Journal of Urban Technology

ISSN: 1063-0732 (Print) 1466-1853 (Online) Journal homepage: www.tandfonline.com/journals/cjut20

Automating Consultation? Urban Planners and Chatbot-Based AI-Mediated Participation

Jonathan Dortheimer, Michelle L. Oren & Aaron Sprecher

To cite this article: Jonathan Dortheimer, Michelle L. Oren & Aaron Sprecher (28 Aug 2026): Automating Consultation? Urban Planners and Chatbot-Based AI-Mediated Participation, Journal of Urban Technology, DOI: 10.1080/10630732.2026.2693907

To link to this article: https://doi.org/10.1080/10630732.2026.2693907

JOURNAL OF URBAN TECHNOLOGY
https://doi.org/10.1080/10630732.2026.2693907

Automating Consultation? Urban Planners and Chatbot-Based AI-Mediated Participation

Jonathan Dortheimer [a], Michelle L. Oren[b], and Aaron Sprecher[c]

[a]School of Architecture, Ariel University, Tel Aviv, Israel; [b]Independent Scholar, Bogota, Colombia; [c]Faculty of Architecture and Town Planning, Technion Israel Institute of Technology, Haifa, Israel

ABSTRACT

This study reports findings of a two-year embedded case study in which planners integrated a WhatsApp-based large language model (LLM) chatbot into a regional strategic planning process. Across two deployments in the Sabana Occidente region of Bogotá, Colombia, the chatbot facilitated 273 citizen-to-AI interactions covering visioning, project exploration, and voting. Drawing on semi-structured interviews with planners, the study examines planners' expectations, experiences, and criteria for evaluating chatbot-mediated engagement. The planners reported that the chatbot reduced participation barriers but raised concerns about misinformation and shallow depth of deliberation. The chatbot expanded access and supported prioritization, yet remained closer to consultation than deliberation. Democratic value depended primarily on governance arrangements—validation, accountable summarization, and traceability—rather than interface novelty.

ARTICLE HISTORY
Received 2 April 2025
Revised 16 April 2026
Accepted 10 June 2026

KEYWORDS
participatory planning; large language model (LLM); chatbot; public participation

Introduction

Information and communication technologies (ICTs) have reshaped urban life, transforming citizens' interaction with public institutions and public participation in decision-making processes (Fu, Li et al., 2025). However, state-of-the-art critical scholarship on the "smart city" warns that such innovations can depoliticize planning, prioritize corporate and technocratic agendas, and reduce participation to largely symbolic forms of engagement (Hollands, 2015; Kummitha, 2025). A key factor responsible for ensuring that this transformation effectively promotes social inclusion, legitimacy, and shared ownership in urban governance is active citizen engagement (Breux et al., 2023; Chang et al., 2021; Joss et al., 2017). Yet traditional participatory methods—such as surveys, public hearings, and workshops remain constrained by limited accessibility, high administrative costs, and low participant turnout (Fainstein, 2000; Forester, 1999; Healey, 1997; Laurian and Shaw, 2009).

To address these limitations, numerous urban planning studies have increasingly explored innovative digital tools (Arciniegas et al., 2013; Jankowski et al., 2022; Manderscheid et al., 2022; Segura-Tinoco et al., 2022; Thoneick, 2021; Wilson and

CONTACT Jonathan Dortheimer jonathand@ariel.ac.il, School of Architecture, Ariel University, Tel Aviv, Israel

Tewdwr-Jones, 2021). Recently, advances in artificial intelligence (AI), particularly large language models (LLMs) have introduced new possibilities for interactive civic engagement (Dortheimer et al., 2024; Williams et al., 2024; Yang et al., 2024; Yun et al., 2024; Zhou et al., 2024). These models, trained on extensive textual corpora, excel at understanding and generating human language, making them suitable for interactive civic engagement tools (Brown et al., 2020; Williams et al., 2024).

However, existing studies examining AI-based participation remain largely limited to controlled experiments, prototypes, or small pilots, all of which provide limited insight into how such systems operate inside real planning organizations. In particular, little is known about how planners evaluate and operationalize AI-mediated participation in practice.

In this context, we address the following two research questions:

RQ1—How do planners evaluate the integration of LLM-based chatbots into participatory urban planning processes?

RQ2—How does AI-mediated engagement compare with established participatory methods?

We conceptualize LLM-based chatbots as socio-technical infrastructures through which public voice is elicited and translated into planning-relevant inputs. Our focus is on governance, validation, and design choices that affect whether engagement remains consultative or enables more substantive engagement.

To address the research questions outlined above, we conducted a two-year case study with an urban planning team developing a strategic plan for the Sabana Occidente region in Bogotá, Colombia. Seven planners designed and implemented citizen interactions using a WhatsApp-based LLM chatbot alongside traditional participatory tools. Next, we investigated planners' expectations, experiences, and evaluations of chatbot-based engagement relative to traditional participation formats using semi-structured interviews that were conducted at three stages of the process.

By foregrounding planners as designers and evaluators of AI-mediated participation, the results of the present study contribute practitioner-centered evidence to urban technology debates on when conversational AI extends participatory capacity—and when institutional conditions risk turning automation into a substitute for deliberation rather than a genuine expansion of democratic governance.

Background

Citizen Participation in Urban Planning

Citizen participation in decision-making is indispensable to human society and predates Greek democracy (Jacobsen, 1943) and is a central feature of contemporary governance and local policy (Forester, 1999; Healey, 1997). Collaborative frameworks are frequently promoted as ways to address fiscal constraint, demographic change, persistent inequalities, fragmentation, and declining trust in institutions (Brandsen et al., 2017; Durose, 2011).

Motivations for participation in co-creation are commonly categorized into personal and circumstantial drivers (Churchman et al., 2017; Van Eijk and Steen, 2014). Overall,

engagement was explained in the literature using the following three recurring factors: (1) transaction costs, understood as the time and effort required to participate; (2) perceived competence and internal efficacy; and (3) the level of professional support, information provision, and mobilization offered by front-line workers and organizations (Verba et al., 1995; Van Eijk and Steen, 2014; De Graaf et al., 2015; Wyborn et al., 2019). In addition, citizens are more likely to invest effort when an issue at stake affects their lives (Simmons and Birchall, 2005; Thijssen and Van Dooren, 2016). Furthermore, structural barriers can hinder citizens' effective engagement: (1) narrow definition of who qualifies as a stakeholder; (2) power imbalances between authorities and the public; (3) trust deficits; (4) communication barriers that make planning discourse inaccessible to lay participants (Arnstein, 1969; Forester, 1999; Fainstein, 2000; Aitken, 2012; Zachrisson et al., 2018). Additional barriers include rigid institutional frameworks, limited resources, inexperienced staff, and tight timelines, among others.

Taken together, the evidence suggests that participatory planning should be understood not merely as an institutional design, but also as a motivational ecosystem where values, social connection, perceived efficacy, accessibility, and issue relevance jointly shape who participates and why.

These challenges frame the context for evaluating digital and AI-mediated participation: any tool that reduces transaction costs, extends spatial and temporal accessibility, or bridges technical language gaps holds theoretical promise—while also introducing new risks of power reconfiguration and algorithmic filtering of public input.

Digital Tools and AI-Mediated Participation in Urban Planning

A growing range of digital tools has been employed in participatory planning, including digital twins, interactive maps, social media platforms, virtual reality, and—more recently—chatbots (Wilson and Tewdwr-Jones, 2021; Du et al., 2024; Ataman and Tuncer, 2022; Dortheimer et al., 2023). Smartphone and Internet-based tools were reported to effectively expand the geographic scope of public participation, attract diverse groups, and enhance transparency in urban and regional planning processes (Griffin and Jiao, 2019; Jankowski et al., 2022; Schroeter and Houghton, 2011). However, evidence on the quality of digital participation remains mixed, with persistent tensions between participation volume and deliberative depth (Arciniegas et al., 2013; Davies et al., 2022).

Recent progress in machine learning has led to the development of transformer-based language models capable of generating human-like text (Brown et al., 2020). A growing body of research has started documenting and evaluating various urban planning LLM-based applications (Fu, Brinkley et al., 2025; Liu et al., 2025). Most of these studies highlighted diverse functionalities of LLMs, ranging from sophisticated scenario simulations to sentiment analysis, simplifying professional language and summarization of citizen feedback (Fu, Li et al., 2025; Li et al., 2024; Zhou et al., 2024; Yun et al., 2024).

Chatbots in Local Government and Planning

Public agencies primarily deploy chatbots for service-oriented tasks, such as information provision, routine transactions, and complaint management (Senadheera et al., 2024).

