advantage of rich language on the topic in the LLM used by the AI platforms. This may limit applicability of the proposed system to domains where significant discussions of the subject have been previously constructed and the LLM has used them in its training.
- AIs have their own biases: ChatGPT’s reputation for moderation and censorship could pose challenges, particularly when engaging in an SDD on value-based issues. One should address the challenge of such biases in LLMs because they could skew the results. If all participants use the same LLM (as in our case here), they might all be subject to the same biases. One idea for future experimentation might be for each participant to use their own LLM, especially now that some (for example, ChatGPT) have memories of previous conversations. In any case, it is crucial to ensure that the human participant is aware of these risks and encouraged to handle the AI’s responses with a critical mind. Limiting the range of responses and introducing biases toward Western civilisation values can be dangerous; especially when human oversight is absent. Such distortion would undermine the fundamental purpose of SDD by stifling diverse voices. This becomes especially contentious when the objective is to explore all perspectives, seek common ground, and achieve consensus.
- Passive, non-critical evaluation of the AI’s response: Ma et al. (2024) argue that during AI-assisted decision-making, humans often review AI’s suggestions passively and decide whether to accept or reject them without engaging in critical or analytical thinking. This could have also happened in our experiments. Ma and colleagues propose a Human-AI Deliberation framework to promote human reflection and discussion on conflicting human opinions in decision-making. Based on theories in

human deliberation, their framework engages humans and AI in dimension-level opinion elicitation, deliberative debate, and decision updates. Their Deliberative AI (leveraging LLMs) enables flexible conversational interactions and faithful information provision, outperforming conventional explainable AI (XAI) assistants in improving humans' appropriate reliance and task performance.

- Participants remain skeptical: our results revealed that one-third of participants remained skeptical of AI's enhancements. Future experiments may attempt to uncover any underlying patterns in the concerns of those participants who remained unconvinced, such as, for example, lack of trust towards AI, unwillingness to let AI (or anyone else) assume some form of control in their thinking, and deliberation process, or any other bias they might have against SDD modifications.
- Interrogative instead of declarative AI engagement: the authors' subjective assessments regarding the AI's responses, were mixed. Some considered the responses simplistic, not making a notable difference in helping them to meet SMART criteria. One improvement, also mentioned above in section 'dependence on the Prompt's syntax,' is to engage with the AI in an interrogative, conversational manner. Humans generally prefer to be asked about what they mean rather than to be 'told' what they mean. Nonetheless, the experiment provided an example how AI can be engaged to enter into a dialogue with participants who themselves are part of an interpersonal dialogue.
- Contributions of AI-assisted participants approximate those of experts: the authors argue that the insertion of AI within the SDD process can empower people from all walks of life to act as experts. While the examples exhibit more elaboration or conciseness, the extended text offered by the AI to justify 'smartness,' in connection with the explicit restructuring to highlight each attribute in SMART, provides knowledge typically owned by experts.
- IPR questions: finally, as one participant noted, ChatGPT's responses are recognisable by standard AI detectors, and in some exceptional cases, they might raise IPR issues.
- Other areas where AI could be useful: one aspect where AI may prove specifically advantageous in the application of SDD is supporting a process for tracking the deliberative arguments presented by participants when they evaluate whether influences among pairs of ideas under discussion are 'significant'. This support can enhance reflection on results constructed in face-to-face SDD dialogue and may frame choices for improved choice-making in asynchronous dialogue.

To conclude, the authors assert that an AI assistant can help SDD participants reflect on their clarifications and make them more concise. When the meaning behind each contribution becomes more spherical, deeper, and 'smarter' early in the SDD process, the quality of the entire dialogue improves.

Acknowledgements

The authors would like to acknowledge the twelve participants of the AI-Repercussions SDD who were willing to participate in these experiments, and in particular Tania Gill, James Jamie Rose, and Kevin Dye for their constructive comments.

Declarations

On behalf of all authors, the corresponding author states that there are no conflicts of interest of any kind.

