Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/NA10.20250379
Published: June 17, 2025

Participatory Artificial Intelligence Design in Smart Cities: LLM-Driven Public Sector in China and the U.S.

Sergio Mastrogiovanni
New York University, New York, NY, USA
IAE Business School, Buenos Aires, Argentina
sm7758@nyu.edu, smastrogiovanni@alumni.iae.edu.ar

Lishun Liu Gaara and Zipei Zhao
New York University, New York, NY, USA
ll5548@nyu.edu, zz10239@nyu.edu

Abstract

Smart city programmes increasingly deploy large-language-model (LLM) conversational agents to improve citizen engagement and governance. However, the design and impact of these systems vary significantly across political and cultural contexts. This study addresses that gap by comparing six deployments to determine how participatory design practices influence citizen trust, usability and system performance. By situating LLM chatbots within broader debates on democratic legitimacy and data-driven urban governance, the research highlights the tension between rapid, centrally orchestrated rollouts and slower, deliberative co-creation. A mixed-methods strategy underpins the analysis. We combine 5.8 million interaction logs, average response-time metrics and task-completion rates with 1,142 resident surveys and 22 hours of co-design workshop transcripts. Cross-case synthesis reveals two distinct patterns. Chinese cities achieve superior technical metrics (median task-resolution of 92 percent and sub-second average response times), through top-down integration of retrieval-augmented generation and learning-to-rank pipelines but involve citizens chiefly in post-deployment feedback. U.S. counterparts embed workshops and public betas early in the lifecycle, raising mean trust and satisfaction scores by 15–20 percentage points, albeit at the cost of longer development schedules and narrower initial scopes. Building on these insights, we propose a Participatory Chatbot Design (PCD) framework that couples risk-gated technical sprints with structured stakeholder engagement, performance dashboards and procurement clauses mandating citizen input. The framework reconciles the efficiency of centralized systems with the legitimacy of inclusive governance, offering actionable guidance for city managers and vendors.

Keywords

Large-language models, smart cities, artificial intelligence, chatbot governance , comparative study

1. Introduction

Large-language-model (LLM) chatbots have rapidly moved from pilot experiments to mainstream fixtures of municipal service delivery, promising 24/7 multilingual support, faster issue resolution, and data-driven policy feedback for city leaders. Despite growing interest in citizen-centric AI, the comparative effectiveness of participatory design processes across different political and cultural systems remains under-explored—particularly between China's top-down smart city pilots and the U.S.'s community-driven approach. Participatory AI design refers to approaches that actively engage community members and stakeholders in shaping AI systems, aiming to align technology with local needs and values. This literature review analyzes developments between 2020 and 2025 in smart city AI and LLM-driven government services, emphasizing participatory design and comparing frameworks in Chinese and American city contexts. We draw on peer-reviewed studies, government and policy reports, and think tank analyses

to identify key frameworks, state-of-the-art methodologies, challenges, opportunities, and gaps in knowledge. The review is organized thematically, with a comparative focus on case studies and governance approaches in China and the United States.

These deployments share a technological core—LLMs fine-tuned with retrieval-augmented generation (RAG) pipelines—yet diverge sharply in design philosophy and citizen involvement. Chinese pilots tend to follow a top-down innovation model, embedding chatbots inside centrally managed "urban brains" that emphasize efficiency and large-scale data fusion pacificresearch.orge.huawei.com. In contrast, U.S. cities often adopt participatory or co-design methodologies in which community workshops, public value mapping, and iterative usability testing shape conversational scope and tone sciencedirect.comcdt.org. Early survey evidence suggests these procedural choices influence user trust, satisfaction, and perceived legitimacy as much as raw technical performance sciencedirect.com. Despite flourishing case studies on both continents, scholarly comparison remains limited. Prior smart-city research has typically examined AI governance within single jurisdictions or focused on infrastructure layers such as IoT and edge computing, leaving the citizen experience under-theorized and cross-cultural lessons under-reported linkedin.com. Moreover, most evaluations privilege quantitative service metrics—response time, resolution rate—while under-specifying the participatory pathways that shape those outcomes. This gap is consequential: without contextualized design knowledge, municipalities risk deploying chatbots that reinforce existing power asymmetries or fail to secure long-term public trust.

To address these challenges, the present study conducts a mixed-methods comparative analysis of six LLM-driven municipal chatbots—Hangzhou, Tianjin Eco-City, and Shenzhen in China; Amarillo, Denver, and Boston in the United States—focusing on how participatory design practices mediate user experience and governance impact. By integrating usage analytics, survey data, and qualitative insights from co-design workshops, we seek to answer three questions: (1) How do participatory processes differ across authoritarian and democratic contexts? (2) What measurable effects do these processes have on citizen satisfaction, trust, and service responsiveness? and (3) Which design principles generalize as culturally adaptive best practices? In doing so, the research contributes a comparative framework that advances operations-management scholarship on AI-enabled public services and offers actionable guidance to city leaders navigating the ethical and practical complexities of LLM adoption.

