

Opportunities for a Collaborative Approach with GenAI Agents in Public Sector Foresight Projects

Lukas Zumbrunn (Lukas.Zumbrunn@unisg.ch; University of St.Gallen & University of Oxford)

Nathan Davies (University of Oxford)

Working Paper for IRSPM Conference 2025, Bologna. **Please do not cite or distribute.**

Abstract

This paper explores the potential integration of Artificial Intelligence (AI) into public sector foresight projects as proxies for citizen participation. Despite motivations to increase participatory approaches, foresight projects have remained constrained by challenges pertaining to scale and confidentiality. Recent advancements in AI technology, particularly in agentic Large Language Models (LLMs), provide new opportunities to simulate diverse public perspectives while managing sensitive strategic information. Through a theory-building approach, we will investigate how AI can enhance openness in foresight without compromising confidentiality. Our conceptual framework integrates insights from foresight, deliberative democracy, and technological affordances literature, emphasizing task-technology-user fit. Preliminary findings indicate significant theoretical and practical opportunities, including enhanced efficiency and inclusivity, alongside challenges such as ethical considerations, data authenticity, and public trust. By analysing expert perspectives across foresight and deliberative democracy fields, the paper contributes to the digital public-citizen innovation discourse and proposes directions for future research on the collaborative potential of human-AI interactions in strategic public administration.

Introduction

In recent years, pressures on the public sector have increased to make strategy processes more open (Hansen et al., 2022). However, sensitive data, especially regarding core assumptions about the future and interactions with partners, reduces the possibility to open strategy fully to the public sphere. To mitigate such pressures and create room to bring the public in, if only through a synthesized version, there can be various approaches.

One such approach focuses on the capacities of Generative Artificial Intelligence (GenAI). Recent developments in GenAI enable technology to simulate interactions in various scenarios, ranging from pandemic simulations (Ibrahim, 2022) to strategic decisions in warfare (Hua et al., 2023). Developments in network learning effects in AI-powered tools make such interactions increasingly impactful (Marro et al., 2024). This raises questions on opportunities and challenges such interactions would bring for open strategy and foresight (Ortner et al., 2024). Against the backdrop of trust in GenAI tools by public servants, these questions are also relevant in the public sector context (Bright et al., 2024).

The goal of this paper is to investigate how an integration of AI-powered agents as a proxy for citizens' participation can bridge the gap between participatory foresight and the exclusive approach through. We engage in a theory-building approach combining strings of theory from technology affordances and urban planning to scope what the future could look like for an interaction between human experts and GenAI Agents in foresight projects. Essentially, we aim to highlight a possibility for public sector foresight to leapfrog other fields in the application of AI tools.

The paper aims to create a research framework for a collaborative approach between humans and GenAI in the public sector, therefore theoretically contributing to the panel and expanding the theory of digital public-citizen innovation (Mu & Wang, 2020).

To best do that, we first introduce the theories of foresight and participation in foresight. We then move to the structuring of participation in deliberative democracy as a foundation for comparison. Then, we introduce AI, LLMs and agentic LLMs, and why affordances matter. We then move to identify the research gap and how we want to approach it. Finally, we discuss implications, especially also concerning ethical questions.

Theory

Foresight

Theory defines foresight as a “distinct process of monitoring prospective oncoming events, analyzing potential implications, simulating alternative courses of action, asking unasked questions, and issuing timely warning to avert a risk or seize an opportunity” (Fuerth & Faber, 2012, p. 7). Furthermore, foresight is a continuous process (Hines & Gold, 2015). Other than general future studies, foresight has the explicit aim to influence strategies (Heo & Seo, 2021).

Issues covered by foresight are manifold. Examples include technological megatrends’ impact on societies (Andronikidis et al., 2025), national tourism (Kenzhebekov et al., 2021), or the common good (Szathmári et al., 2024). Foresight practices are employed for diverse purposes, including creating an early warning system, the prioritizing of resources in contingency planning, or the advancing of societal learning (Johnston & Calof, 2012).

Foresight is conducted using a variety of methods. The most well-known include scenario planning (Ramírez & Wilkinson, 2016), horizon scanning (Ansoff, 1975), and Delphi studies (Kishita, 2021). The decision on which method is used is context and case specific (Popper, 2008). Generally, in foresight, perspectives of multiple stakeholders are considered to understand the organizational environment, both on a transactional and a contextual level (Pouru-Mikkola et al., 2023). However, citizen participation in public sector foresight can vary.

Citizen Participation in Public Sector Foresight

In both literature and practice, foresight can both be conceptualized as an inclusive or an exclusive process. This relates to research on the dimensions of participation (Fung, 2006). The relevant axis of participation for the present paper is the inclusion of participants, irrespective of both their authority and the decision mode.

Prior research has shown a tendency to establish a link between participation and foresight. For example, when foresight is explicitly related to as a “[...] participatory [...] process aimed at present-day decisions and mobilising joint actions. [...] It brings together key agents of change and various sources of knowledge in order to develop strategic visions [...]” (Gavigan et al., 2001, p. 4). By including citizens in foresight projects, the aim can be less elitist and improve values such as openness and fairness, which are valued elements in public sector practices (Berze, 2014).