References

- Althaus, C., Ball, S., Bridgman, P., Davis, G. and Threlfall, D. (2022) *The Australian Policy Handbook: A Practical Guide to the Policymaking Process*, Routledge, <https://doi.org/10.4324/9781003351993>.
- Argyle, L.P., Bail, C.A., Busby, E.C., Gubler, J.R., Howe, T., Rytting, C., Sorensen, T. and Wingate, D. (2023) 'Leveraging AI for democratic discourse: chat interventions can improve online political conversations at scale', *Proceedings of the National Academy of Sciences*, Vol. 120, No. 41, p.e2311627120, <https://doi.org/10.1073/pnas.2311627120>.
- Bohm, D. (1987) *Unfolding Meaning: A Weekend of Dialogue with David Bohm*, Routledge, New York.
- Chambers, S. (2009) 'Rhetoric and the public sphere: has deliberative democracy abandoned mass democracy?', *Political Theory*, Vol. 37, No. 3, pp.323–350, <https://doi.org/10.1177/0090591709332336>.
- Chiang, C.W., Lu, Z., Li, Z. and Yin, M. (2024) 'Enhancing AI-assisted group decision making through LLM-powered Devil's advocate', in *Proceedings of the 29th International Conference on Intelligent User Interfaces*, March, pp.103–119, <https://doi.org/10.1145/3640543.3645199>.
- Christakis, A.N. and Bausch, K.C. (Eds.) (2006) *How People Harness Their Collective Wisdom and Power: to Construct the Future in Co-Laboratories Of Democracy*, Information Age Publishing, Leeds, UK.
- Dewey, J. (1922) *Human Nature and Conduct*, Henry Holt and Co., New York.
- Dryzek, J.S. (2002) *Deliberative Democracy and Beyond: Liberals, Critics, Contestations*, Oxford University Press, <https://doi.org/10.1093/019925043X.001.0001>.
- Dye, K.M. and Conaway, D.S. (1999) *Lessons Learned from Five Years of Application of the CogniScope™ Approach to the Food and Drug Administration*, CWA Report, CWA Ltd.
- Fesmire, S. (2003) *John Dewey and Moral Imagination: Pragmatism in Ethics*, Indiana University, Press, Indiana, USA.
- Flanagan, T. (2006) 'Co-laboratories of democracy: How people harness their collective wisdom to create the future', *Systems Research and Behavioral Science*, Vol. 23, No. 6, pp.845–847, <https://doi.org/10.1002/sres.767>.
- Flanagan, T.R. and Christakis, A.N. (2021) *The Talking Point: Creating an Environment for Exploring Complex Meaning*, 2nd ed., Information Age Publishing, Leeds, UK.
- Flood, R. (1988) 'Situational complexity, systems modelling and methodology', *Transactions of the Institute of Measurement and Control*, Vol. 10, No. 3, pp.122–129, <https://doi.org/10.1177/014233128801000303>.