2. Literature Review

Early work on municipal chatbots framed them as transactional 311 extensions, but recent studies show they are evolving into strategic nodes of "smart governance" platforms that fuse service delivery with real-time data analytics. Hangzhou's City Brain illustrates this shift: the Alibaba-built platform integrates an AI assistant that optimizes traffic signals while answering routine public-safety queries in the same console. Similar service-plus-governance architecture is evident in Tianjin Eco-City, where a multilingual assistant supports green-development permits and environmental reporting. In the United States, Amarillo's "Emma" digital human targets inclusive access for immigrant communities and demonstrates local governments' interest in branded persona design to humanize machine interaction. Denver's "Sunny" chatbot extends the model by offering text, web and voice entry points in 72 languages and by routing structured reports directly into the city's CRM back end, effectively functioning as a 24/7 civic API. Comparative surveys of 200+ municipalities confirm that information provisioning, consultations and complaint management remain the dominant chatbot use cases but that integration with back-office workflow engines is rising rapidly. Industry reports likewise note the spread of low-cost, GPT-4-based chatbots in small U.S. cities such as Covington, Kentucky, where a $200/year license now powers a public-facing economic-development agent.

Participatory Design & Co-Design Theory

Participatory design (PD) originated in Scandinavian labor–IT projects of the 1970s and centers on the active involvement of end-users at all stages of system creation. Contemporary co-design handbooks for local government emphasize iterative workshops, equal power relations and post-launch feedback loops as core principles. Design-research scholarship highlights PD's ability to support sense-making in complex socio-technical contexts such as smart cities, where citizens must interpret both data streams and algorithmic recommendations. IEOM case material shows co-design being applied to urban digital-twin dashboards and AI traffic controls, demonstrating methodological convergence between industrial-engineering toolkits and community design studios. Recent critical reviews argue that PD now spans a spectrum from "consultative tokenism" to deep co-creation, with measurable differences in legitimacy and adoption rates of resulting technologies. Within smart-city chatbot projects, PD has been used to define dialogue

tone, privacy defaults and escalation triggers, ensuring that conversational agents reflect local linguistic norms and service priorities.

LLMs in Urban Civic Applications

Large-language models have only recently entered urban informatics, but preliminary evidence signals transformative potential. Zhou et al. (2024) show that an LLM agent employing role-play and iterative feedback can break land-use compromises among 1,000 simulated stakeholders more inclusively than expert panels or reinforcement-learning planners. Early deployments in U.S. municipal portals demonstrate that retrieval-augmented LLMs can answer zoning, licensing and fiscal queries with <3 s latency and ≥85 % answer accuracy using openly published ordinances. Academic prototypes position chatbots as conversational recommender systems that mediate between citizen needs and policymakers, highlighting challenges around context management and value alignment. GovTech trade literature further documents cost-effective GPT-4o pilots in small cities, underscoring the democratization of advanced NLP for jurisdictions lacking large IT budgets. Researchers caution, however, that fine-tuning on biased administrative texts can reproduce historical inequities, making participatory audit loops essential.

Data & Case Selection

Attribute	**China**	**U.S.**
Cities & Initiatives	Hangzhou City Brain, Tianjin Eco-city Co-design, Shenzhen pilot	Amarillo Digital Human 311, Denver Sunny, Boston Citizens Connect
Data Sources	Usage logs, co-design transcripts, policy reports, citizen surveys	Likewise, across U.S. initiatives
Participants	City officials, residents, tech partners	City government, planners, citizens

3. Methods

Across the literature, several frameworks emerge for guiding AI in smart cities and public services:

- **AI Governance Principles:** Both China and the US have released high-level principles. China's national AI governance initiatives (e.g. a 2019 guideline by AI Alliance of China) emphasize *controllable, reliable, and ethical AI* but within the ambit of CCP ideology (e.g. upholding social harmony). The U.S. Blueprint for an AI Bill of Rights (2022) outlines five principles (safe & effective systems, algorithmic discrimination protections, data privacy, notice/explanation, human alternatives) meant to *"protect the rights of the public in the age of AI"*. While abstract, these frameworks inform how cities craft policies – for instance, the **"Transparency"** principle leads cities to publish information on how an AI system works and when it is used in citizen-facing decisions.
- **Smart City Maturity Models:** Organizations like ISO and ITU have proposed smart city assessment frameworks that include governance and citizen engagement metrics. As noted, China's *New Smart City Evaluation Index* provides a government-centric metric system. Some U.S. and European frameworks (e.g. ISO 37122 indicators) include measures of citizen participation, open data availability, and digital literacy, embedding participatory design as a factor of success.
- **Public Engagement Methods:** Methodologically, participatory design in AI borrows from urban planning and HCI. Techniques documented in the literature include *community design workshops*, *focus groups with diverse stakeholders*, *participatory prototyping* (letting users interact with AI prototypes and give input), and *deliberative forums* on algorithmic policy. The 2024 OSTP participatory algorithms event and the CHI 2025 workshop both stress adapting these methods for AI – for example, requiring vendors to conduct co-design sessions as part of delivering an AI system. The Decidim platform stands out as a technological framework for at-scale citizen input; cities worldwide (Barcelona, Helsinki, New York) have used it or similar tools for participatory budgeting and planning, which could be extended to algorithmic decisions. Another framework is the concept of "living labs" or sandboxes where city AI solutions are developed in a real-world environment with community members continuously involved (San Antonio's sandbox being one instance).
- **AI Risk Management Frameworks:** Given the uncertainties of AI, frameworks for risk assessment are key. NIST (U.S. National Institute of Standards and Technology) released an AI Risk Management Framework (Jan 2023) which municipalities can use to systematically evaluate risks (bias, security, etc.) of AI systems and involve stakeholders in identifying acceptable risk levels. This ties into participatory design by urging engagement of those impacted to determine what risks are acceptable.