A concrete case of continuously high citizen participation in foresight is the project “Lithuania 2030”. With the institutionalization of a national day of ideas and the implementation of roundtables to discuss visions of the future, Lithuania embarked on a highly inclusive path of foresight (Paliokaitė & Sadauskaitė, 2023). Given the relative youth of the Lithuanian democracy, this process aims to consolidate various public perceptions.¹

The implementation is structured in three parts. First, policy continuity is ensured by integrating the foresight results into the existing strategic planning and government decision making system; then, success stories are shared to focus on mutual learning; and third, citizens and organizations were tasked to ensure the continuous input of societal needs (Paliokaitė & Sadauskaitė, 2023). There is not an expiry date to the project, but there is an annual process.

However, citizen participation should not be seen as a *conditio sine qua non*, as research has also identified dilemmas arising from participation in foresight projects. For example, the involved actors do not lose their intrinsic self-interest, or challenges of differing knowledge and backgrounds that can influence the direction of the discussions (van der Helm, 2007).

One foresight approach with lower citizen participation is – rather ironically – called “Public Open Foresight” (POF). POF is defined as “an inter-organizational discussion [...] process of (a) various public-sector organizations regarding future developments [...] and (b) in which decisions related to future strategies and innovations are taken collectively” (Schmidhuber & Wiener, 2017, p. 95). POF, however, excludes participation by citizens or non-organized stakeholder groups and focuses on only integrating people on the same hierarchical level (Schmidhuber & Wiener, 2017). There still is the aim to include differing perspectives, while potentially mitigating challenges of knowledge-differences.

Concrete cases of foresight projects that do not involve public participation include projects led by Ministries of foreign affairs. There, debates arise around the confidentiality of the contents of what is being discussed in the context of plausible futures, as well as the necessity of creating a safe space (Jütersonke & Munro, 2024). Such safe spaces allow for more open discussions, the discussion of different perspectives, and avoiding group think or social pressure (Ramírez & Wilkinson, 2016).

¹ See “Lithuania’s Progress Strategy—Lithuania 2030”

Participation in Deliberative Democracy

Public sector foresight and deliberative democracy have distinctive foci. The former is frequently focused on anticipating long-term challenges using a small number of stakeholders or experts, and the latter on structuring inclusive decision-making in the present. However, as some foresight practices increasingly seek to incorporate participatory engagement, the boundaries between these domains blur. Lithuania 2030, exemplifies this convergence, embedding deliberative mechanisms into strategic policy visioning (Paliokaitė & Sadauskaitė, 2023). This foresight initiative closely mirrors Ackerman and Fishkin's proposed, but never realized, 'Deliberation Day' for a nationwide, structured deliberation process before elections (J. Fishkin et al., 2025).

Deliberative democracy offers a framework for citizen participation that emphasises reasoned discussion and mutual justification as central to political legitimacy (Habermas, 1981; Cohen, 1989). Unlike traditional democratic models focused on aggregating preferences through voting, deliberative approaches foster structured dialogue where participants critically engage with diverse perspectives before decision-making (Bächtiger et al., 2018; Mansbridge, 1999). This paradigm has led to institutional innovations such as Citizens' Assemblies, which engage representative cross-sections of citizens in informed deliberation and collective recommendation formulation (McKinney, 2024).

Modern deliberative processes rely on careful participant selection through quasi-random sampling, ensuring representation across socio-demographic dimensions while mitigating self-selection bias (Gašiorowska, 2023; Ryan & Smith, 2014). The growth of these mini-publics across governance levels signals increasing efforts to restore democratic legitimacy, yet they face a fundamental challenge: scale. While effective in groups of dozens or hundreds, expanding to thousands or millions of citizens introduces significant coordination challenges (J. Fishkin et al., 2025). As participation increases, maintaining quality deliberation becomes increasingly resource-intensive, creating practical barriers to large-scale implementation.

Motivations for deliberative processes vary considerably. Some practitioners view deliberation primarily as a means of achieving more robust consensus, where reasoned discussion converges toward shared understanding (J. S. Fishkin, 2018). Others value it for systematically mapping disagreements, making underlying divisions more explicit (Lafont, 2019; Landemore, 2022). A third motivation centres on legitimacy, where the process itself confers democratic validity regardless of specific outcomes, enhancing public acceptance even among those who disagree with final decisions (Curato et al., 2023).

These challenges and varying motivations have led practitioners to explore technological solutions, with AI emerging as a promising but complex approach to addressing scale limitations while accommodating different deliberative aims.

LLMs and Agentic LLMs

AI encompasses various approaches to creating computational systems that perform tasks requiring human intelligence. Machine learning, a subset of AI focusing on algorithms that improve through experience, has evolved from rule-based expert systems to sophisticated statistical methods identifying patterns in complex datasets. Three key factors have driven recent AI advancements: increasing computational power, improved algorithms, and expanding datasets. The field has moved beyond traditional supervised learning to more advanced techniques, laying the groundwork for language-based and agentic systems.

LLMs emerged as transformative AI technology with the introduction of the Transformer architecture (Vaswani et al., 2017), which enabled parallel data processing and significantly enhanced efficiency. LLMs undergo a three-phase development: pre-training on vast datasets, fine-tuning for specific tasks, and inference to generate responses. The development of reinforcement learning with human feedback (RLHF) further refined outputs by aligning them with human preferences (Christiano et al., 2017). ChatGPT's 2022 release marked a turning point in accessibility, driving rapid adoption across academic and professional fields.