- Fung, A. (2006) 'Varieties of participation in complex governance', *Public Administration Review*, Vol. 66, No. s1, pp.66–75, <https://doi.org/10.1111/j.1540-6210.2006.00667.x>.
- Gutmann, A. and Thompson, D. (2004) *Why Deliberative Democracy?* Princeton University Press, New Jersey, USA.
- Ji, Y., Wang, M. and Wang, S. (2023) *Building an AI Reflection Agent for Policy Deliberation* [online] <https://bpb-usw2.wpmucdn.com/voices.uchicago.edu/dist/9/3887/files/2024/02/FINALWangJiWangXLabSRF23-AIReflec8onAgentPolicyDelibera8on-1608fccc214321a8.pdf> (accessed 8 October 2025).
- Johnson, M. (2019) 'Dewey's radical conception of moral cognition', *The Oxford Handbook of Dewey*, p.175, <https://doi.org/10.1093/oxfordhb/9780190491192.013.6>.
- Joss, S. and Durant, J. (Eds.) (1995) *Public Participation in Science: The Role of Consensus Conferences in Europe*, NMSI Trading Ltd., London, UK.
- Kalliola, S. and Nakari, R. (Eds.) (1999) *Resources for Renewal: a Participatory Approach to the Modernization of Municipal Organizations in Finland*, Vol. 10, John Benjamins, <https://doi.org/10.1075/dowi.10>.
- Kim, S., Eun, J., Seering, J. and Lee, J. (2021) 'Moderator chatbot for deliberative discussion: Effects of discussion structure and discussant facilitation', *Proceedings of the ACM on Human-Computer Interaction*, (CSCWI), Vol. 5, pp.1–26, <https://doi.org/10.1145/3449161>.
- Landemore, H. (2020) *Open Democracy: Reinventing Popular Rule For The Twenty-First Century*, Princeton University Press, <https://doi.org/10.2307/j.ctv10crczs>.
- Laouris, Y. (2012) 'The ABCs of the science of structured dialogic design', *International Journal of Applied Systemic Studies*, Vol. 4, No. 4, pp.239–257, <https://doi.org/10.1504/IJASS.2012.052235>.
- Laouris, Y. (2015) 'Reengineering and reinventing both democracy and the concept of life in the digital era', in Floridi, L. (Ed.): *The Onlife Manifesto: Being Human in a Hyperconnected Era*, pp.125–142, Springer Nature, https://doi.org/10.1007/978-3-319-04093-6_7.
- Laouris, Y. (2022) 'Managing large-scale societal change', in Márquez, F.P.G. (Ed.): *Operations Management and Management Science*, IntechOpen, <https://doi.org/10.5772/intechopen.96833>.
- Laouris, Y. (2023a) 'Method to integrate asynchronously produced individual influence maps into an extrapolated population influence map following the face-to-face stage of a structured democratic dialogue', *Systems Research and Behavioral Science*, Vol. 40, No. 3, pp.437–450, <https://doi.org/10.1002/sres.2877>.
- Laouris, Y. (2023b) *Structured Dialogical Design Frameworks for Addressing Complexity in Systems Engineering*, Doctoral Dissertation, University of Portsmouth, University of Portsmouth Research Portal [online] <https://researchportal.port.ac.uk/en/studentTheses/structured-dialogical-design-frameworks-for-addressing-complexity> (accessed 8 October 2025).
- Laouris, Y. and Christakis, A.N. (2007) 'Harnessing collective wisdom at a fraction of the time using structured dialogic design process in a virtual communication context', *International Journal of Applied Systemic Studies*, Vol. 1, No. 2, pp.131–153, <https://doi.org/10.1504/IJASS.2007.015585>.
- Laouris, Y. and Dye, K. (2024) 'Multi-stakeholder structured dialogues: five generations of evolution of dialogic design', *Systems Research and Behavioral Science*, Vol. 41, No. 2, pp.368–389, <https://doi.org/10.1002/sres.2971>.
- Laouris, Y. and Michaelides, M. (2018) 'Structured democratic dialogue: an application of a mathematical problem structuring method to facilitate reforms with local authorities in Cyprus', *European Journal of Operational Research*, Vol. 268, No. 3, pp.918–931, <https://doi.org/10.1016/j.ejor.2017.04.039>.
- Laouris, Y. and Midgley, G. (2025) 'Contemporary and future challenges for structured democratic dialogue: a conversation between Laouris and Midgley', *Systems Research and Behavioral Science* (in press).