A systematic review of chatbot use in local government found that most deployments focus on improving service efficiency rather than enabling citizen influence over policy (Senadheera et al., 2024). In this context, "consultation" is often operationally defined as AI-powered information retrieval. This differs significantly from planning theory, where consultation describes an active solicitation of public input to shape policy before decisions are finalized (Fung, 2006; Rowe and Frewer, 2005).

Within Arnstein's (1969) Ladder of Citizen Participation, such consultation is viewed as a "tokenistic" or intermediate level of engagement, where authorities seek feedback but retain all decision-making power; higher levels—like partnership or co-creation—require citizens to meaningfully influence outcomes (Fung, 2006). While chatbots may broaden outreach and accessibility (Liu et al., 2025), their current role in service delivery primarily facilitates navigation rather than the democratic shaping of policy—a distinction influenced by the system's perceived "humanness" and its capacity for genuine deliberation (Senadheera et al., 2024).

Unlike rule-based chatbots that are based on scripted decision trees and pre-programmed responses, LLM-based chatbots are capable of open-ended natural dialog. Recent research also explored using LLM-based chatbots to support democratic deliberation by assisting individuals with diverse viewpoints in identifying common ground. Tessler et al. (2024) found that participants preferred AI-mediated statements over human-mediated alternatives, and that the system reduced polarization while incorporating minority perspectives. However, risks related to inaccuracies, hallucinations, and embedded biases may undermine trust and distort public understanding (Williams et al., 2024). Without adequate human oversight, such systems may inadvertently propagate misinformation, reinforce existing inequalities, or misrepresent critical community perspectives (Obermeyer et al., 2019). Accordingly, the integration of LLMs and chatbots into urban planning processes requires proactive mitigation strategies, including transparent system design, systematic bias detection, and sustained human oversight (Williams et al., 2024). Building on the above, in the present study, we examine the use of an LLM-based chatbot as a participatory tool. By embedding the system within an ongoing institutional planning process and placing its design, use, and evaluation in the hands of professional planners, we seek to provide evidence on how AI-mediated participation operates in practice and how the opportunities and risks identified in the literature materialize when LLMs are applied to support public engagement in complex urban governance contexts.

Methodology

In addressing our research questions, the research team collaborated with the non-profit association ProBogotá Región, which was leading the preparation of the non-statutory Strategic Plan for the Western Savannah of Bogotá 2051. This plan addresses territorial challenges across five municipalities—El Rosal, Facatativá, Funza, Madrid, and Mosquera—covering approximately 932 km² and home to nearly 600,000 residents. Local challenges include limited connectivity, restricted economic opportunities, housing deficits, environmental degradation, congestion, informality, erosion of local identity, and marginalization of vulnerable populations.

Conventional Participation Methods

ProBogotá Región traditionally used a structured, top-down participatory methodology based on stakeholder engagement tools, including online surveys, workshops, hearings, rating and voting exercises, as well as both in-person and virtual meetings. A structural limitation of these formats is their reliance on small, invited stakeholder groups—typically 20 to 45 participants per session—selected to represent regional interests, with limited opportunity for open or scaled engagement. The strategic plan was developed in four phases (See overview in Figure 1) with conventional in-person and virtual participatory initiatives incorporated into each of the first three phases.

Chatbot Design and Deployment

In parallel to the conventional strategic planning process, we integrated a chatbot into the first two planning phases, thereby establishing a new and direct communication channel with the general public. The chatbot was designed through a structured process. During the ideation stage, the planners defined the specific participatory objectives of each chatbot deployment, including the planning topics to be addressed, the intended user groups, and the type of public input to be collected. The design stage involved translating these objectives into conversational flows, question sequences, and response logics, as well as defining the tone, language, and interaction style of the chatbot.

The research team developed the chatbot software (See the Appendix for further details). The chatbot, branded Charli (derived from the Spanish verb charlar, "to chat"), was integrated into WhatsApp, an instant messaging platform, to provide a familiar and accessible user experience. WhatsApp's voice message functionality made the tool accessible to older adults, low-literacy users, and visually impaired participants.

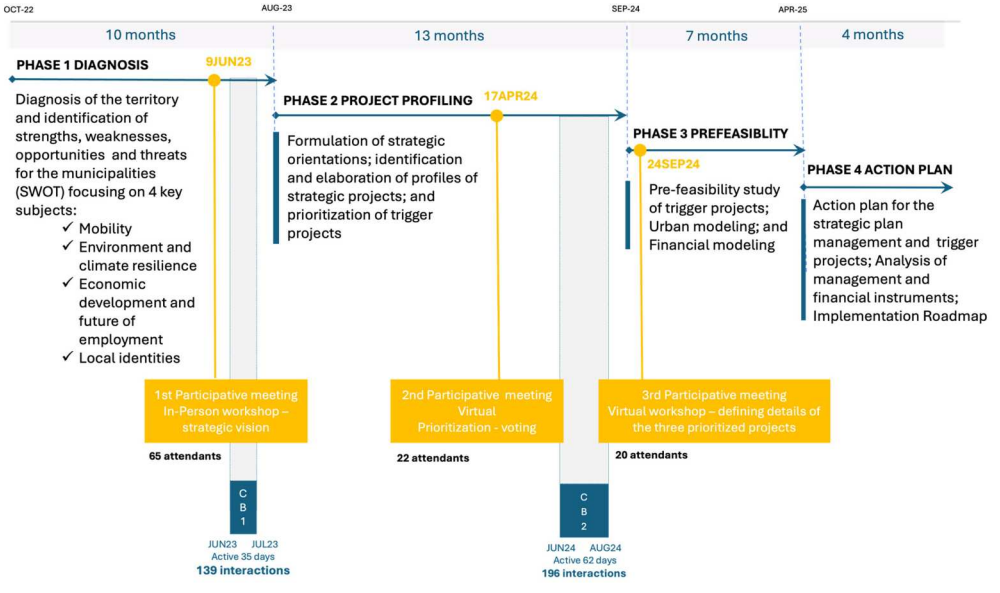

Figure 1. Phases of ProBogota's strategic planning process

Its inclusion in mobile data plans in Colombia also enabled participation without additional data costs.

Chronology of the Planning Process

The strategic plan was developed in four phases. The chatbot was activated during the first and second planning phases (See Figure 2).

Phase 1: From Diagnosis to Strategic Orientations. Beginning October 2022, planners identified key strengths, weaknesses, opportunities, and threats related to connectivity, economic development, environmental sustainability, and social inclusion, producing strategic guidelines to shape a long-term regional vision. The outcomes of this phase were subsequently shared with stakeholders for subsequent review and validation.

In March 2023, the planning team held an open online informational webinar on AI and LLMs to introduce, among other topics, the proposed use of a chatbot as a planning tool for the forthcoming public participation experiment.

In-person workshop: In June 2023, the planning team held its first in-person participatory workshop with 57 participants from local government, academia, and the private sector to collaboratively develop the region's strategic vision and guiding orientations. Participation was by invitation and required prior registration. The workshop brought together a total of 15 representatives from local governments, eight from the central government, six from academia, 28 participants from the private sector, and members of the planning team.

Chatbot interaction was designed to elicit the participants' visions for the region's future and, after each response, expressed interest and encouraged deeper reflection through follow-up questions (See example in Figure 3). Concurrently, the first chatbot (CB1) was launched to broaden public reach beyond the invited workshop participants. Over 35 days, 133 users exchanged 1,910 messages (M = 14.36, SD = 11.68). A Design Brief Generator synthesized the conversations into seven action themes—including