Despite these advancements, significant limitations remain. One influential critique characterised LLMs as “stochastic parrots” that produce superficially coherent text by mimicking patterns without genuine understanding, highlighting risks of environmental costs, inscrutability, and perpetuated biases (Bender et al., 2021). While developments in chain-of-thought prompting and recursive reasoning have enhanced problem-solving capabilities (Wei et al., 2022), fundamental concerns persist. These systems inherit and potentially amplify biases from training data, disproportionately reflecting dominant narratives while marginalising minority perspectives (Blodgett et al., 2020). They also suffer from “hallucination,” generating plausible-sounding but factually incorrect information, particularly problematic in high-stakes policy environments (Huang et al., 2025). Their black-box nature makes tracing outputs and identifying error sources difficult (Rudin, 2019), underscoring the need for structured oversight as AI integration in public sector decision-making increases.

Applications of LLMs

LLMs have been shown to offer potential solutions to automating aspects of deliberative democratic processes. Building on the use of machine learning in Polis, an open-source deliberation tool that clusters viewpoints, some have sought to use LLMs to address the persistent challenge of scaling deliberation. Notably, DeepMind's Habermas Machine has demonstrated the ability to mediate deliberation more effectively than human facilitators. By analysing participants' contributions and refining statements through successive iterations, the AI system increased agreement among discussants while preserving minority perspectives (Tessler et al., 2024). Deliberative democracy facilitators are currently experimenting with a range of AI-powered tools developed by startups and non-profits, such as Dembrane, Harmonica, DeliberAId, and the AI Objectives Institute's Talk to the City, to analyse multiple inputs (text, audio, and video). These tools have been used to map differences and build consensus in large-scale deliberation, as they attempt to identify priority concerns while preserving perspective diversity that might be otherwise lost through manual synthesis (McKinney, 2024).

However, AI-mediated deliberation introduces significant normative and practical concerns. Critics worry that AI systems risk overdetermining discussion outcomes, subtly shaping deliberation toward pre-determined consensus rather than allowing organic, open-ended debate (Landemore, 2022). Additionally, while AI can enhance inclusivity by scaling engagement beyond traditional mini-publics, it also introduces new forms of opacity, with "black boxes" making it difficult to trace how deliberative outputs are generated and whether biases in training data affect the representativeness of synthesized arguments (McKinney, 2024). Furthermore, AI-assisted deliberation may struggle with legitimacy concerns: If AI-generated statements gain undue influence, do they still reflect genuine citizen preferences or merely algorithmic approximations? Empirical studies suggest that AI-facilitated discussions can increase participant satisfaction and reduce polarization, but their long-term impact on public trust in deliberative democracy remains an open question (Fishkin et al., 2025).

Also in the public sector, LLMs and their applications are not completely alien. Current research has focused on several concrete applications, such as political learning (Taylor & Richey, 2024), public trust (Robles & Mallinson, 2023; Wang et al., 2024), or public servant's willingness to adopt AI (Barodi & Lalaoui, 2025) and follow AI recommendations (Selten et al., 2023). Also in the base level task of understanding documents to enhance productivity, there have been studies narrowing down on the public sector (Fitzpatrick et al., 2025).

In public sector foresight practices, there are only few instances in which LLMs have found application. One example explores pathways for LLMs to support public sector tasks that relate to anticipatory tasks (Allaham & Diakopoulos, 2024). Also in the practice of scenario methods, there have been studies exploring the potential to use AI (Trujillo-Cabezas, 2024). Still, however, both research and applications are rather underdeveloped in foresight as compared to other fields such as deliberative democracy.

Thus, we recognize the gap between experts in deliberative democracy, where applications of AI are already more widely used, and the experts in foresight, where such tools are just starting to find first points of application. Hence, as shown in Figure 1, we distinguish between the expected usefulness and the experienced usefulness of AI tools.




Figure 1. Differences in Understanding

Why an Agent and not just a Chatbot?

There is a further frontier emerging in the field of Computer Science: the integration of AI agents. AI agents are systems designed to perform tasks autonomously within an organization, acting as if they are agents capable of completing various complex tasks on demand (Bullock & Chen, 2024). These agents can be programmed to imitate and support human-like decision-making processes, utilizing data analysis to inform their actions and responses (Johnson et al., 2022; Xiao et al., 2023). They can contribute to discussions and even add deviating opinions or perspectives (Baltaji et al., 2024). Ultimately, AI agents aim to enhance the efficiency and effectiveness of public administration by automating routine tasks and providing valuable insights for decision-making.

In this capacity, AI agents can be applied in situations that focus on participation. They can be integrated into strategic decision making situations, where they can provide novel points with the goal of improving the final outcome (Houde et al., 2025). Further, they can simulate different personas, thus showing different positions on matters (Sun et al., 2024). In this, they can act as proxies for the different simulated groups. This is also the case for LLMs, which adds the possibility for an output that is easily usable for humans (Salminen et al., 2020). Such personas can analyse values of the groups they are standing in for, and this has received a further push through recent developments in LLMs (Choi et al., 2025). This can be used to see and model macro-social trends (Zeng et al., 2025). Previous research has shown promising results in participatory urban planning (Zhou et al., 2024) and regeneration (Quan & Lee, 2025).