- Laouris, Y. and Michaelides, M. (2007) 'What obstacles prevent practical broad-band applications from being produced and exploited?', in Patrick, R. (Ed): *Towards an Inclusive Future Impact and Wider Potential of Information and Communication Technologies*, European Cooperation in Science and Technology, Chapter 7, pp.281–299 [online] <https://op.europa.eu/en/publication-detail/-/publica8on/c8297d4c-4f84-4651-8832-04ba3511a575>.
- Laouris, Y. and Romm, N.R. (2022a) 'Structured dialogical design as a problem structuring method illustrated in a re-invent democracy project', *European Journal of Operational Research*, Vol. 301, No. 3, pp.1072–1087, <https://doi.org/10.1016/j.ejor.2021.11.046>.
- Laouris, Y. and Romm, N.R. (2022b) 'African Youth's Visioning for re-inventing democracy in the digital era: a case of use of structured dialogical design', *World Futures*, Vol. 78, No. 1, pp.18–61, <https://doi.org/10.1080/02604027.2021.2014112>.
- Laouris, Y., Dye, K.M.C., Michaelides, M. and Christakis, A.N. (2014) 'Co-laboratories of democracy: best choices for designing sustainable futures', in Metcalf, G. (Ed.): *Social Systems and Design. Translational Systems Sciences*, Vol. 1, Springer, Tokyo https://doi.org/10.1007/978-4-431-54478-4_7.
- Laouris, Y., Emiliani, P.L. and Roe, P. (2017) 'Systemic evaluation of actions toward developing practical broadband applications for elderly and people with disabilities', *Universal Access in the Information Society*, Vol. 16, pp.247–255, <https://doi.org/10.1007/s10209-015-0441-0>.
- Laouris, Y., Laouri, R. and Christakis, A.N. (2008a) 'Communication praxis for ethical accountability: the ethics of the tree of action: dialogue and breaking down the wall in Cyprus', *Systems Research and Behavioral Science*, Vol. 25, No. 2, pp.331–348, <https://doi.org/10.1002/sres.890>.
- Laouris, Y., Michaelides, M. and Sapio, B. (2008b) 'A systemic evaluation of obstacles preventing the wider public benefiting from and participating in the broadband society', *Observatorio J.*, Vol. 5, pp.21–31 [online] <https://obs.obercom.pt/index.php/obs/article/view/199/165>, <https://doi.org/10.1002/1646-5954/ERC123483/2008021>.
- Laouris, Y., Erel, A., Michaelides, M., Damdelen, M., Taraszow, T., Dagli, I. and Christakis, A. (2009a) 'Exploring options for enhancement of social dialogue between the Turkish and Greek communities in Cyprus using the structured dialogic design process', *Systemic Practice and Action Research*, Vol. 22, No. 5, pp.361–381, <https://doi.org/10.1007/s11213-009-9134-z>.
- Laouris, Y., Michaelides, M., Damdelen, M., Laouri, R., Beyatli, D. and Christakis, A. (2009b) 'A systemic evaluation of the state of affairs following the negative outcome of the referendum in Cyprus using the structured dialogic design process', *Systemic Practice and Action Research*, Vol. 22, No. 1, pp.45–75, <https://doi.org/10.1007/s11213-008-9111-y>.
- Lara, F. (2021) 'Why a virtual assistant for moral enhancement when we could have a socrates?', *Science and Engineering Ethics*, Vol. 27, No. 4, p.42, <https://doi.org/10.1007/s11948-021-00318-5>.
- Linstone, H.A. and Turoff, M. (1975) *The Delphi Method: Techniques and Applications*, Addison-Wesley.
- Ma, S., Chen, Q., Wang, X., Zheng, C., Peng, Z., Yin, M. and Ma, X. (2024) 'Towards human-AI deliberation: Design and evaluation of LLM-empowered deliberative AI for AI-assisted decision-making', arXiv preprint arXiv:2403.16812, <https://doi.org/10.48550/arXiv.2403.16812>.
- Mansbridge, J., Bohman, J., Chambers, S., Christiano, T., Fung, A., Parkinson, J., ... and Warren, M.E. (2012) 'A systemic approach to deliberative democracy', in Parkinson, J. and Mansbridge, J. (Eds.): *Deliberative Systems: Deliberative Democracy at the Large Scale*, pp.1–26, Cambridge University Press, <https://doi.org/10.1017/cbo9781139178914.002>.
- Michaelides, M. and Laouris, Y. (2024) 'A cascading model of stakeholder engagement for large-scale regional development using structured dialogical design', *European Journal of Operational Research*, Vol. 315, No. 1, pp.307–323, <https://doi.org/10.1016/j.ejor.2023.11.050>.
- Noveck, B.S. (2018) *Solving Public Problems: A Practical Guide to Civic Innovation*, Yale University Press.