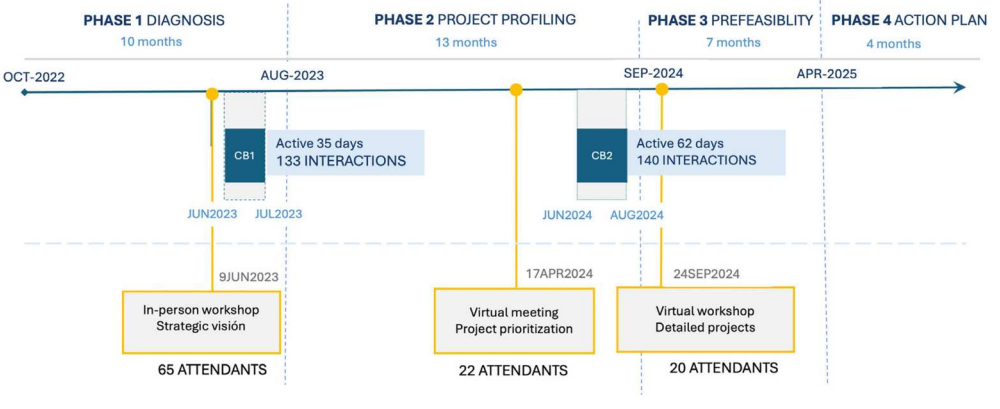

Figure 2. Timeline of participative interactions

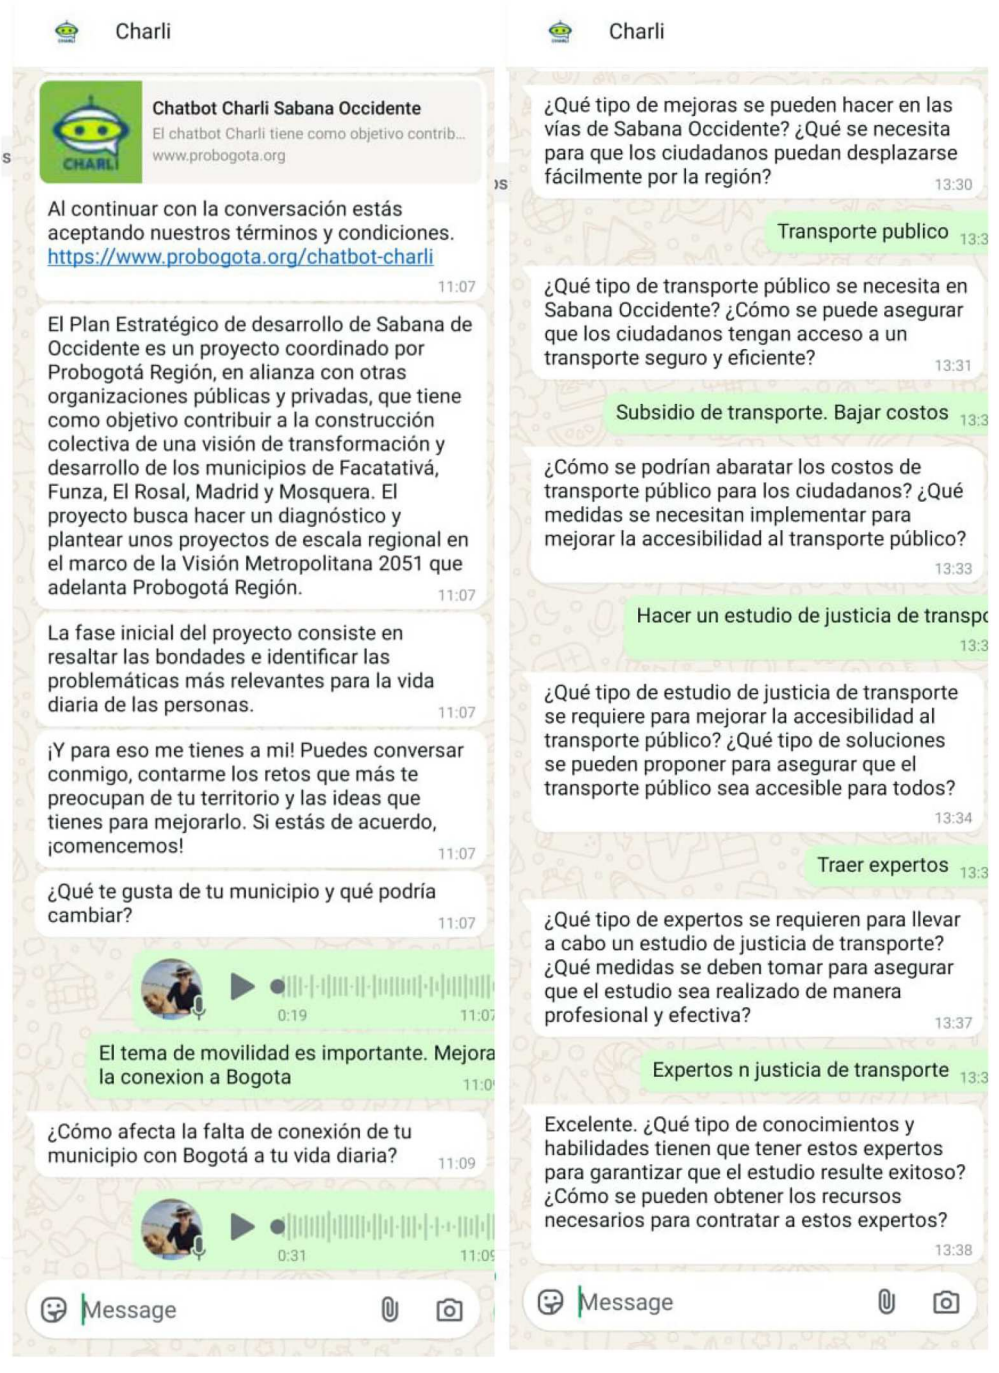

Figure 3. Chatbot user interface example screens from chatbot interaction

mobility, renewable energy, education, and security, the last of which had not surfaced in the in-person workshop. Anonymized raw data were shared with planners to allow independent verification.

Phase 2: From Project Profiling to Prioritization. In September 2023 the planning team brainstormed project ideas to address the challenges and formulated 10 proposed projects grouped into four categories: habitat, knowledge seeding, logistics, and mobility.

Virtual meeting: In April 2024, an online meeting was convened to evaluate and prioritize the proposed projects, bringing together a total of 22 participants from local government, academia, the private sector, and the planning team. The prioritization process applied an evaluation methodology that identified three projects for detailed design, with two projects receiving closely comparable scores.

The second chatbot interaction (CB2) was used to gather input from citizens and experts, contributing to the prioritization of projects (See Figure 4). The chatbot recorded a total of 140 interactions, including contacts from the first phase. Each interaction averaged 2.15 messages (SD = 4.13) and lasted on average 1.87 days (SD = 4.7). After CB2 completed its 62-day cycle, CB2 votes exceeded those cast in a parallel online expert meeting (26 votes) and produced a clear ranking: the Regional Food Logistics Center, the Integrated Mobility System, and Urban Mobility Infrastructure were selected for development.

Phases 3 and 4: Feasibility to Action Plan. The three prioritized projects underwent technical, urban, and financial feasibility analysis, and were incorporated into a final Strategic Plan document. An action plan was then developed to integrate the strategy into Bogotá's land-use policies (POT), alongside a proposed subregional governance framework for long-term implementation.

Data Collection

Our primary data source was a series of semi-structured interviews conducted with seven members of the planning team responsible for designing and implementing the chatbot-based participatory processes (See Table 1). To capture the participants' evolving perceptions, expectations, and evaluations, the interviews with them were conducted at three key moments in the planning cycle: prior to the first chatbot deployment (pre-CB1), prior to the second deployment (pre-CB2), and after the completion of the second deployment (post-CB2) (See overview in Figure 5). All interviews were administered individually via video conferencing in Spanish, audio-recorded with informed consent, transcribed verbatim, translated to English when necessary, and anonymized prior to the subsequent analysis (See the Appendix for further detail).

To contextualize interview accounts, the research team also conducted participant observation throughout chatbot design and deployment, maintained a reflexive researcher diary, and collected system-generated interaction logs.

Analysis

Reflexive thematic analysis (RTA) was used to analyze the interview data. Analysis was conducted using Braun and Clarke's (2006) six-phase approach by two coders. The team undertook repeated readings of the transcripts to achieve data familiarization and maintained reflexive memos to document analytic insights and positionality throughout the process. Initial codes were systematically generated across the dataset,

Figure 4. Chatbot user interface example screens from CB2, allowing project browsing by category and project selection

Table 1. Interview participants

ID	Role	Seniority	Interview Stages
Planner 1	Analyst	Senior	Pre-CB1, Pre-CB2, Post-CB2
Planner 2	Researcher	Junior	Pre-CB1
Planner 3	Coordinator	Senior	Pre-CB1
Planner 4	Analyst	Early-career	Pre-CB1, Pre-CB2
Planner 5	Director	Senior	Pre-CB2
Planner 6	Assistant	Junior	Post-CB2, Pre-CB2
Planner 7	Assistant	Junior	Pre-CB2, Post-CB2

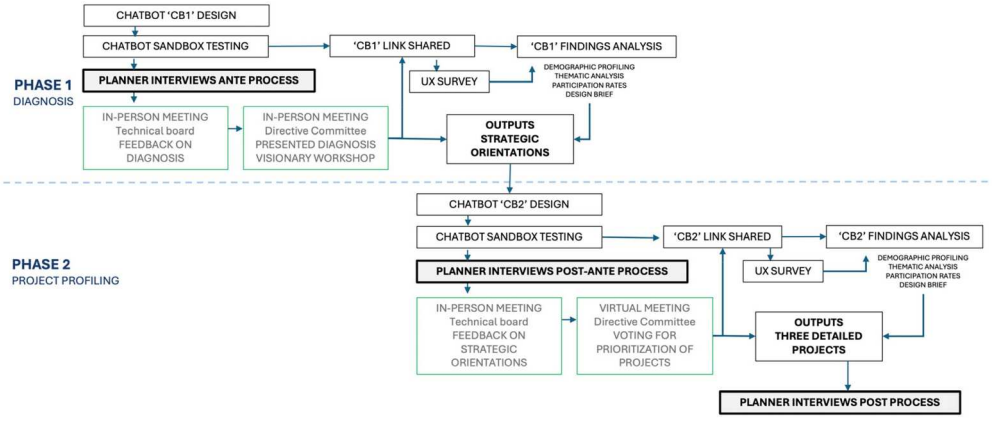

Figure 5. Research methodology flow

and candidate themes were developed through the identification of patterns. The identified themes were iteratively reviewed, refined, and defined to ensure coherence and conceptual distinction, and were interpreted on both semantic and, where relevant, latent levels in alignment with our research questions.