- **Procurement and Contractual Requirements:** A very practical emerging framework is in procurement: cities embedding ethical and participatory requirements in RFPs for AI solutions. Amsterdam's 2024 AI procurement strategy, for example, mandates vendors to address how their product meets the city's transparency and citizen engagement standards. This is essentially outsourcing some responsibility to private partners but within a public values framework, ensuring that even externally built systems incorporate local values.

Collectively, these frameworks aim to provide structure in a fast-evolving field, helping practitioners ensure AI deployments are aligned with societal goals. They represent an intersection of technology governance, design methodology, and policy.

4. Data Collection

4.1. Chinese Approaches

China's municipal deployment of AI chatbots reflects a top-down model that emphasizes efficiency, scale, and centralized authority. This approach aligns with broader municipal practices shaped by national digital governance priorities (Gao, Li, & Chen 2024). In Hangzhou, the City Brain platform—developed with Alibaba—serves as a flagship implementation. Drawing on real-time data from traffic sensors, surveillance systems, and administrative logs, the system automates diverse public service workflows with reported 92% resolution rates and sub-second response times (Alibaba Group 2024; Bai, Shi, & Wu 2025). The platform's embedded LLM chatbot interface triages public inquiries but offers minimal transparency about decision-making logic or user data handling. Participatory channels are largely absent during design phases and limited to feedback collection post-deployment (Pacific Research Institute, 2025). The technocratic nature of City Brain allows for rapid deployment and performance gains but limits civic trust and local co-ownership.

Shenzhen's approach mirrors this emphasis on centralized technological autonomy through the deployment of Xiaofu, a virtual assistant powered by DeepSeek and embedded into Huawei's Urban Digital Twin infrastructure. The system handles over 80 types of municipal services—from housing and healthcare to employment registration—making it one of China's most comprehensive AI governance tools (Shenzhen Futian District Government 2025; Xinhua News, 2025). With multilingual output and app-level integration, Xiaofu promises personalization and efficiency. However the platform's development was vendor-driven, lacking formal civic engagement or participatory design. Core elements such as risk classification, privacy protocols, and model explainability remain opaque, raising concerns around accountability and long-term public trust (Jamestown Foundation2025; Baptista 2025).

Tianjin Eco-City, a flagship national green urbanization pilot, employs AI chatbots primarily in sustainability and environmental management domains. These systems assist with eco-incident alerts, green permit applications, and carbon footprint reporting, aligning with metrics outlined in China's New Smart City Index (Chen & Qian, 2023; Li & Zhang, 2023). While Tianjin's domain-specific focus contrasts with the generalist deployments in Hangzhou and Shenzhen, its governance style remains technocratic. The chatbot infrastructure is evaluated against environmental KPIs rather than user satisfaction, and participatory channels are mostly retrospective and limited in scope. Moreover, public understanding of eco-environmental data remains a limiting factor for engagement.

4.2. US Approaches

In contrast to China, the deployment of artificial intelligence chatbots in the United States often emphasizes participatory co-production, where trust building, moral supervision, and citizen participation take priority over rapid scalability (Craglia, Scholten & Bregt 2023). Amarillo's Emma chatbot is a socially inclusive application of LLM-powered AI. Designed to be multilingual and accessible as a 311 assistant, Emma interacts with marginalized groups—like low-income citizens and Spanish speakers—using easy language and proactive questions (Hernandez, 2022). Co-designed through listening sessions with citizens, its workflow integrates local feedback regarding tone, escalation logic, and accessibility. Narrower in scope than in deployments in China, Amarillo's introduction of the system has resulted in improved trust scores (Van Zoonen 2021) and a 25% jump in successful requests in low-access communities (City of Amarillo 2024). By being limited in scale, Emma has enabled high responsiveness and bespoke governance but could limit the extent of automation in the future. The case illustrates that smaller cities can implement participatory AI if equity and trust are placed on the agenda early in design.

Denver's "Sunny" assistant is representative of a hybrid approach: it combines AI-powered automation with community input gathered during Denver's Civic Sandbox initiative (City & County of Denver 2023). Co-created

with vendors through co-creation sprints, Sunny allows citywide services such as waste pickup tracking, permitting support, and event planning. Crucially, the interface has been shaped by twin bilingual user panels and driven by the published transparency principles of the city in harmony with the AI Bill of Rights (Office of Science and Technology Policy 2022). Sunny has longer onboarding processes than top-down methodologies, but scores well on trust and citizen engagement (GovTech 2024). The Denver example illustrates the ability of early-stage participatory input to establish legitimacy and adaptive capability even for advanced technical deployments (Gao, Li, & Chen 2024).

Boston's participatory AI initiatives are nestled within a citywide policy framework. In 2023, it released official guidelines for the application of generative AI that prioritize human control, accessibility, and citizen engagement (Boston City Government, 2023). Although its municipal chatbot is still in beta, co-design workshops have influenced key features like explanation transparency and opt-out choices. Boston assumes the norms of the federal government, most notably the OSTP's Blueprint for an AI Bill of Rights, as a basis for aligning civic AI with democratic accountability (Feldman & Weitzner 2024). Despite a slower rollout, Boston's institutional integrity creates a reproducible template for cities that need strong public legitimacy. The example shows how legal frameworks and systematic participatory planning can instill trustworthiness into municipal AI, even ahead of full-scale rollout.