Affordances of LLMs as Proxy

The concept of affordances focuses on the perceived capabilities of a technology that support the users, or as research puts it: “Affordances imply the complementarity of the acting organism and the acted-upon environment. Most fundamentally, affordances are properties of the world that make possible some action to an organism equipped to act in certain ways.” (Gaver, 1991).

This work has tended to only consider, perceived affordances, as traditionally discussed in design and human-computer interaction (Norman, 1999), referring to the set of actions that users believe a technology enables based on its interface, design cues, and prior experiences. Others have put forward a complementary notion of “imagined affordances” (Nagy & Neff, 2015), expanding beyond user perceptions to also include the ways that expectations, cultural narratives and design intentions shape technological engagement. According to this view, users’ expectations, whether accurate or not, shape how tools are used. Affordances, further, have previously been explored in the context of public governance (Mergel, 2024), and also AI has been analysed from an affordances perspective (Park, 2024).

For us, the concept of affordances is important, as it is intended to go beyond the acceptance of a certain technology into the actual use and usefulness thereof. We emphasize the importance of the user, the task and the context in which a technology is used.

Geography and urban studies have coined the terminology of Planning Support Systems (PSS). These are tools used by planners especially in the context of urban development to analyse demands, allocate resources and thereby deal with the inherent uncertainty and complexity presented by the future. At their core, such tools should present relevant information for the planning process (Klosterman, 1997).

Such systems have also been analysed for their usefulness, with a focus first on the fit of the task and technology (Pelzer, 2017), followed by a further integration of the user and the context (Jiang et al., 2020). In this framework, the “task-technology fit indicates to the appropriateness of the technology to handle the task at hand; while the user-technology fit indicates to the goodness of fit between the capabilities of the user and the functionalities offered by the technology” (Jiang et al., 2020).




Figure 2: Framework of PSS usefulness by Jiang et al. (2020)

There are remarkable similarities between core assumptions held in PSS and in foresight. For PSS, used in urban strategy making, there is a notable description:

“[I]t involves actions taken by some to affect the use of land controlled by others, following **decisions taken by third parties** based on **values not shared by all concerned**, regarding **issues no one fully comprehends**, in an attempt to guide events and processes that very likely will **not unfold in the time, place, and manner anticipated**” (Couclelis, 2005, p. 1355)

As such, this is a useful framework for us to borrow from PSS given its core similarities. We adopt the PSS framework and expand it to reflect both the two-sidedness as well as the differences between expectation and experience. This further developed model is shown in Figure 3.




Figure 3: Expanded Research Model

Research Gap and Method

Our model allows us to ask several research questions. We have identified the core questions to be in the Utility and Usability frames within the framework. We, thus, come to the following proposed future research questions (PFRQ) that we hope to answer:

PFRQ₁: How do public sector foresight people imagine the affordances of AI Agents as proxies in foresight projects?

PFRQ₂: How do these imagined affordances compare to currently perceived affordances of AI aggregators as proxies in Deliberative Democracy?

PFRQ₃: How do the perceived and imagined affordances vary in Deliberative Democracy?

PFRQ₄: What are implications of the results of PFRQ3 on the future alignment and development of AI Agents in Foresight?

PFRQ₅: How do the respective contexts influence the usefulness of AI Agents?

The research is conducted using expert interviews. The target groups are twofold, to cover both the expectations of the experts in foresight and the experience and expectations of the experts in deliberative democracy.

We have had initial discussions with experts and academics and also attended conferences and events in the relevant fields. Thereby, we have identified the relevant experts using convenience sampling, with the authors complimenting each other's expertise and having each grown a network in the respective field. Hence, by conducting surveys, we aim to gain the necessary insights and identify potentially further people in the network through the snowballing method. The organizations and fields of expertise of our network are outlined in Table 1:

Table 1: Experts to be contacted for interviews

Organization	Expertise
Federal Department of Foreign Affairs, CH	Foresight (non-participatory)
Federal Department of Defence, CH	Foresight (non-participatory)
Cantonal Department of Strategy, CH	Foresight (mixed participation)
Prime Minister's Office, SG	Foresight (non-participatory)
OECD	Foresight (participatory)
International Atomic Energy Agency	Foresight (non-participatory)
University of Oxford	Foresight and Deliberative Democracy
AI Deliberation Startup	Deliberative Democracy
AI Deliberation Startup	Deliberative Democracy
Meta	Deliberative Democracy
Deliberative Democracy Practitioner	Deliberative Democracy
Deliberative Democracy Practitioner	Deliberative Democracy

While partly explained by convenience sampling of expert groups, the variety of our participants also has the benefit of including a diversity of perspectives within foresight, with varying degrees of participation, and also within deliberative democracy, with a range of direct experience with AI. We acknowledge the differences within foresight being located at different levels. This is due to our initial scoping, where we keep the option to follow up about variation in subsequent research.

Discussion and Expected Contributions

There is one key distinction between the applications in deliberative democracy and foresight, leading to a situation in which some limitations do not apply. The focus of people working in deliberative democracy is on ensuring that the voices of the people are not flattened and are replicated as close as possible, as anything else could be seen as antidemocratic or against the core purpose of deliberative democracy as an ideal type. This is drastically different for certain foresight engagement, especially the ones actively excluding public participation. There, a flattening of perspectives can be seen as acceptable, as any “voice of the public” adds to the polyvocality in the project.