Findings

Across the two-year case study, planners evaluated the integration of an LLM-based chatbot primarily through a pragmatic planning lens, as an added participation channel that could extend reach, lower barriers to contribution, and make inputs more workable as "data" (RQ1). In the pre-CB1 interviews, they positioned existing participation as structurally constrained by limited representativeness, politicized, monopolized in-person dynamics, and persistent translation gaps between technical planning language and citizens' everyday concerns. They, therefore, expected the chatbot to help with scale, candor, centralization, and synthesis of dispersed feedback. After the CB2, the chatbot was judged effective for structured tasks such as project prioritization, supporting clear, fast voting, and helping resolve ties. However, it was also seen as less capable of producing the deliberative richness and shared closure associated with workshops and in-person meetings, with outputs described as fragmented data points rather than collective meaning-making (RQ2). Planners further identified design and governance tensions: conversational flows were too complex to sustain user engagement, and the LLM's generative prompting could steer agendas in ways that complicated transparency and control. Overall, AI-mediated engagement was valued as a complementary, efficient and scale-oriented component of a hybrid participatory approach, but not as a substitute for established methods when the goals are co-creation, deliberation, and consensus-building.

Limitations of Existing Participation Methods

Before experiencing AI-mediated participation via the bot, the planners described participation primarily as structured data gathering and project prioritization embedded

in institutional workflows. The tools—including Kahoot, Miro, Mentimeter, and online voting platforms—were described not as innovative civic technologies, but as extensions of traditional planning participation practices, adapted incrementally to digital environments (especially after COVID-19). The tools appeared in the interviews as supportive, rather than transformative tools. One planner provided the following account: "the people who show up are the same interest groups that always participate in the same spaces, not open communities, and there's not such a broad, general interest in participation" (Planner 4).

Furthermore, several planners pointed out that methodologies for public participation remained quite similar regardless of the specific project size. One planner confirmed, "Yes, even for a building. While we do not specifically handle buildings here, in general, the participation processes are quite similar" (Planner 3). Another interviewee noted:

> Participation was dominated by the same interest groups, who often claim to represent the larger population of localities with up to a million residents, yet meetings would have only about 45 attendees. When participation is scaled up to city-wide planning, we still rely on voluntary workshops and discussion panels over 12 to 18 months, which is not sufficient. Major decisions require longer and more extensive processes to ensure legitimacy …
>
> Participation is far from perfect. When applied to broader urban planning discussions, our current methods become inadequate. They are neither innovative nor flexible enough to engage a wider audience. Even if someone's participation is minimal, we should have better strategies to include them in the process. (Planner 4)

The interviewed planners consistently described current participation and data-collection practices as difficult to scale, cognitively exclusionary, and institutionally fragmented. A recurrent concern that emerged in the interview data analysis was the limited reach of participation, particularly in digital formats. One planner noted:

> Online surveys had very low response rates despite considerable efforts. Reaching targets proved extremely difficult. Even when participation occurred, it was often restricted to a small group of invited individuals and characterized by uneven engagement, where some are highly engaged, while others contribute very little. This constrained participation resulted in non-representative data, limited in analytical value. (Planner 1)

Another widely reported challenge concerned the technical and communicative barriers embedded in current participation methods. The interviewees emphasized that planning discourse remained largely inaccessible to lay public, explaining that "experts and planners use specialized terminology that is hard for citizens to understand," and that many planning topics are "highly technical and difficult for the average citizen to grasp" (Planner 1). As a result, responses were frequently misaligned with institutional planning logics. One interviewee observed that, while governments prioritized "a long-term, city-wide vision," citizens tended to focus on "immediate, local issues," such as "fixing the pavement outside their home" (Planner 1). This mismatch was perceived as undermining both the relevance and the usefulness of the collected data.

Furthermore, the interviewed planners also highlighted systemic and infrastructural limitations, particularly the absence of integrated data environments. As one participant explained:

> Various data sources exist, but they are not systematically integrated, requiring planners to rely on manual compilation and interpretation, which makes the process inefficient, more complex, and less effective. (Planner 2)

Other interviewees criticized institutional participation models as fundamentally outdated, arguing that participation is still treated as mere information extraction rather than genuine dialogue. As one planner put it: "I think there is a conversation in Charli, but it's not real, because it's a conversation to extract more information. It's not a conversation that invites you to question things or open your mind or see another point of view … " (Planner 1). Traditional instruments—such as surveys and QR-code-based public meetings—were described as producing "little follow-up" and remaining largely "superficial," despite experimental efforts such as gamification intended to improve understanding and engagement.

Finally, the interviewed planners noted a broader conceptual ambiguity surrounding participation itself. For instance, one of the respondents provided the following account:

> Citizens often complain about a lack of participation but then fail to engage when given the opportunity, while institutions struggle to define the limits of public influence, including whether citizens should have full decision-making power or whether expert planners should retain strategic control. These unresolved tensions were seen as further weakening the participatory process. (Planner 1)

The interviewed planners' overall evaluations of current data-collection practices reflected these challenges. Existing methods were generally rated as moderately to weakly effective. While some planners acknowledged that current tools "provide valuable insights into the population" (Planner 2), one interviewee summarized, "There are efforts, but participation remains misunderstood and ineffective" (Planner 3), thus aptly capturing the prevailing assessment across the interviews.

Taken together, these assessments establish the baseline against which planners subsequently evaluated the chatbot: a participatory landscape characterized by limited scale, unequal access, and a persistent gap between technical planning discourse and citizens' communicative capacities. The subsections that follow examine whether, and under what conditions, AI-mediated engagement addressed these deficits.

Expanding Scale and Reducing Barriers

The overall perception of the planning staff involved in the project was that the chatbot significantly contributed to the planning process by providing a broader perspective and bringing new information to the design process. For instance, one interviewee noted:

> Charli helped us understand the territory and its issues from a broader perspective. Many challenges can only be truly grasped by those who are part of the territory, the inhabitants. Studying the territory academically or through satellite plans differs greatly from experiencing it firsthand. (Planner 5)

Furthermore, from the planners' perspective, the chatbot helped to validate some ideas discussed in the meetings. One planner stated: "Charli helped validate many of the discussions held in the steering committees. Many topics that arose in the chatbot discussions aligned with those discussed in the Zoom and in-person meetings" (Planner 1).

In addition, as per several respondents, the chatbot's conversational interface improved interaction, providing more detailed insight than online surveys. As Planner 2 phrased it, "Charli has a great virtue: it deepens conversations and introduces different tensions depending on the interlocutor. It does not provide static responses. Rather, it evolves within the conversation itself."

According to the interviewed planners, the chatbot played a significant role in democratizing the planning process by facilitating open discussions with residents. As several interviewees mentioned, the chatbot helped to broaden participation by being highly accessible and offering a more enjoyable and comprehensible experience. Similar to other online tools, it was available anytime and anywhere, but its use of natural language conversations made interactions more enjoyable and comprehensible. Accordingly, one planner remarked, "Unlike traditional participatory processes that rely on public officials or designated facilitators, Charli allows flexible engagement. Additionally, being an innovative WhatsApp-based tool, it reaches people who usually aren't interested in traditional participation formats" (Planner 2).

Another point made by the planners interviewed was their satisfaction with the extended participation achieved by overcoming the time and space limitations of in-person meetings. In this respect, Planner 4 noted: "It reached 133 users. While this may not be a large percentage of the population, it was a significant achievement in terms of outreach" (Planner 4).

Furthermore, the chatbot's use of a popular instant messaging platform enabled easy and casual engagement, reaching a broader audience. This point is captured in the following interview excerpt: "Previously, these processes were managed in-house with technical teams and expert validation. With Charli, we opened this process to non-specialized and expert validation" (Planner 3).