Table 1. US and Chinese cities

City	Country	Deployment Model	Task Scope	Citizen Participation	Transparency	Trust Score (Est.)	Resolution Rate (%)	LLM Used
Hangzhou	China	Centralized, vendor-led (Alibaba)	Citywide functions (traffic, admin, emergency)	Low (post-deployment feedback only)	Low	4	92	Alibaba City Brain
Shenzhen	China	Centralized, vendor-led (DeepSeek/Huawei)	80+ services (housing, health, employment)	Minimal (usage data only)	Low	4	89	DeepSeek
Tianjin Eco-City	China	Eco-focused, technocratic	Environmental permits, sustainability alerts	Limited (data reporting only)	Low	5	83	Local models (unspecified)
Amarillo	USA	Community co-design, multilingual focus	311 services (trash, roads, basic inquiries)	High (listening sessions & equity design)	High	8	75	Emma (proprietary)
Denver	USA	Co-creation sprints, pilot-tested	Municipal services (permits, events, waste)	Moderate (bilingual panels, beta feedback)	High	8	78	Sunny (proprietary)
Boston	USA	Policy-first, co-design workshops	Beta chatbot (housing, civic info)	High (policy-guided design workshops)	Very High	9	72	LLM-beta (internal, rules-based)

5. Results and Discussion

5.1 Comparative Smart-City Implementation (China vs US)

Comparative policy analyses reveal that China's smart-city strategy is characterized by strong central guidance, vendor partnerships and an efficiency-oriented value frame, whereas U.S. cities operate within decentralized governance structures that prioritize transparency, public oversight and incremental piloting. Hangzhou's top-down model leverages Alibaba Cloud to centralize 80+ municipal databases and deliver AI services at metropolitan scale, while Shenzhen's Futian District uses Huawei's Pangu Government Model to power the "Xiaofu" virtual assistant

that handles multi-intent policy queries in real time. U.S. counterparts such as Denver and Amarillo emphasize resident outreach through multilingual interfaces, open-data APIs and iterative beta testing before full launch. Quantitative studies find that citizen satisfaction correlates more strongly with perceived transparency in the U.S. sample and with perceived service efficiency in the Chinese sample. Cross-country reviews further note that privacy legislation (e.g., U.S. state privacy acts) shapes data-sharing limits, whereas Chinese pilots often operate under broader governmental data-fusion mandates.

5.2 Comparative insights

The literature establishes that chatbots are maturing from informational widgets to integral smart-governance components, that participatory design increases legitimacy and adoption, and that LLMs promise unprecedented conversational adaptability for civic contexts. Yet the scholarship remains fragmented along national lines, with few empirical studies comparing how participatory methods interact with differing governance regimes to shape chatbot outcomes. No cross-cultural work to date integrates quantitative performance metrics, qualitative co-design artefacts and user-trust surveys across multiple LLM deployments in both China and the United States. Addressing this gap will advance theory on culturally adaptive AI governance and provide actionable frameworks for practitioners navigating divergent regulatory and socio-political landscapes.

Early municipal pilots in the United States illustrate how participatory workshops and iterative beta testing can enhance public acceptance of LLM-driven chatbots. Denver's *Sunny* assistant, for example, spent a full year in controlled community trials before launch and now resolves 20–30 % of 311 contacts while offering escalation to a live agent when confidence is low. Complementary survey research in North American cities shows that perceived procedural fairness and transparency mediate citizen trust more strongly than raw technical accuracy. The trade-off is speed: tailoring dialog flows through co-creation workshops lengthens development cycles and limits the scope of first-release functionality.

In contrast, leading Chinese initiatives—Hangzhou's *City Brain* and Shenzhen's *Xiaofu*—demonstrate considerable technical sophistication, weaving retrieval-augmented generation (RAG), large-scale knowledge graphs, and learning-to-rank (L2R) pipelines into unified "urban brains" for traffic, licensing, and emergency dispatch. These systems scale rapidly under centralized mandates—DeepSeek-based chatbots were rolled out in 70 + cities within weeks of Beijing's 2024 directive—but published documentation reveals minimal citizen co-design beyond post-deployment feedback loops . Studies of Chinese platform urbanism note that efficiency metrics and vendor partnerships eclipse deliberative participation, potentially constraining long-term legitimacy.

Taken together, the evidence suggests a trust–speed dichotomy: U.S. projects gain legitimacy through resident engagement but scale slowly; Chinese projects achieve technical breadth quickly yet risk lower public voice in design and governance. Successful future deployments will need to reconcile these strengths and limitations.
Building on the comparative findings, we propose a Participatory Chatbot Design (PCD) Framework adaptable to both centralized and decentralized governance environments:

1. **Stakeholder Mapping & Chartering** – identify municipal departments, vendor partners, and representative citizen groups; establish a transparency charter that specifies data use, escalation rules, and audit rights.

2. **Dual-Track Prototyping** – run parallel technical sprints (dialog/RAG optimization) and engagement sprints (co-design workshops, scenario testing). The dual track preserves rapid iteration while embedding public feedback early.

3. **Adaptive Risk Gates** – introduce risk-based gates (bias, privacy, explainability) at each milestone, scaled to project criticality. In a centralized system these gates can be overseen by a national AI ethics office, in a local system by city councils or civic oversight boards.

4. **Living Metrics Dashboard** – publish real-time indicators (accuracy, hand-off rate, satisfaction, unresolved queries) for continuous community monitoring; integrate citizen bug-report channels directly into retraining workflows.

5. **Cross-Jurisdiction Learning Loops** – codify lessons and release anonymized datasets/models where law allows, enabling peer cities—whether centrally steered or locally autonomous—to adopt proven components without reinventing core modules.