- Parkinson, J. and Mansbridge, J. (Eds.). (2012) *Deliberative Systems: Deliberative Democracy At The Large Scale*, Cambridge University Press, <https://doi.org/10.1080/19460171.2013.799894>.
- Pietroski, P. (2021) 'Chomsky on meaning and reference', in Allott, N., Lohndal, T. and Rey, G. (Eds.): *A Companion to Chomsky*, pp.404–415, <https://doi.org/10.1002/9781119598732.ch25>.
- Pinka, R. (2021) 'Synthetic deliberation: can emulated imagination enhance machine ethics?', *Minds and Machines*, Vol. 31, No. 1, pp.121–136, <https://doi.org/10.1007/s11023-020-09531-w>.
- Pruitt, B. and Thomas, P. (2007) *Democratic Dialogue: A Handbook for Practitioners*, International IDEA, Washington, DC [online] <https://www.idea.int/sites/default/files/publications/democratic-dialogue-a-handbook-for-practitioners.pdf>.
- Roe, P., Gill, J., Allen, B., Boyle, B., Heck, H., Shitta, G. and Laouris, Y. (2011) 'Towards a technology transfer roadmap from the coordination action in R&D in accessible and assistive ICT (CARDIAC)', *Technology and Disability*, Vol. 23, No. 3, pp.171–181, <https://doi.org/10.3233/TAD-2011-0325>.
- Rowe, G. and Frewer, L.J. (2000) 'Public participation methods: a framework for evaluation', *Science, Technology, and Human Values*, Vol. 25, No. 1, pp.3–29, <https://doi.org/10.1177/016224390002500101>.
- Rowe, G. and Wright, G. (1999) 'The Delphi technique as a forecasting tool: issues and analysis', *International Journal of Forecasting*, Vol. 15, No. 4, pp.353–375.
- Sabel, C.F. and Zeitlin, J. (2008) 'Learning from difference: the new architecture of experimentalist governance in the EU', *European Law Journal*, Vol. 14, No. 3, pp.271–327.
- Sunstein, C.R. (2006) *Infotopia: How Many Minds Produce Knowledge*, Oxford University Press.
- Todorov, T. (1984) *Mikhail Bakhtin: The Dialogical Principle*, Vol. 13, Manchester University Press.
- Tsivacou, I. (1997) 'The rationality of distinctions and the emergence of power: a critical systems perspective of power in organizations', *Systems Research and Behavioral Science*, Vol. 14, No. 1, pp.21–34, [https://doi.org/10.1002/\(SICI\)1099-1743\(199701/02\)14:1<21::AID-SRES117>3.0.CO;2-U](https://doi.org/10.1002/(SICI)1099-1743(199701/02)14:1<21::AID-SRES117>3.0.CO;2-U).
- Turrisi, P.A. (1997) *Pragmatism As A Principle And Method of Right Thinking: The 1903 Harvard Lectures on Pragmatism*, State University of New York Press, Albany.
- Vygotsky, L.S. (1978) *Mind in Society: the Development of Higher Psychological Processes*, Harvard University Press.
- Yeo, S., Lim, G., Gao, J., Zhang, W. and Perrault, S.T. (2024) 'Help me reflect: leveraging self-reflection interface nudges to enhance deliberativeness on online deliberation platforms', in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA, May, pp.1–32.
- Zeitlin, J. (2005) 'Social Europe and experimentalist governance: towards a new constitutional compromise?', *European Governance Papers (EUROGOV)*, No C-05-04 [online] https://edoc.vifapol.de/opus/volltexte/2011/2456/pdf/egp_connex_C_05_04.pdf.
- Zhang, A., Walker, O., Nguyen, K., Dai, J., Chen, A. and Lee, M.K. (2023) 'Deliberating with AI: improving decision-making for the future through participatory AI design and stakeholder deliberation', *Proceedings of the ACM on Human-Computer Interaction, CSCW1*, Vol. 7, pp.1–32, <https://doi.org/10.1145/3579601>.