One aspect that was particularly appreciated by the planners was the chatbot's accessible language, as it translated technical terms into simpler expressions, making complex planning concepts more understandable. In this connection, one of the interviewees said, "I see it as a great tool for simplifying complex planning concepts and helps translate technical language used in strategic plans into more accessible conversations" (Planner 3).

In addition, one planner noted that young people tend to avoid participating due to a lack of interest: "young people could be much more involved because I do feel that [in-person spaces] are spaces that are like … for young people … it's like going to class … " (Planner 1). Yet several interviewees shared the belief that the chatbot successfully broadened participation among younger individuals. For instance, Planner 2 mentioned that, "The chatbot's ability to engage young people, who typically do not participate in these processes due to time constraints or a lack of interest in attending in-person meetings, was particularly noteworthy" (Planner 2).

Structured Tasks Versus Deliberation

At the concluding interview, when asked if the outcomes of chatbot performance were comparable to those of the traditional participatory tools, a planner provided the following detailed account:

> I believe Charli had two functions in this second phase.

> First, in our voting system for ProBogotá, Charli did influence the final decision on project selection. It significantly contributed to determining the three projects that were ultimately chosen. This was because, in our usual voting scheme, there were many ties, and Charli helped resolve them quantitatively by prioritizing projects. Second, beyond Charli itself, this experience helped us at Probogota understand our potential influence in initiatives like this—both in terms of communication and engagement.
>
> Charli also had a role in encouraging participation, serving as a gauge of how well we could engage people using chatbot-based tools like the ones we promote at ProBogotá. I don't think Charli can be compared to traditional tools we have been using for similar purposes in the sense that the conclusions drawn in other scenarios are much richer due to the conversations that take place around people's participation. This allows for extensive discussions about what people are thinking and creates an environment of consensus, regardless of the topic.
>
> With Charli, on the other hand, I feel like we get a lot of very specific data points, but it's difficult to grasp the real outcome of the exercise. The results become just small, isolated responses. It kind of shuts down the debate. (Planner 1)

This account positions the chatbot firmly at Arnstein's rung 4—consultation: individual preferences were elicited and aggregated, but citizens had no opportunity to hear or contest one another's views. The observation that outputs felt like "small, isolated responses" that "shut down the debate" speaks directly to RQ2—where workshops produce shared closure through discussion, the chatbot produced discrete data points that foreclosed collective reasoning. Crucially, a tool that resolves institutional ties through quantitative aggregation displaces deliberative judgment without making that displacement visible to participants—a governance concern, not merely a design limitation.

Overall, planners perceived the chatbot as a powerful tool amplified their informational advantage. Planner 1 noted, "Planners have more power because we have more information, and information is power." This sentiment was echoed by Planner 2, who said, "I think Charli is a tool that amplifies this power and, at the same time, gives citizens impressive access to information and power." Yet this framing reveals an asymmetry central to RQ1—planners gained a richer dataset and a tie-breaking instrument, while citizens gained channel access but not visibility into the reasoning process that converted their inputs into decisions.

Misinformation, Inconsistency, and Governance Deficits

Despite the strengths of the chatbot that were explicitly mentioned during the interviews, some planners were concerned with its technological aspects. While most conversations were enjoyable, they sometimes failed to cover a broader spectrum of the required information. The open-ended nature of discussions posed a distinct challenge of recognizing the stakeholders' needs. As one of the planners put it, "We wanted to understand how they perceived this diagnosis and the region's needs. However, achieving this awareness was challenging because the conversation was so open-ended" (Planner 1).

Another significant concern that emerged in the interview data analysis was misinformation. The interviewed planners worried about the chatbot communicating incorrect information, which could compromise decision-making. As one planner warned,

"Charli's biggest threat is when people ask something, and Charli responds with information that may not be entirely true. This is a threat because incorrect information can spread easily" (Planner 1).

This is not merely a technical risk but a legitimacy risk: in civic contexts, citizens reasonably expect factual accuracy about local conditions, and inaccurate responses can deepen rather than bridge the information asymmetry between institutions and the public. From a socio-technical infrastructure perspective, the chatbot's generative layer is an active mediator of knowledge whose reliability depends on governance arrangements—prompt design, human oversight, and output auditing—that were only partially in place during these deployments.

The challenge of misinformation can also be attributed to generating design briefs. The research team used two different prompts to create separate briefs from the same dataset. This approach led to notable discrepancies that raised concerns among the planners. The expectation was that tools would provide consistent outputs and ensure data traceability. The discrepancy between two briefs generated from the same dataset using different prompts illustrates that chatbot outputs are co-created by summarization logic, not simply reflective of citizen inputs. Traceability—the ability to audit how raw interactions became planning artifacts—is, therefore, a structural requirement for democratic accountability, not an optional quality check.

Discussion

In this study, LLM-based chatbots were conceptualized not merely as digital tools for data collection, but rather as socio-technical participatory infrastructures capable of mediating, structuring, and translating public voices into forms of knowledge actionable within institutional planning processes. Accordingly, rather than functioning as passive interfaces, the chatbots tested in the present study operated as algorithmic arenas of participation that shaped how participation occurred, what types of contributions could be expressed, and how those contributions were reformulated into planning-relevant outputs. Viewed from this perspective, the chatbot constituted a new participatory layer alongside traditional mechanisms such as workshops, public hearings, and surveys.

These processes converted diverse, unstructured expressions of opinions, narratives, preferences, and ideas into standardized planning artifacts, including strategic orientations, ranked projects, and design briefs. Through this translation, the chatbot transformed informal public voices into legible and actionable planning knowledge.

In summary, the chatbot operated as a mediating layer between the public and planning institutions, shaping the format of participatory outputs. This conceptualization positions AI-mediated participation not as a supplementary communication channel, but rather as a governing participatory infrastructure that actively organizes civic input, embeds algorithmic logics into decision-making, and restructures institutional participation pathways.

Enhancing Participation and Engagement

A significant contribution of the chatbot identified in the interview data was its ability to expand participation, supporting findings from other studies (Liu et al, 2025). Unlike in-

person and online meetings, which the interviewed planners perceived as resource-intensive and time-constrained, the chatbot provided a more accessible and flexible platform for 24/7 engagement, even after the meeting, especially since it was accessible through a widely used messenger service. Not only did the planners find the chatbot more enjoyable, but also the collected data provided, in the interviewees' view, a more nuanced view of public opinion. Indeed, the chatbot's conversational interface facilitated richer and more exciting interactions than conventional online tools (e.g., surveys).

While AI-mediated participation radically expanded the scale, accessibility, and continuity of public engagement, its use in the present study simultaneously introduced a new layer of infrastructural power that was neither neutral nor merely technical. Overall, by mediating how public voice is elicited, structured, summarized, and rendered legible to institutions, LLM-based chatbots actively participate in the production of "what counts" as public input. This shifts participation from an open-ended communicative practice into an algorithmically curated knowledge pipeline, where visibility, salience, and influence get increasingly shaped by prompt design, data preprocessing, summarization logics, and model architectures. As a result, power is partially displaced from participatory forums and deliberative publics into invisible socio-technical infrastructures governing inclusion, omission, amplification, and abstraction.

That said, without transparent governance, contestable design processes, and institutional accountability, such systems risk expanding engagement while concentrating interpretive power in the prompt and summarization pipeline. In our case, the interviewed planners emphasized that what became actionable knowledge depended on design choices and on whether outputs could be audited back to raw logs—particularly when summaries were used to justify planning priorities.

Hence, the real constraint, as grew explicit in our results, was not technology—it was governance. The true limits were defined by the answers to the following questions: Who controls the prompts? How is aggregation done? What becomes "official" knowledge? How is algorithmic bias governed? In essence, these are political constraints, not technical ones.

Information Collection and Validation

The information collected by the chatbot was broadly comparable to that obtained through in-person and virtual meetings, and in some respects more expansive. CB1's 1,910 messages across 133 participants produced a thematic map that largely corroborated the strategic orientations developed through the invitation-based workshop. The collected information validated topics such as mobility, renewable energy, and education while also surfacing a new concern about security that the workshop did not raise. This empirical result is theoretically significant: the security theme's emergence in the open-access chatbot channel, but not in the curated workshop, suggests that invitation-based participation systematically excludes certain publics and their priorities, and that chatbot-mediated open access can partially compensate for this structural exclusion. CB2's voting outcomes likewise corroborated and extended the expert prioritization meeting, producing a clear ranking that aligned with expert assessments while drawing on a larger and more diverse respondent base. This finding is largely consistent with

previous research showing that chatbots produce an amount of information similar to surveys (Dortheimer et al., 2024).