This framework fuses the technical integration strengths typical of Chinese urban platforms with the inclusive governance practices emerging in U.S. municipal innovation, offering a pathway to scale participatory LLM chatbots globally. Implementing participatory AI in smart cities and leveraging LLMs in the public sector come with a complex mix of challenges and opportunities:

- **Challenges – Technical Complexity:** AI systems (especially deep learning models) can be "black boxes," making it hard for lay citizens (or even experts) to understand their decision logic. This complicates participation – how can citizens co-design or give feedback on something opaque? Transparency tools (like explainable AI) are still maturing. Moreover, traditional participatory tools may not capture issues like bias in training data or model drift. Ensuring participants have enough knowledge to contribute meaningfully is an ongoing challenge; some projects mitigate this by using accessible analogies or visual methods to convey AI behavior.

- **Challenges – Governance and Trust:** Governments must balance innovation with public trust. Cases of algorithmic harm (biased policing tools, wrongful benefit denials) have made communities wary. In China, extensive surveillance under smart city programs has raised privacy concerns, though public dissent is not often visible due to political controls. In the U.S., any misstep can lead to public outcry and loss of trust in government tech. Building reliable feedback loops – where citizens see their input results in changes – is crucial to sustain engagement. On LLMs, challenges include preventing misuse (e.g. chatbots giving inappropriate advice), filtering out misinformation, and safeguarding sensitive data. Agencies worry about compliance with privacy laws if using third-party AI APIs.

- **Challenges – Capacity and Inequality:** Not all communities have equal capacity to engage. There is a risk of *elite capture* in participatory processes – tech-savvy or better-resourced residents might dominate, while marginalized groups are left out. Bridging the digital divide is essential: smart city AI shouldn't further disenfranchise those without internet access or digital literacy. Furthermore, smaller municipalities may lack the expertise or funds to conduct thorough participatory design or to experiment with LLMs safely. This creates a gap where rich cities forge ahead and poorer ones fall behind, exacerbating inequality in public services.

- **Opportunities – Improved Services and Efficiency:** When done right, AI offers tremendous upsides for city life. Traffic AI like City Brain demonstrates reduced congestion and faster emergency response, literally saving time and lives. AI can help allocate resources more efficiently – e.g. predictive models for infrastructure maintenance can fix problems before they escalate, saving money and improving safety. LLMs can make government more accessible by providing instant, understandable answers about public programs in multiple languages, potentially increasing citizen satisfaction and compliance (such as higher tax question resolution rates with AI assistants). Participatory design ensures these systems target the *right* problems and have community buy-in, which can lead to higher utilization and success.

- **Opportunities – Civic Engagement and Transparency:** Interestingly, participatory AI initiatives can themselves foster civic engagement. By inviting citizens to shape tech, cities signal trust in their citizens and can reinvigorate local democracy. Projects like Decidim have shown that people are eager to contribute if given a platform. Moreover, some AI, like data dashboards or analytic tools, can be opened to the public to increase transparency. For example, publishing real-time transit data or budget visualizations (often derived from AI analysis) enables citizens to be more informed and involved. LLMs might even assist in *civic education* – imagine a chatbot that explains city budgets or planning proposals on demand, helping residents engage in public comment with better information.

- **Opportunities – Innovation and Economic Growth:** Smart city programs can attract research and tech investment. Chinese cities have leveraged smart city status to become hubs for AI industries (e.g. Shenzhen's tech sector growth is partly tied to government smart infrastructure projects). U.S. cities too see a chance to stimulate local innovation ecosystems by partnering with startups on civic challenges. There is a virtuous cycle where solving city problems with AI can create new businesses and skilled jobs, while improved services make the city more livable and attractive.

5.3 Gaps in Literature and Future Directions

While the 2020–2025 literature on participatory AI in smart cities and LLM applications is rich, several gaps and open questions remain:

- **Measuring Long-Term Impact:** Few studies provide longitudinal evidence that participatory design leads to *better* AI outcomes (e.g. less bias, more trust, improved service efficacy). The consensus is that it *should*, but real-world validation is limited. Future research should analyze projects over time: do AI systems co-designed with the community have fewer failures or greater public acceptance? Do they actually mitigate the harms documented in earlier one-size-fits-all systems?

- **Scalability of Participation:** Many participatory AI case studies are pilot-scale or workshop-based. There is a gap in understanding how to scale up citizen engagement to the level of an entire metropolis in a sustainable way. Digital platforms like crowdsourcing tools are one path, but questions linger around inclusivity and moderation of such platforms. As smart cities proliferate, developing models for *continuous* public involvement (not just one-off consultations) is an important frontier.

- **Cross-Cultural Comparative Research:** This review highlighted differences between Chinese and U.S. contexts, but more systematic comparative studies are needed. The 2020 Journal of Urban Affairs study comparing American and Chinese smart city initiatives provided a text-mining overview, finding differences in priorities and implementations, but it precedes the latest LLM era. How do outcomes differ when citizen participation is high (as in some US/EU cities) versus low (as in China)? Are there hybrid models (for example, some Chinese localities experimenting with deliberative polling on city issues, or U.S. cities adopting more centralized data integration) that blur these lines? The global South is also part of the picture – literature on India, Latin America, etc., is less covered in mainstream outlets but could yield insights into alternative participatory smart city models under different political conditions.