Additionally, hallucinations are not a negative feature of AI when working in foresight, depending on the framing. Of course, when the focus of the tool is on staying as close to the dataset as possible, this is inherently negative. However, foresight per definition is considering futures that are outside of our current imagination. Hence, some practitioners have called for the consideration of hallucination as a potential upside to break with fixed images of the future (Polchar & Karle, 2023). These contextual differences open an opportunity for practitioners of foresight to leapfrog the current work of practitioners of deliberative democracy from a technological perspective.

However, there are several limitations that must be considered. The first is that it is not necessarily the case that the public trusts uses of AI or wants the public sector to use such tools. There is already research on when the public trusts AI applications in the public sector (Kleizen et al., 2023; Wang et al., 2024). Issues that apply to the future may add further scrutiny due to the heightened uncertainty presented by such topics.

A second limitation lies in the willingness of the users, concretely public servants, to apply the tools in their work (Ahn & Chen, 2022). Usually, the tools used are still voluntary for the individual employee. The success, however, also depends heavily on the active implementation in the routines. Especially given the comparatively higher scrutiny of the public sector, previous research would attribute a lower risk tolerance in public organizations (Callender, 2001), thus reducing the willingness to experiment with novel tools. Therefore, we assume that it is important to explore the side of public servants and their imagined benefits of using agentic AI to support strategic foresight work.

Finally, there is a limitation that is specific to the public sector and may not apply directly to practices of deliberative democracy. Given the high societal relevance of foresight projects in the public sector, a higher ethical scrutiny may be necessary to weight against the perceived usefulness of the agents in the decision making. Despite only limited direct impact on policies, foresight projects and their results are seen as an important information source for policymakers (Yoda, 2011). Thus, when algorithms, the foundation of agents, are called “moral zombies” (Véliz, 2021), this should not be neglected. There are high stakes associated with certain issues considered in public sector foresight, thus requiring a heightened level of attention.

Additionally, we understand the importance in distinguishing between the differences in data that can be used for foresight and for deliberative democracy. While it is impertinent for data in deliberative democracy to be *pure*, meaning that it is not crossed with artificial data but reflecting the true views of the members of the data set, this is different for the necessary data set in foresight. For questions that are debated in exclusion of the public, the data has to be bootstrapped and generated. This means that the dataset is used to train and generate new but similar data, thereby introducing a *pollution*. We will, in our future research, put a key emphasis on this challenge.

This work is supposed to be an outline for the further development of research papers targeting:

1. Scoping the use of AI Agents in foresight practices in the public sector
2. Comparing the use cases between the practices of deliberative democracy and the public sector
3. Experimental approach with public sector foresight practitioners, engaging them in workshops with existing AI tools used in deliberative democracy

The various research questions should be sequentially answered and put into the context of the respective literature. By presenting this to IRSPM, we hope to gather important field specific feedback and see how the understanding in our community is for the continuation of such research approaches. We greatly appreciate your insights.

References

- Ahn, M. J., & Chen, Y.-C. (2022). Digital transformation toward AI-augmented public administration: The perception of government employees and the willingness to use AI in government. *Government Information Quarterly*, 39(2), 101664. <https://doi.org/10.1016/j.giq.2021.101664>
- Allaham, M., & Diakopoulos, N. (2024). Supporting Anticipatory Governance using LLMs: Evaluating and Aligning Large Language Models with the News Media to Anticipate the Negative Impacts of AI. *arXiv Preprint arXiv:2401.18028*.
- Andronikidis, A., Kouskoura, A., Kalliontzi, E., & Bakouros, I. (2025). Foresight study for addressing megatrends in information and communication technology (ICT). *Journal of Innovation and Entrepreneurship*, 14(1), 5. <https://doi.org/10.1186/s13731-025-00466-z>
- Ansoff, H. I. (1975). Managing Strategic Surprise by Response to Weak Signals. *California Management Review*, 18(2), 21–33. <https://doi.org/10.2307/41164635>
- Bächtiger, A., Dryzek, J. S., Mansbridge, J., & Warren, M. (2018). Deliberative Democracy: An Introduction. In A. Bächtiger, J. S. Dryzek, J. Mansbridge, & M. Warren (Eds.), *The Oxford Handbook of Deliberative Democracy* (p. 0). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780198747369.013.50>
- Baltaji, R., Hemmatian, B., & Varshney, L. (2024). Conformity, Confabulation, and Impersonation: Persona Inconstancy in Multi-Agent LLM Collaboration. *Proceedings of the 2nd Workshop on Cross-Cultural Considerations in NLP*, 17–31. <https://doi.org/10.18653/v1/2024.c3nlp-1.2>
- Barodi, M., & Lalaoui, S. (2025). Civil servants' readiness for AI adoption: The role of change management in Morocco's public sector. *Problems and Perspectives in Management*, 23(1), 63–75. [https://doi.org/10.21511/ppm.23\(1\).2025.05](https://doi.org/10.21511/ppm.23(1).2025.05)
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. <https://doi.org/10.1145/3442188.3445922>
- Berze, O. (2014). *Mapping Foresight Practices Worldwide*.