However, the study reveals a limitation: the chatbot's summarization outputs were not stable across prompting conditions. When two different prompts were applied to the same CB1 dataset, the resulting design briefs differed in ways that planners found substantively concerning. This reproducibility problem has particular consequences in planning contexts, where the transformation of raw citizen inputs into strategic artifacts is expected to be auditable and defensible.

Misinformation and Trust

Trust, which plays a significant role in participatory planning, particularly in the interactions between planners and residents (Aitken, 2012; Zachrisson et al., 2018), is essential for fostering a participatory process that encourages mutual learning, deliberation, and consensus. Interestingly, mistrust can also drive greater participation as stakeholders seek to influence outcomes.

Trust and misinformation constitute a central theme in AI, especially in LLM-based technologies such as chatbots (Liu et al., 2025). When incorporating LLM chatbots into this process, planners and participants must trust the technology to reliably communicate ideas (Williams et al., 2024). In the present study, the chatbot facilitated this by providing conversation logs that the interviewed planners could review to ensure the chatbot was accurately representing residents' opinions. This traceability of information back to its source is a crucial aspect of trustworthy AI systems.

However, the participatory planning process examined in this study was characterized by limited transparency by design, as information was shared only at specific, predefined stages when participation was formally solicited, and otherwise primarily through the final report dissemination at the end of the process. Although the original design presupposed that the participants' contributions would appear in real time as part of an automatically generated and continuously evolving, albeit supervised, "idea cloud" visible on a public website, this live visualization component was never implemented due to technical complexity. As a result, the citizens had little visibility into how their inputs related to those of others, or how ideas evolved throughout the process.

Conversely, chatbot technology could keep participants more informed throughout the planning process, enhancing transparency and building trust for current and future initiatives. For instance, if a user-generated idea was implemented, the chatbot promptly notified the participant, acknowledging their valuable contribution.

Current Limitations of Chatbot-Mediated Participation

Integrating the chatbot within participatory planning in the present study did not facilitate a genuine dialogue among stakeholders and planners. This limitation is critical, as the essence of participatory planning is that stakeholders not only express their individual interests, but also understand and respond to the perspectives of others. Conventional participatory frameworks rely on discussions, allowing the participant parties to challenge, refine, and collaboratively learn from one another. Such a deliberation process can reach a consensus by examining and reconciling diverse viewpoints. However, the

current chatbot did not mediate differing stakeholder opinions or present counterarguments among participants, which considerably limits its deliberation potential.

Overall, the rationale behind participatory planning is not simply to collect data or opinions, but to actively involve participants in shaping opinions and decision-making. While the chatbot broadened such participation, which can be regarded as a positive step towards user-centered design, genuine participation should extend beyond mere consultation, aligning with democratic ideals such as those embodied in the Colombian constitution, which emphasizes citizen involvement in governmental processes.

Positioned within Arnstein's (1969) framework, the chatbot deployments remained at rung 4—consultation. Citizens could share preferences through open channels, but they had no formal influence over how their inputs were interpreted, summarized, or used. Three mechanisms kept participation at this level. First, participants could not see or evaluate how 1,910 messages were turned into thematic orientations by the Design Brief Generator and the planning team. Second, they had no way to challenge summaries that omitted or misrepresented their views. Third, they had no role in deciding whether chatbot outputs would be adopted, ignored, or overridden; that remained entirely with the planning team. As a result, the chatbot functioned mainly as an efficient input-gathering tool rather than a deliberative forum.

To move toward partnership, future deployments would need transparent processes, mechanisms for contesting AI-generated outputs, and clear rules for how chatbot results are weighted in planning decisions. Reaching delegated power would require citizens to have a formal say over whether and how those outputs shape decisions. These changes are less technical than institutional: they depend on planning agencies giving citizens real interpretive authority. Advanced chatbot features may support discussion and learning, but they cannot by themselves move a deployment higher on Arnstein's ladder.

Study's Contributions

This study makes three contributions to debates on AI-mediated participation in urban planning.

Theoretically, we extend Arnstein's (1969) ladder of participation to AI-mediated contexts by showing that LLM chatbots expand the breadth of participation by reaching more people and surfacing more issues without increasing the depth of that participation. Participants had no deliberative role, decision-making power, or ability to challenge how their input was summarized. This suggests that consultation-mode AI systems are not neutral tools but socio-technical infrastructures that shape what counts as public knowledge. The key analytical shift is from the chatbot interface to the underlying algorithmic pipeline and who controls it.

We suggest three policy implications: First, public agencies should be transparent about the LLM logic used, to ensure accountability. Second, AI-generated themes used in planning should meet standard evidentiary and auditability requirements. Third, chatbot-based participation should be classified as consultation, not co-creation, unless citizens can meaningfully contest, revise, and influence how their inputs are used.

For practitioners, the findings suggest that LLM chatbots work best for bounded tasks such as issue-raising, preference elicitation, and prioritization, rather than open-ended

deliberation. Their outputs should be treated as supplementary signals, not direct expressions of public opinion, and should be interpreted and verified by planners.

Limitations and Future Research

A limitation of this study is its focus on a single, context-specific planning process, which allows for an in-depth and nuanced analysis; future research could build on these insights by examining additional cases across diverse contexts. This focus is inherent to studying an embedded, real-world intervention, and the case study format aligns well with the study's aim of explaining how planners evaluate and integrate LLM-based chatbots within ongoing participatory planning practice, rather than measuring population preferences or estimating causal effects.

Furthermore, our analysis focused on the planners' perspectives and did not systematically evaluate citizens' experiences, trust, or perceived legitimacy of the chatbot channel. Finally, chatbot performance depended on the specific LLM configuration, so prompts and summarization pipelines used with different models or guardrails may yield different results.

Future research should therefore: (1) evaluate citizen-side legitimacy, trust, and perceived efficacy in AI-mediated channels; (2) test AI interactions that move beyond consultation toward deliberation; and (3) develop and assess governance frameworks for "accountable summarization," that could include traceability, reproducibility, and protocols clarifying how AI-derived outputs are validated and used.

Conclusion

This paper explored the integration of an LLM-based chatbot as an active participation channel within a regional strategic planning process. Across two deployments, the instant-messaging chatbot facilitated citizen-AI interactions that complemented traditional invitation-based workshops and virtual meetings. By focusing on planners as both institutional designers and evaluators, the study examined how they assess and integrate LLM chatbots into participatory planning in comparison with established methods, providing insights into the practical opportunities and challenges of AI-mediated engagement. Three overarching conclusions follow:

(1) Planners valued the chatbot for expanding access via asynchronous messaging. Its conversational format elicited more situated, place-based concerns than typical digital tools.
(2) The chatbot scaled consultation and prioritization well but did not support deliberation. It produced many individual data points, aligning with "consultation" on Arnstein's ladder rather than shared reasoning or decision power.
(3) Legitimacy depended on governance and accountability of the AI pipeline. Issues like misinformation, inconsistent summaries, and weak traceability made control, validation, and auditability central to turning input into planning knowledge.

Overall, the case indicates that LLM chatbots can materially expand participation and provide actionable insights for planners. However, without robust governance and

validation, the same automation that expands access can also narrow democratic capacity by turning participation into efficient extraction rather than collective reasoning. The central implication is thus institutional: the democratic potential of AI-mediated participation hinges on making the AI layer transparent, contestable, and accountable within planning decision processes.

Acknowledgments

We express our deepest gratitude to ProBogotá Región for their exceptional dedication throughout this project. We are especially grateful to their core team of planners whose complete engagement was vital to the process: María Carolina Castillo (President), Bibiana Rodríguez (Director of Sustainable Urban Development), Martin Anzellini (Director of Smart Cities), Cristina Hermida (Urban Projects Coordinator), and Ricardo García (Designer). Their invaluable contributions spanned every stage of the study, from designing Charlie's graphic representation and developing personality interactions, to facilitating outreach through their official spokespersons and social networks. Additionally, we extend our sincere thanks to Nicolás Gutiérrez (Coordinator) and Natalia Silvia Mora (Subsecretary of Territorial Planning of Bogotá) for their distinct technical contributions to this research.

"Charla con Charli": AI chatbot for Bogota's Occidental Savannah's inclusive Strategic Plan won the first place in the 2024 GO SMART Award, presented by the Global Organization of Smart Cities (GO SMART), during the Smart City Summit & Expo (SCSE) held in Taipei, Taiwan, on March 12, 2024.

We thank the anonymous reviewers for their valuable feedback.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the Israeli Ministry of Innovation, Science and Technology under Grant No. 0004675 / 2034517. The authors gratefully acknowledge this support.

Notes on Contributors

Jonathan Dortheimer is a senior lecturer at the School of Architecture, Ariel University, where he heads the Architectural AI Research Laboratory (aairl.com), and is a guest scholar at the Chair of Architectural Informatics, Technical University of Munich. His current work develops urban digital twins incorporating synthetic populations to predict the social outcomes of design and planning decisions, building on his earlier research on AI-augmented, community-engaged design processes.