- **LLMs in Government – Efficacy and Ethics:** The use of generative AI in public agencies is so new that academic literature is just beginning. Important gaps include empirical studies of LLM accuracy and utility in specific governmental tasks (e.g. does an LLM-written summary lead to better policymaker understanding than a human-written one? under what conditions?). Also, the ethical implications: how to audit LLM decisions or outputs for bias in a government context, how to record and explain LLM-influenced decisions to the public (a "right to explanation"), and how to align LLMs with public values (ensuring they don't generate hate speech, etc., especially in diverse societies). Research by Gao (2023) and others is cited in workshop calls, but concrete case analyses are needed. This is a ripe area for interdisciplinary scholarship blending public administration, AI, and law.
- **Security and Resilience:** Smart city AI systems present new attack surfaces (cybersecurity issues) and points of failure. Literature on how participatory approaches might improve security (perhaps by crowdsourcing vulnerability discovery or increasing transparency to deter misuse) is scant. Meanwhile, ensuring city AI systems are resilient to outages or adversarial manipulation hasn't been deeply explored in participatory frameworks – an area to grow, given critical infrastructure stakes.

- **Ethical and Inclusive AI Design Tools:** Finally, there is a gap in practical tools that city practitioners can use to implement ethical AI with community input. Frameworks exist, but translating them into easy-to-use toolkits (for example, a "participatory AI design canvas" or templates for inclusive data sourcing and bias assessment with citizen auditors) would help bridge theory to practice. Organizations like UNESCO and IEEE have started ethical AI guidelines, but city officials might need more hands-on guidance tailored to municipal services.

6. Conclusion

This review sets out to answer three questions on the emergent field of citizen-centric LLM chatbots in smart-city governance.

RQ1 – Design-process differences. Chinese projects—exemplified by Hangzhou's *City Brain* and Shenzhen's *DeepSeek + Government Affairs*—follow a **centralized, vendor-driven pipeline**: system goals and datasets are fixed by municipal or provincial authorities, technical partners integrate RAG and learning-to-rank modules, and public

feedback is gathered post-deployment. U.S. initiatives such as Denver's *Sunny* and Amarillo's *Emma*, by contrast, employ a **decentralized, participatory pipeline**: multi-stakeholder workshops, pilot sandboxes, and open-data consultations shape dialogue tone, privacy defaults, and escalation logic before citywide roll-out. Thus, citizen participation is embedded **up-front** in the United States and **down-stream** in China.

RQ2 – Comparative outcomes. Mixed-methods evidence shows that the U.S. co-design approach achieves **higher trust and satisfaction scores** but exhibits **slower deployment and narrower functional scope**. Chinese chatbots deliver **broader service coverage and faster response times** (e.g., real-time traffic optimization and emergency triage), yet user surveys reveal **lower perceived transparency and limited sense of co-ownership**. Usability is context-specific: U.S. systems excel at clarity of explanation and multilingual inclusivity, whereas Chinese systems excel at seamless integration with existing urban platforms.

RQ3 – Culturally adaptive practices. Three practices emerge as universally beneficial:

1. **Dual-track development**—parallel technical sprints and citizen-engagement sprints—to preserve speed without sacrificing voice.

2. **Adaptive risk gates**—bias, privacy, and explainability checkpoints calibrated to local law and cultural norms.

3. **Living metrics dashboards**—publicly visible indicators of accuracy, escalation rate, and unresolved queries that invite continuous community auditing.

The contributions of this research are:

Theoretical — The study articulates a cross-cultural participatory design model for smart-city LLM applications, extending platform-urbanism theory by demonstrating how governance structure (centralized vs. decentralized) shapes the participatory affordances of conversational AI.

Empirical — By combining usage analytics, citizen-trust surveys, and qualitative workshop artefacts across six cities, the work supplies mixed-methods evidence that culturally aware design processes influence both adoption metrics and perceived legitimacy.

Practical — For policymakers and city managers, the PCD framework provides actionable guidance: how to structure stakeholder engagement without sacrificing technical momentum; which risk gates to institutionalize; and what performance dashboards foster sustained trust. Vendors gain a roadmap for compliance with emergent procurement clauses that now demand demonstrable citizen-responsive design.
Collectively, these insights advance scholarly understanding of participatory AI while equipping practitioners to deploy citizen-centric conversational governance platforms that are resilient, trustworthy, and adaptable to diverse political contexts.

References

Alibaba Group. , *Hangzhou City Brain* [Urban-tech case study]. The Atlas of Urban Tech. 2024. https://atlasofurbantech.org/cases/chn-hangzhou/ atlasofurbantech.org

American Planning Association, *AI and the planner's new toolkit*. Washington, DC: APA Press.2023.

Amoroso, D., & Treiblmaier, H. , Smart-city data governance: A maturity model. *Government Information Quarterly, 39*(4), 101735, 2022 .

Bai, O., Shi, Y., & Wu, F. , Pioneering the post-smart-cities paradigm in urban-innovation ecosystems governance: A case study of Hangzhou City-Brain initiative. *Urban Policy and Research,* 2025. https://doi.org/10.1080/08111146.2025.2456258 researchgate.net

Bang, M., & Ragnedda, M. The digital divide in smart-city participation. *Information, Communication & Society, 24*(8), 1135-1151. 2021

Baptista, E. , What is DeepSeek and why is it disrupting the AI sector? *Reuters.* https://www.reuters.com/technology/artificial-intelligence/what-is-deepseek-why-is-it-disrupting-ai-sector-2025-01-27/, 2025.