- Blodgett, S. L., Barocas, S., III, H. D., & Wallach, H. (2020). *Language (Technology) is Power: A Critical Survey of “Bias” in NLP* (No. arXiv:2005.14050). arXiv. <https://doi.org/10.48550/arXiv.2005.14050>
- Bright, J., Enock, F. E., Esnaashari, S., Francis, J., Hashem, Y., & Morgan, D. (2024). Generative AI is already widespread in the public sector. *arXiv Preprint arXiv:2401.01291*.
- Bullock, J. B., & Chen, Y.-C. (2024). The brave new world of AI: Implications for public sector agents, organisations, and governance. *Asia Pacific Journal of Public Administration*, 46(4), 321–325. <https://doi.org/10.1080/23276665.2024.2356540>
- Callender, G. (2001). Public Sector Organizations. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social & Behavioral Sciences* (pp. 12581–12585). Pergamon. <https://doi.org/10.1016/B0-08-043076-7/04253-4>
- Choi, Y., Kang, E. J., Choi, S., Lee, M. K., & Kim, J. (2025). *Proxona: Supporting Creators’ Sensemaking and Ideation with LLM-Powered Audience Personas* (No. arXiv:2408.10937). arXiv. <https://doi.org/10.48550/arXiv.2408.10937>
- Christiano, P., Leike, J., Brown, T. B., Martic, M., Legg, S., & Amodei, D. (2017). *Deep reinforcement learning from human preferences*. <https://arxiv.org/abs/1706.03741v4>
- Couclelis, H. (2005). “Where has the Future Gone?” Rethinking the Role of Integrated Land-Use Models in Spatial Planning. *Environment and Planning A: Economy and Space*, 37(8), 1353–1371. <https://doi.org/10.1068/a3785>
- Curato, N., Chalaye, P., Conway-Lamb, W., De Pryck, K., Elstub, S., Morán, A., Oppold, D., Romero, J., Ross, M., Sanchez, E., Sari, N., Stasiak, D., Tilikete, S., Veloso, L., von Schneidemesser, D., & Werner, H. (2023). *Global Assembly on the Climate and Ecological Crisis Evaluation Report*. University of Canberra.
- Fishkin, J., Bolotnyy, V., Lerner, J., Siu, A., & Bradburn, N. (2025). Scaling Dialogue for Democracy: Can Automated Deliberation Create More Deliberative Voters? *Perspectives on Politics*, 1–18. <https://doi.org/10.1017/S1537592724001749>
- Fishkin, J. S. (2018). *Democracy When the People Are Thinking: Revitalizing Our Politics Through Public Deliberation* (J. S. Fishkin, Ed.). Oxford University Press. <https://doi.org/10.1093/oso/9780198820291.003.0001>

- Fitzpatrick, T., Kelly, S., Carey, P., Walsh, D., & Nugent, R. (2025). *Assessing Generative AI value in a public sector context: Evidence from a field experiment* (No. arXiv:2502.09479). arXiv. <https://doi.org/10.48550/arXiv.2502.09479>
- Fuerth, L. S., & Faber, E. M. H. (2012). *Anticipatory Governance Practical Upgrades: Equipping the Executive Branch to Cope with Increasing Speed and Complexity of Major Challenges*. National Defense University, Institute for National Strategic Studies.
- Fung, A. (2006). Varieties of participation in complex governance. *Public Administration Review*, 66, 66–75. <https://doi.org/10.1111/j.1540-6210.2006.00667.x>
- Gąsiorowska, A. (2023). Sortition and its Principles: Evaluation of the Selection Processes of Citizens' Assemblies. *Journal of Deliberative Democracy*, 19(1), Article 1. <https://doi.org/10.16997/jdd.1310>
- Gaver, W. W. (1991). Technology affordances. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Reaching through Technology - CHI '91*, 79–84. <https://doi.org/10.1145/108844.108856>
- Gavigan, J., Scapolo, F., Keenan, M., Miles, I., Farhi, F., Lecoq, D., Capriati, M., & Di Bartolomeo, T. (2001). FOREN (Foresight for Regional Development Network) a practical guide to regional Foresight. *Publications of the European Communities*.
- Hansen, J. R., Pop, M., Skov, M. B., & George, B. (2022). A review of open strategy: Bridging strategy and public management research. *Public Management Review*, 26(3), 678–700. <https://doi.org/10.1080/14719037.2022.2116091>
- Heo, K., & Seo, Y. (2021). Anticipatory governance for newcomers: Lessons learned from the UK, the Netherlands, Finland, and Korea. *European Journal of Futures Research*, 9(1). <https://doi.org/10.1186/s40309-021-00179-y>
- Hines, A., & Gold, J. (2015). An organizational futurist role for integrating foresight into corporations. *Technological Forecasting and Social Change*, 101, 99–111. <https://doi.org/10.1016/j.techfore.2014.04.003>
- Houde, S., Brimijoin, K., Muller, M., Ross, S. I., Moran, D. A. S., Gonzalez, G. E., Kunde, S., Foreman, M. A., & Weisz, J. D. (2025). *Controlling AI Agent Participation in Group Conversations: A Human-Centered Approach*. <https://doi.org/10.1145/3708359.3712089>