Michelle L. Oren is a licensed urban planner and scholar whose research examines the socio-technical dynamics of municipal governance, infrastructure, and human rights. She earned her doctorate from the Faculty of Architecture and Town Planning at the Technion-Israel Institute of Technology and has lectured extensively on public asset management, smart mobility, and urban innovation. Dr. Oren is the author of the handbook *Promoting Local Innovation for Inclusion of People in Vulnerable Situations* and *Leaving No One Behind*, published as part of the United Nations Department of Economic and Social Affairs (UNDESA) Handbook Series.

Aaron Sprecher is a scholar in the field of computational architecture and serves as a faculty member at the Technion - Israel Institute of Technology Faculty of Architecture and Town

Planning. He leads the Material Topology Research Laboratory (www.mtrl.tech),and is the academic head of the Technion Advanced Construction Center. In addition to his academic endeavors, he is a co-founder and partner at N-Form, a transdisciplinary architecture practice that fosters collaboration among scholars and practitioners in industrial design, architecture, and engineering.

ORCID

Jonathan Dortheimer http://orcid.org/0000-0002-7464-8526

References

D. Aitken, "Trust and Participation in Urban Regeneration," *People, Place and Policy Online* 6: 3 (2012) 133–147.

G. Arciniegas, R. Janssen, and P. Rietveld, "Effectiveness of Collaborative Map-Based Decision Support Tools: Results of an Experiment," *Environmental Modelling & Software* 39 (2013) 159–175.

S. R. Arnstein, "A Ladder of Citizen Participation," *Journal of the American Institute of Planners* 35: 4 (1969) 216–224.

C. Ataman and B. Tuncer, "Urban Interventions and Participation Tools in Urban Design Processes: A Systematic Review and Thematic Analysis (1995–2021)," *Sustainable Cities and Society* 76 (2022) 103462.

T. Brandsen, W. Trommel, and B. Verschuere, "The State and the Reconstruction of Civil Society," *International Review of Administrative Sciences* 83: 4 (2017) 676–693.

V. Braun and V. Clarke, "Using Thematic Analysis in Psychology," Qualitative Research in Psychology 3: 2 (2006) pp. 77–101.

S. Breux, J. Diaz, and H. Loiseau, "The Smart City—Does the Individual Matter?" Journal of Urban Technology 30: 4 (2023) pp. 33–49.

T. B. Brown et al., "Language Models are Few-Shot Learners," arXiv preprint arXiv:2005.14165 (2020) <https://arxiv.org/abs/2005.14165>.

I.-C. C. Chang, S.-C. Jou, and M.-K. Chung, "Provincialising Smart Urbanism in Taipei: The Smart City as a Strategy for Urban Regime Transition," *Urban Studies* 58: 3 (2021) 559–580.

A. Churchman, E. Wiesenfeld, and E. Sadan, "On Participation and Participatory Practices," in M. A. Bond, I. Serrano-Garcia, C. B. Keys, and M. Shinn, eds., *APA Handbook of Community Psychology: Theoretical Foundations, Core Concepts, and Emerging Challenges* (Washington, DC: American Psychological Association, 2017) 315–334.

J. Davies, M. Arana-Catania, and R. Procter, "Embedding Digital Participatory Budgeting within Local Government: Motivations, Strategies and Barriers Faced," paper presented at the 15th International Conference on Theory and Practice of Electronic Governance (Guimaraes, Portugal, October 2022) 98–104.

L. De Graaf, M. Van Hulst, and A. Michels, "Enhancing Participation in Disadvantaged Urban Neighbourhoods," *Local Government Studies* 41: 1 (2015) 44–62.

J. Dortheimer, S. Yang, Q. Yang, and A. Sprecher, "Conceptual Architectural Design at Scale: A Case Study of Community Participation Using Crowdsourcing," *Buildings* 13: 1 (2023) 222.

J. Dortheimer et al., "Evaluating Large-Language-Model Chatbots to Engage Communities in Large-Scale Design Projects," Artificial Intelligence for Engineering Design, Analysis and Manufacturing 38 (2024) e4.

J. Du, X. Ye, P. Jankowski, T. W. Sanchez, and G. Mai, "Artificial Intelligence Enabled Participatory Planning: A Review," *International Journal of Urban Sciences* 28: 2 (2024) 183–210.

C. Durose, "Revisiting Lipsky: Front-Line Work in UK Local Governance," *Political Studies* 59: 4 (2011) 978–995.

S. S. Fainstein, "New Directions in Planning Theory," *Urban Affairs Review* 35: 4 (2000) 451–478.

J. Forester, *The Deliberative Practitioner: Encouraging Participatory Planning Processes* (Cambridge, MA: MIT Press, 1999).
X. Fu, C. Li, S. J. Quan, T. Yigitcanlar, and D. Wasserman, "Large Language Models in Urban Planning," *Nature Cities* 2: 7 (2025) 585–592.
X. Fu, C. Brinkley, T. W. Sanchez, and C. Li, "Text Mining Public Feedback on Urban Densification Plan Change in Hamilton, New Zealand," *Environment and Planning B* 52: 3 (2025) 646–666.
A. Fung, "Varieties of Participation in Complex Governance," *Public Administration Review* 66: s1 (2006) 66–75.
G. P. Griffin and J. Jiao, "The Geography and Equity of Crowdsourced Public Participation for Active Transportation Planning," *Transportation Research Record: Journal of the Transportation Research Board* 2673: 1 (2019) 460–468.
P. Healey, *Collaborative Planning* (London: Macmillan Education UK, 1997).
R. G. Hollands, "Critical Interventions into the Corporate Smart City," *Cambridge Journal of Regions, Economy and Society* 8: 1 (2015) 61–77.
T. Jacobsen, "Primitive Democracy in Ancient Mesopotamia," *Journal of Near Eastern Studies* 2: 3 (1943) 159–172.
P. Jankowski, K. Forss, M. Czepkiewicz, H. Saarikoski, and M. Kahila, "Assessing Impacts of PPGIS on Urban Land Use Planning: Evidence from Finland and Poland," *European Planning Studies* 30: 8 (2022) 1529–1548.
S. Joss, M. Cook, and Y. Dayot, "Smart Cities: Towards a New Citizenship Regime? A Discourse Analysis of the British Smart City Standard," *Journal of Urban Technology* 24: 4 (2017) 29–49.
R. K. Kummitha, "Smart City Governance: Assessing Modes of Active Citizen Engagement," *Regional Studies* 59: 1 (2025) 2399262.
L. Laurian and M. M. Shaw, "Evaluation of Public Participation: The Practices of Certified Planners," *Journal of Planning Education and Research* 28: 3 (2009) 293–309.
Z. Li, L. Xia, J. Tang, Y. Xu, L. Shi, L. Xia, D. Yin, and C. Huang, "UrbanGPT: Spatio-Temporal Large Language Models," paper presented at the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Barcelona, Spain, August 2024) 5351–5362.
K. Liu, T. Yigitcanlar, R. Mehmood, J. Corchado, and X. Fu, "Large Language Models in Urban Planning: A Systematic Review and Conceptual Framework," *Journal of Urban Technology* (2025) 1–44. <https://www.tandfonline.com/doi/full/10.108010630732.2025.2556551> Accessed January 18, 2026.
M. Manderscheid, V. Fiala, F. Edwards, B. Freyer, and I. Saumel, "Let's Do It Online?! Challenges and Lessons for Inclusive Virtual Participation," *Frontiers in Sustainable Food Systems* 6 (2022).
Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations," *Science* 366: 6464 (2019) 447–453.
G. Rowe and L. J. Frewer, "A Typology of Public Engagement Mechanisms," *Science, Technology, & Human Values* 30: 2 (2005) 251–290.
R. Schroeter and K. Houghton, "Neo-Planning: Location-Based Social Media to Engage Australia's New Digital Locals," *Australian Planner* 48: 3 (2011) 191–202.
A. Segura-Tinoco, A. Holgado-Sanchez, I. Cantador, M. E. Cortes-Cediel, and M. P. Rodriguez Bolivar, "A Conversational Agent for Argument-Driven E-Participation," paper presented at DG.O 2022: The 23rd Annual International Conference on Digital Government Research (Virtual Event, Republic of Korea, June 2022) pp. 191–205.
S. Senadheera, T. Yigitcanlar, K. C. Desouza, K. Mossberger, J. Corchado, R. Mehmood, R. Y. M. Li, and P. H. Cheong, "Understanding Chatbot Adoption in Local Governments: A Review and Framework," *Journal of Urban Technology* 32: 3 (2024) 35–69.
R. Simmons and J. Birchall,: "A Joined-Up Approach to User Participation in Public Services: Strengthening the 'Participation Chain,'" *Social Policy & Administration* 39: 3 (2005) 260–283.
M. H. Tessler, M. A. Bakker, D. Jarrett, H. Sheahan, M. J. Chadwick, R. Koster, G. Evans, L. Campbell-Gillingham, T. Collins, D. C. Parkes, M. Botvinick, and C. Summerfield, "AI Can Help Humans Find Common Ground in Democratic Deliberation," *Science* 386: 6719 (2024) eadq2852.