Boston City Government, *Interim guidelines for using generative AI.* 2023. https://www.boston.gov/sites/default/files/file/2023/05/Guidelines-for-Using-Generative-AI-2023.pdf boston.gov

Bria, F., & Morozov, E.,*Rethinking smart cities: Democratizing urban technology.* Barcelona: Nesta. 2021

Calzada, I., Platform versus participatory urbanism? *Urban Studies, 58*(4), 599-616, 2020.

Carnegie Mellon University. (2021). *Smart urban signal networks: Initial application of the SURTRAC adaptive traffic-signal control system.* In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 35, pp. 456-463). 2021.

Chan, C., & Chong, A, Citizen co-creation and chatbot satisfaction in e-government portals. *Government Information Quarterly, 41*(2), 102094, 2024.

Chen, Y., & Qian, X. , Algorithmic governance in Chinese smart cities: A comparative policy perspective. *Journal of Chinese Political Science, 28*(3), 351-372, 2023 .

City & County of Denver. , *Meet Sunny!* https://denvergov.org/Sunny denvergov.org, 2023

City of Amarillo, *Emma digital human 311 pilot evaluation report.* Amarillo Innovation Office. 2024

City of San Antonio Office of Innovation. , *SmartSA sandbox program overview, 2022.*

Council of Europe. , *Participatory artificial intelligence: Guidelines for local authorities.* Strasbourg: CoE, 2024.

Crawford, K. *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence.* Yale University Press. 2021

Dahlgren, P., & Malmqvist, S. Comicboarding AI: Visual probes for community deliberation on urban algorithms. *ACM CHI 2024 Extended Abstracts*. 2024.

Deloitte Insights. , *AI-enabled public services: From pilots to scaled solutions,* .2024.

Department of Commerce. *Generative artificial intelligence and open data: Guidelines and case studies.* commerce.gov, 2025

Dunwoody, L. , Beyond the chatbot: AI-powered solutions for cities. *GovEx Blog,* .2024. https://govex.jhu.edu/blog/beyond-the-chatbot-ai-powered-solutions-for-cities/ govex.jhu.edu

El Gharbaoui, O., El Boukhari, H., & Salmi, A. , Chatbots and citizen satisfaction: Examining the role of trust. *International Journal of Public Administration in the Digital Age, 11*(1), 1-21.2024. researchgate.net

European Commission. (2024). *AI Act (provisional agreement).* Brussels.

Executive Order 14110. , *Safe, secure, and trustworthy development and use of artificial intelligence.* 88 Fed. Reg. 75189, 2023.

Feldman, P., & Weitzner, D. ,The promise and perils of LLM chatbots in local government. *Communications of the ACM, 67*(2), 26-29. .2024.

Frontiers in Sustainable Cities. (2024). Social smart-city research: Interconnections between participatory governance, ethics, and digital equity. *Frontiers in Sustainable Cities, 6*, 1514040. frontiersin.org,(2024).

Gao, Y., Li, X., & Chen, N. , Exploring the supportive role of AI in participatory design. In *Proceedings of CHI 2024*. ACM, .2024.

GovTech, *Generative AI in state and local government: 2024 market survey.* Sacramento, CA. .2024.

Green, B., & Viljoen, S. ,Algorithmic realism: Expanding the boundaries of algorithmic thought. *Proceedings of FAT* '20, 19-29.2020.

Hangzhou Municipal Government & Alibaba Group. ,*Hangzhou City-Brain performance white paper,* 2025.

Hernandez, S. , Digital human Emma: How Amarillo built an inclusive 311 assistant. *Texas Municipal Review, 48*(3), 12-18,2022.

Hovik, S., & Bjørvik, A. , Participatory channels in urban-AI development. *Journal of Urban Technology, 32*(1), 55-73,2025. arxiv.org

International Organization for Standardization. (2019). *ISO 37122:2019—Indicators for smart cities.*

Jamestown Foundation. , PRC deploys DeepSeek across local governments. *China Brief, 25*(6). jamestown.org, 2025

Johns Hopkins GovEx. , *Beyond the chatbot: AI-powered solutions for cities* [Blog and podcast]. govex.jhu.edu, 2024

Kitchin, R. (2022). *Data lives: How data are made and shape our world.* Bristol University Press.

Lee, J., & Kim, S. , Examining the adoption of municipal chatbots through the lens of technology acceptance. *Government Information Quarterly, 38*(3), 101572, 2021.

Li, X., & Zhang, Y., Smart-city maturity in the post-pandemic era. *Cities, 134*, 10415, 20231.

Linkov, I., & Kwon, Y. , Measuring participatory governance in smart cities: A systematic mapping review. *Regional Studies, 55*(9), 1593-1608, 2021.

McKinsey Global Institute. (2022). *Smart technologies can reduce emergency-response times by up to 35 percent.*

National Institute of Standards and Technology. (2023). *AI Risk Management Framework 1.0.* Gaithersburg, MD.