- Hua, W., Fan, L., Li, L., Mei, K., Ji, J., Ge, Y., Hemphill, L., & Zhang, Y. (2023). War and peace (waragent): Large language model-based multi-agent simulation of world wars. *arXiv Preprint arXiv:2311.17227*.
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., & Liu, T. (2025). A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Transactions on Information Systems*, 43(2), 1–55. <https://doi.org/10.1145/3703155>
- Ibrahim, S. (2022). *A review of agent-based model simulation for covid 19 spread*. 585–602.
- Jiang, H., Geertman, S., & Witte, P. (2020). Ignorance is bliss? An empirical analysis of the determinants of PSS usefulness in practice. *Computers, Environment and Urban Systems*, 83, 101505. <https://doi.org/10.1016/j.compenvurbsys.2020.101505>
- Johnson, M., Albizri, A., Harfouche, A., & Fosso-Wamba, S. (2022). Integrating human knowledge into artificial intelligence for complex and ill-structured problems: Informed artificial intelligence. *International Journal of Information Management*, 64, 102479. <https://doi.org/10.1016/j.ijinfomgt.2022.102479>
- Johnston, R., & Calof, J. L. (2012). Developing the capacity to assess the impact of foresight. *Foresight*, 14(1), 56–68. <https://doi.org/10.1108/14636681211210369>
- Jütersonke, O., & Munro, E. (2024). *Developing Anticipatory Governance Capacities in Ministries of Foreign Affairs* (Strategic Security Analysis). Geneva Centre for Security Policy.
- Kenzhebekov, N., Zhailauov, Y., Velinov, E., Petrenko, Y., & Denisov, I. (2021). Foresight of tourism in Kazakhstan: Experience economy. *Information*, 12(3), 138.
- Kishita, Y. (2021). Foresight and Roadmapping Methodology: Trends and Outlook. *Foresight and STI Governance*, 15(2), 5–11. <https://doi.org/10.17323/2500-2597.2021.2.5.11>
- Kleizen, B., Van Dooren, W., Verhoest, K., & Tan, E. (2023). Do citizens trust trustworthy artificial intelligence? Experimental evidence on the limits of ethical AI measures in government. *Government Information Quarterly*, 40(4), 101834. <https://doi.org/10.1016/j.giq.2023.101834>
- Klosterman, R. E. (1997). Planning Support Systems: A New Perspective on Computer-Aided Planning. *Journal of Planning Education and Research*, 17(1), 45–54. <https://doi.org/10.1177/0739456X9701700105>

- Lafont, C. (2019). *Democracy without Shortcuts: A Participatory Conception of Deliberative Democracy*. Oxford University Press.
- Landemore, H. (2022). *Can AI bring deliberative democracy to the masses?* [Working Paper]. <https://www.law.nyu.edu/sites/default/files/Helen%20Landemore%20Can%20AI%20bring%20deliberative%20democracy%20to%20the%20masses.pdf>
- Mansbridge, J. J. (1999). Everyday Talk in the Deliberative System. In S. Macedo (Ed.), *Deliberative Politics: Essays on Democracy and Disagreement*, (pp. 211–242). Oxford University Press.
- Marro, S., La Malfa, E., Wright, J., Li, G., Shadbolt, N., Wooldridge, M., & Torr, P. (2024). A Scalable Communication Protocol for Networks of Large Language Models. *arXiv Preprint arXiv:2410.11905*.
- McKinney, S. (2024). Integrating Artificial Intelligence into Citizens' Assemblies: Benefits, Concerns and Future Pathways. *Journal of Deliberative Democracy*, 20(1). <https://doi.org/10.16997/jdd.1556>
- Mergel, I. (2024). Social affordances of agile governance. *Public Administration Review*, 84(5), 932–947. <https://doi.org/10.1111/puar.13787>
- Mu, R., & Wang, H. (2020). A systematic literature review of open innovation in the public sector: Comparing barriers and governance strategies of digital and non-digital open innovation. *Public Management Review*, 24(4), 489–511. <https://doi.org/10.1080/14719037.2020.1838787>
- Nagy, P., & Neff, G. (2015). Imagined Affordance: Reconstructing a Keyword for Communication Theory. *Social Media + Society*, 1(2), 205630511560338. <https://doi.org/10.1177/2056305115603385>
- Norman, D. A. (1999). Affordance, conventions, and design. *Interactions*, 6(3), 38–43. <https://doi.org/10.1145/301153.301168>
- Ortner, T., Hautz, J., Stadler, C., & Matzler, K. (2024). Open strategy and digital transformation: A framework and future research agenda. *International Journal of Management Reviews*. <https://doi.org/10.1111/ijmr.12379>
- Paliokaitė, A., & Sadauskaitė, A. (2023). Institutionalisation of participative and collaborative governance: Case studies of Lithuania 2030 and Finland 2030. *Futures*, 150, 103174. <https://doi.org/10.1016/j.futures.2023.103174>