P. Thijssen and W. Van Dooren, "Who You Are/Where You Live: Do Neighbourhood Characteristics Explain Co-Production?" *International Review of Administrative Sciences* 82: 1 (2016) 88–109.

R. Thoneick, "Integrating Online and Onsite Participation in Urban Planning: Assessment of a Digital Participation System," *International Journal of E-Planning Research* 10: 1 (2021) 1–20.

C. J. A. Van Eijk and T. P. S. Steen, "Why People Co-Produce: Analysing Citizens' Perceptions on Co-Planning Engagement in Health Care Services," *Public Management Review* 16: 3 (2014) 358–382.

S. Verba, K. L. Schlozman, and H. E. Brady, *Voice and Equality: Civic Voluntarism in American Politics* (Cambridge, MA: Harvard University Press, 1995).

S. Williams, S.: Beery, C. Conley, M. L. Evans, S. Garces, E. Gordon, N. Jacob, and E. Medina, "People-Powered Gen AI: Collaborating with Generative AI for Civic Engagement," An MIT Exploration of Generative AI (2024) <https://mit-genai.pubpub.org/pub/6uejzuox> Accessed September 3, 2024.

A. Wilson and M. Tewdwr-Jones, *Digital Participatory Planning: Citizen Engagement, Democracy, and Design* (New York: Routledge, 2021).

C. Wyborn, A. Datta, J. Montana, M. Ryan, P. Leith, B. Chaffin, C. Miller, and L. Van Kerkhoff, "Co-Producing Sustainability: Reordering the Governance of Science, Policy, and Practice," *Annual Review of Environment and Resources* 44: 1 (2019) 319–346.

S. Yang, J. Dortheimer, A. Sprecher, and Q. Yang, "When Design Workshops Meet Chatbots: Meaningful Participation at Scale?" *International Journal of Architectural Computing* 22: 2 (2024) 216–237.

L. Yun, S. Yun, and H. Xue, "Improving Citizen-Government Interactions with Generative Artificial Intelligence: Novel Human-Computer Interaction Strategies for Policy Understanding through Large Language Models," PLOS ONE 19: 12 (2024) pp. 1–18.

A. Zachrisson, T. Bjarstig, and K. Eckerberg, "When Public Officers Take the Lead in Collaborative Governance: To Confirm, Consult, Facilitate or Negotiate?" Scandinavian Journal of Public Administration 22: 4 (2018) pp. 21–44.

Z. Zhou, Y. Lin, and Y. Li, "Large Language Model Empowered Participatory Urban Planning," arXiv preprint arXiv:2402.01698 (2024) <https://arxiv.org/abs/2402.01698>.

Appendices

Chatbot Software.

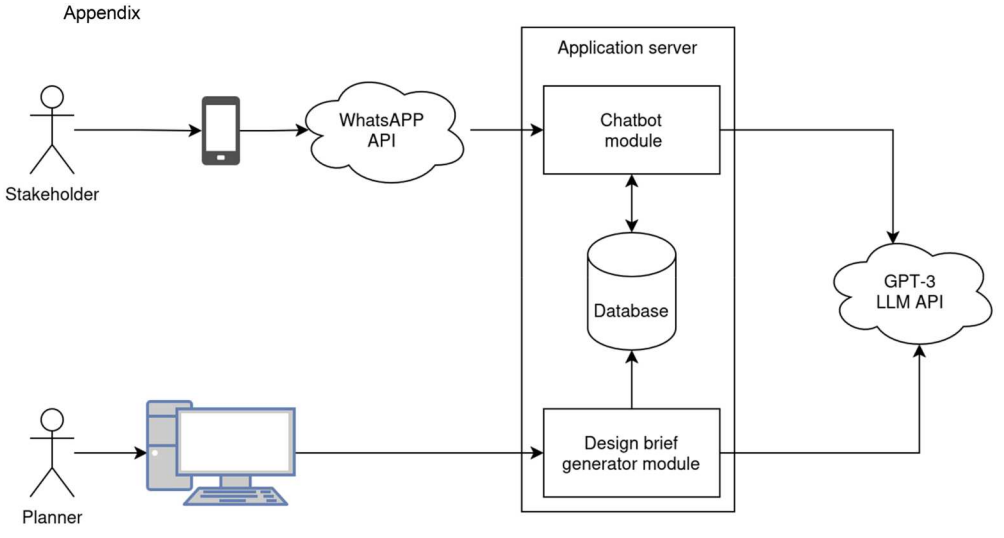

Figure A1. Software architecture overview: Chatbot and design brief generator

The chatbot application architecture is illustrated in Figure 5. The application received messages from the user through the WhatsApp application programming interface (API) and compiled a "prompt" sent to OpenAI's GPT-3 LLM API. A prompt is a text string that includes all the necessary information for the LLM to respond in the expected manner. The prompt consists of an internal prompt with all the chatbot's modifying settings, such as planning project data, persona, example questions, requested behavior, and tone. Then, a shared prompt is added, including the initial introduction messages the chatbot sends, which sets the stage for the interaction. This shared prompt is provided to the human participant and the LLM. Then, the prompt included all the conversation messages.

Furthermore, all the messages generated by the participants and the LLM were stored in a local database for further analysis. Figure 4 presents examples of performed human-machine interactions.

Since WhatsApp restricts sending messages outside of a 24-hour frame, it limits the ability of the chatbot to initiate conversations with existing users over the various phases of the planning process.

The planning team was interested in having the ability to initiate conversations with existing participants to engage with them as the planning process progressed.

To overcome this, the software implemented outgoing WhatsApp templates with fixed text messages approved by WhatsApp (Meta) to actively initiate a conversation with the participants.

Once a project phase was completed (out of four phases and two chatbots—CB1 and CB2), another software component analyzed all the conversations to summarize each conversation. This was done by compiling a specific prompt for the LLM, requesting the conversation to be summarized into a list of action points (the output). Later, all the conversation output lists were compiled into another prompt instructed to generate a design brief document.

During the project, the research team refined the prompts and validated and monitored the process, ensuring that the LLM did not omit data, which occasionally occurred.

Semi-Structured Interviews. Pre-Deployment Planner Interview (Pre-CB) consisted of eight open-ended questions. Respondents were also asked to provide basic professional background information, including profession, role in the strategic planning process, and years of experience in the field. The interview protocol explored:

(1) Existing tools and methods used to collect stakeholder input in regional strategic planning, including typical stakeholder groups
(2) Differences in participatory practices between small- and large-scale projects
(3) Prior experience with digital and computer-based participation tools (e.g., online meetings)
(4) Perceived limitations of current participation methods and self-assessment of existing practices relative to desired standards
(5) Critical and measurable objectives for effective input collection
(6) Understandings and practices of co-creation with communities
(7) Perceptions of chatbots and their potential utility in planning practice
(8) Expectations regarding the use of chatbot tools and their anticipated impact on planners' professional capacities

Post-First Chatbot Deployment Planner Interview (Post-CB1). This second semi-structured interview was conducted individually with each participating planner after completion of Phase I (Diagnostic Phase) of the planning process. This interview aimed to capture planners' reflections on the use of the chatbot as a tool for generating conversational input, strategic orientations, and preliminary design syntheses, and learning towards the design of CB2. The interview protocol explored five principal dimensions:

(1) Perception of the Tool
(2) Perception of Tool Performance in Diagnostic Functions
(3) Perception of Tool Performance in Synthesis and Integration
(4) Evaluation of Design Briefs
(5) Perception of Power Balance and Co-Creation.

Post-Second Chatbot Deployment Planner Interview (Post-CB2). A third semi-structured interview was conducted with participating planners following completion of Phase II of the planning process. It consisted of eleven open-ended questions and aimed to examine planners' perceived experiences with the chatbot in its role as a participatory voting and feedback instrument, as well as their broader reflections on the implications of AI-mediated participation for planning practice. The interview protocol explored three principal dimensions:

(1) Perception of the Tool and its Design for Phase 2
(2) Perceived Performance in Voting and Decision Support
(3) AI as a Capacity-Extending (Augmented) Planning Actor.