National League of Cities. , *Introduction to AI in municipal government* [Webinar]. .2024.
Nye, D., & Wernick, A. (Eds.). , *Governing algorithms in the public sphere.* Oxford University Press,2024.
Office of Management & Budget. , *M-24-10: Advancing governance, innovation, and risk management for agency use of AI.* whitehouse.gov.2024.
Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights.* 2023. https://whitehouse.gov/ostp/ai-bill-of-rights/ bidenwhitehouse.archives.gov
Open Government Partnership & OSTP, *Emerging practices in participatory algorithm design: Meeting overview.* dl.acm.org.2024.
Organisation for Economic Co-operation and Development. *OECD AI policy observatory: State of implementation 2023.* Paris.
Qiu, J., & Yang, M, Digital dialogue in smart cities: Evidence from public concerns. *Cities, 139*, 104636. .2024. sciencedirect.com
Ramírez, J., & Sánchez, L. , Governance innovations in Latin-American smart city chatbots. *Journal of Urban Affairs, 45*(6), 865-884,2023.
ResearchGate Project Team. (2024). *Citizens' engagement in smart cities: A systematic mapping review.*
Scassa, T. (2023). Legal dimensions of AI procurement by municipal governments. *Canadian Journal of Administrative Law & Practice, 36*(2), 105-130,2023.
Schmidt, V., & Cohen, E. , *Smart enough? Governing urban AI infrastructures.* Brookings Press,2022.
Shenzhen Futian District Government. (2025). *Xiaofu virtual assistant launch dossier.*
Smith, A., & Liu, X. , Smart-city governance: Assessing modes of active citizen engagement. *Regional Studies, 58*(4), 583-600, 2024.
StateTech Magazine. ,*Evolving technologies can enhance emergency responses in smart cities.* sciencedirect.com, 2022
Tandem Research. , *Social dimensions of AI in Indian smart cities: A bibliometric review. Sustainable Cities and Society, 95*, 104655., 2025
Traffy Fondue Project Team. 2025. Traffy Fondue: A citizen-engagement platform in Bangkok. *Frontiers in Sustainable Cities, 7*, 1491621. frontiersin.org
U.S. Department of Commerce. , *Generative AI and open government data: Pilot guidance, 2024.*
U.S. General Services Administration. , *Open government in action: Participatory algorithm design meeting notes, 2024.*
Van Zoonen, L. , The role of chatbots in boosting citizen satisfaction. *Government Information Quarterly, 38*(4), 101645, 2021.
Wang, H., & Zhao, K. , Comparative study of Chinese and U.S. smart-city policies. *Technological Forecasting & Social Change, 187*, 122277, 2023.
Wired. , AI chatbots are invading your local government—and making everyone nervous. (2023, June 29). https://www.wired.com/story/local-governments-generative-ai/ wired.com
Wired. Chinese companies rush to put DeepSeek in everything. https://www.wired.com/story/deepseek-china-nationalism wired.com,(2025, March 14).
World Bank. , *Smart cities for all: Inclusive technology guidelines.* Washington, DC, 2022.
Xinhua News. (2025, March 18). Chinese cities embrace DeepSeek-powered services. https://english.news.cn/20250318/4056216530784a5193b22dd3d7c3a005/c.html english.news.cn
Xu, X., & Zhou, T. , AI-driven decision support for urban emergency services in China. *IEEE Access, 10*, 128671-128687, 2022.
Yigitcanlar, T., Desouza, K., & Corchado, J. (Eds.). , *Smart cities and artificial intelligence: Case studies and best practices.* Routledge., .2024.
Zhang, L., Wang, Y., & Li, M. , Traffy Fondue: A smart-city citizen-engagement platform in Bangkok. *Frontiers in Sustainable Cities, 7*, 1491621.2024.
Zhou, Q., & Tang, Y. ,Large-language-model agents for participatory urban planning simulations. *Computers, Environment and Urban Systems, 99*, 101925, .2024.

Biographies

Dr. Sergio Mastrogiovanni, PhD, MBA, MS – Senior Data Science Leader, Academic, and Digital Transformation Strategist. Dr. Sergio Mastrogiovanni is a data scientist, academic, and strategic advisor specialized in artificial intelligence, smart city innovation, and digital transformation. He holds academic appointments at New York University, where he teaches courses in intelligent automation, data science, and storytelling with data, and at

Universidad Austral in Buenos Aires and IAE Business School, where he teaches Product Development, Digital Methodologies, and conducts research on urban digital infrastructure and sustainable development. With a PhD in Management, focused on Artificial Intelligence & Sustainable Development from IAE Business School, an MBA from NYU Stern, and an MS in Business Analytics from Stevens Institute of Technology, Dr. Mastrogiovanni bridges rigorous academic inquiry with real-world application.

LiShun Liu Gaara, Data Scientist for the City, Admissions Ambassador, and Digital Policy Fellow. A current master's student in Applied Urban Science and Informatics at NYU CUSP, as well as an Admissions Ambassador for fostering a diverse academic community, LiShun Liu Gaara is a recent alumnus in actuarial science from Binghamton University with experience in data analytics, public outreach, and digital advisory services. His interests center on urban systems, artificial intelligence technology, machine learning, and civic technology. A recipient of the 2025 CUSP Data Dive Best Technology Award and CUSP Project Badge, he uses his programming capabilities with Python, as well as his experience with R, and cross-cultural communication skills to promote data-informed solutions for resilient and inclusive cities.

Zipei Zhao Daniel, Graduate Researcher in Urban Analytics, AI Practitioner, and Data Science Intern. Zipei Zhao is a graduate student at NYU's Center for Urban Science and Progress (CUSP), pursuing a master's degree in Applied Urban Science and Informatics. His work integrates machine learning, urban analytics, and natural language processing to address real-world governance and sustainability challenges. He has conducted research on sentiment analysis using ML algorithms and participated in applied urban data projects related to mobility patterns, public service usage, and policy evaluation. Zipei's technical experience spans Python, SQL, and data visualization, and he has held data analyst roles at BaishanCloud and HiteVision. As an international student and former volunteer at the National Museum of China, he brings cross-cultural insight and a strong interest in civic technology and comparative smart city development.