- Park, H. E. (Grace). (2024). The double-edged sword of generative artificial intelligence in digitalization: An affordances and constraints perspective. *Psychology & Marketing*, 41(11), 2924–2941. <https://doi.org/10.1002/mar.22094>
- Pelzer, P. (2017). Usefulness of planning support systems: A conceptual framework and an empirical illustration. *Transportation Research Part A: Policy and Practice*, 104, 84–95. <https://doi.org/10.1016/j.tra.2016.06.019>
- Polchar, J., & Karle, J. P., Claire. (2023, December 4). *Instant Scenarios: How to fast-forward your foresight and skip straight to strategy - Observatory of Public Sector Innovation*. <https://oecd-opsi.org/blog/instant-scenarios/>, <https://oecd-opsi.org/blog/instant-scenarios/>
- Popper, R. (2008). How are foresight methods selected? *Foresight*, 10(6), 62–89. <https://doi.org/10.1108/14636680810918586>
- Pouru-Mikkola, L., Minkkinen, M., Malho, M., & Neuvonen, A. (2023). Exploring knowledge creation, capabilities, and relations in a distributed policy foresight system: Case Finland. *Technological Forecasting and Social Change*, 186. <https://doi.org/10.1016/j.techfore.2022.122190>
- Quan, S. J., & Lee, S. (2025). Enhancing participatory planning with ChatGPT-assisted planning support systems: A hypothetical case study in Seoul. *International Journal of Urban Sciences*, 29(1), 89–122. <https://doi.org/10.1080/12265934.2025.2462823>
- Ramírez, R., & Wilkinson, A. (2016). *Strategic reframing: The Oxford scenario planning approach* (First edition.). Oxford University Press.
- Robles, P., & Mallinson, D. J. (2023). Artificial intelligence technology, public trust, and effective governance. *Review of Policy Research*. <https://doi.org/10.1111/ropr.12555>
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- Ryan, M., & Smith, G. (2014). Defining Mini-Publics. In K. Grönlund, A. Bächtiger, & M. Setälä, *Deliberative mini-publics: Involving citizens in the democratic process* (pp. 9–24). ECPR Press.

- Salminen, J., Guan, K., Jung, S.-G., Chowdhury, S. A., & Jansen, B. J. (2020). A Literature Review of Quantitative Persona Creation. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–14. <https://doi.org/10.1145/3313831.3376502>
- Schmidhuber, L., & Wiener, M. (2017). Aiming for a sustainable future: Conceptualizing public open foresight. *Public Management Review*, 20(1), 82–107. <https://doi.org/10.1080/14719037.2017.1293145>
- Selten, F., Robeer, M., & Grimmelikhuijsen, S. (2023). ‘Just like I thought’: Street-level bureaucrats trust AI recommendations if they confirm their professional judgment. *Public Administration Review*, 83(2), 263–278. <https://doi.org/10.1111/puar.13602>
- Sun, G., Zhan, X., & Such, J. (2024). Building Better AI Agents: A Provocation on the Utilisation of Persona in LLM-based Conversational Agents. *ACM Conversational User Interfaces 2024*, 1–6. <https://doi.org/10.1145/3640794.3665887>
- Szathmári, A., Köves, A., & Gáspár, J. (2024). Human-centred decision support for the common good: A combination of participatory foresight methods. *Journal of Decision Systems*, 1–15. <https://doi.org/10.1080/12460125.2024.2345946>
- Taylor, J. B., & Richey, S. (2024). AI chatbots and political learning. *Journal of Information Technology & Politics*, 1–11. <https://doi.org/10.1080/19331681.2024.2422929>
- Tessler, M. H., Bakker, M. A., Jarrett, D., Sheahan, H., Chadwick, M. J., Koster, R., Evans, G., Campbell-Gillingham, L., Collins, T., Parkes, D. C., Botvinick, M., & Summerfield, C. (2024). AI can help humans find common ground in democratic deliberation. *Science*, 386(6719), eadq2852. <https://doi.org/10.1126/science.adq2852>
- Trujillo-Cabezas, R. (2024). Exploring the link between foresight and artificial intelligence methods to strengthen collective future-building in contexts of social instability. *Foresight, ahead-of-print*(ahead-of-print). <https://doi.org/10.1108/FS-11-2023-0231>
- van der Helm, R. (2007). Ten insolvable dilemmas of participation and why foresight has to deal with them. *Foresight*, 9(3), 3–17. <https://doi.org/10.1108/14636680710754138>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). *Attention Is All You Need* (No. arXiv:1706.03762). arXiv. <https://doi.org/10.48550/arXiv.1706.03762>
- Véliz, C. (2021). Moral zombies: Why algorithms are not moral agents. *AI & SOCIETY*, 36(2), 487–497. <https://doi.org/10.1007/s00146-021-01189-x>

- Wang, Y.-F., Chen, Y.-C., Chien, S.-Y., & Wang, P.-J. (2024). Citizens' trust in AI-enabled government systems. *Information Polity*, 29, 293–312. <https://doi.org/10.3233/IP-230065>
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2022). *Chain-of-Thought Prompting Elicits Reasoning in Large Language Models* (No. arXiv:2201.11903). arXiv. <https://doi.org/10.48550/arXiv.2201.11903>
- Xiao, B., Yin, Z., & Shan, Z. (2023). *Simulating Public Administration Crisis: A Novel Generative Agent-Based Simulation System to Lower Technology Barriers in Social Science Research* (No. arXiv:2311.06957). arXiv. <https://doi.org/10.48550/arXiv.2311.06957>
- Yoda, T. (2011). Perceptions of domain experts on impact of foresight on policy making: The case of Japan. *Technological Forecasting and Social Change*, 78(3), 431–447. <https://doi.org/10.1016/j.techfore.2010.08.005>
- Zeng, F., Blank, G., & Schroeder, R. (2025). Using AI to model future societal instability. *Futures*, 166, 103543. <https://doi.org/10.1016/j.futures.2025.103543>
- Zhou, Z., Lin, Y., Jin, D., & Li, Y. (2024). *Large Language Model for Participatory Urban Planning* (No. arXiv:2402.17161). arXiv. <https://doi.org/10.48550/arXiv.2402.17161>