

Article

Not peer-reviewed version

Smart Cities in the Agentic AI Era: Three Vectors of Urban Transformation

[Esteve Almirall](#)*

Posted Date: 1 April 2026

doi: 10.20944/preprints202604.0004.v1

Keywords: smart cities; agentic AI; autonomous mobility; urban robotics; cognitive government; mirroring hypothesis; cumulative recursive hybridization; urban innovation ecosystems; chatbots; public sector



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Smart Cities in the Agentic AI Era: Three Vectors of Urban Transformation

Esteve Almirall

Center for Innovation in Cities, Esade Law & Business School (Ramon Llull University), Barcelona, Spain; esteve.almirall@esade.edu

Abstract

Agentic artificial intelligence—systems capable of reasoning, anticipating, and acting autonomously on behalf of citizens and institutions—is converging with electric and autonomous mobility and urban robotics to reshape how cities govern, move, and maintain their physical environments. This paper examines three interconnected vectors of AI-driven urban transformation: (1) the evolution of public-sector conversational AI from informational chatbots toward cognitive, agentic government; (2) the emergence of autonomous electric mobility—robotaxis, on-demand transit, and autonomous logistics—that is fundamentally altering urban spatial structure, cost, and connectivity; and (3) the deployment of intelligent robotics and city brain platforms that automate the physical management of urban space. We extend the mirroring hypothesis (Conway, Colfer and Baldwin) in two directions: dynamically, arguing that organizations and ecosystems converge toward the best strategic configurations that new technologies make possible; and ontologically, arguing that agentic AI introduces non-human agents as first-class participants in organizational architectures, requiring hybrid human-AI coordination structures. We further propose the concept of cumulative recursive hybridization—a dynamic in which the three vectors interact through data, regulatory, infrastructure, and talent feedback loops within specific urban ecosystems, generating compounding returns analogous to those observed during the Industrial Revolution. Drawing on comparative international evidence from over twenty governance chatbot deployments, the rapidly scaling autonomous mobility ecosystems of the United States and China, and emerging urban robotics landscapes, we find that advanced deployments concentrate in cities—not nations—that combine regulatory agility, talent ecosystem density, institutional willingness to redesign, and tolerance for experimental iteration. The paper concludes that the cities which will lead the next era of urban transformation are those that pursue simultaneous deployment across all three vectors, redesign their institutional architectures to mirror the possibilities of the agentic era, and actively orchestrate the cross-domain ecosystems in which cumulative innovation takes hold.

Keywords: smart cities; agentic AI; autonomous mobility; urban robotics; cognitive government; mirroring hypothesis; cumulative recursive hybridization; urban innovation ecosystems; chatbots; public sector

1. Introduction

Generative artificial intelligence (AI) is precipitating a step change in how societies organize work, coordinate decisions, and provision services. Since its public breakout in late 2022, adoption has progressed from experimentation to widespread use at a pace comparable to—yet faster than—previous general-purpose technologies such as electricity or the internet [1,2]. Crucially, these systems no longer merely classify or predict; they reason over open-ended tasks, draft and translate, synthesize evidence, and increasingly act within workflows, making them candidates for a new operating layer of governance, mobility, and urban infrastructure [3,4].

This inflection point is especially consequential for cities. Throughout history, cities have served as the primary arenas for social, economic, and technological experimentation—from the Greek polis

to the industrial metropolis to the contemporary smart city [5]. Each wave of general-purpose technology has reshaped the urban form: the railway redefined the Victorian city, the automobile created suburbia, and the internet enabled the networked metropolis. Agentic AI—systems capable not only of generating content but of learning, anticipating, and acting on behalf of citizens and institutions—promises a transformation of comparable magnitude, touching simultaneously how cities govern, how people and goods move, and how physical urban space is maintained and managed [6,7].

Yet the diffusion of AI across urban domains is neither automatic nor uniform. While many private organizations have already reorganized around AI-augmented workflows, public administrations face structural constraints—rigid hierarchies, procurement and hiring frictions, and incentive misalignments—that slow endogenous change [8]. At the same time, citizen expectations have been irreversibly reset by private digital services that deliver immediacy, personalization, and reliability [2,9]. The result is a widening gap between what is technologically possible and what public institutions are organizationally prepared to deliver. Closing this gap requires not simply “adding AI” to legacy systems, but re-architecting governance toward agentic models: interoperable, data-secure, auditable systems of human–AI collaboration that can learn, anticipate, and act on citizens’ behalf within democratically bounded mandates [10,11].

This paper argues that three interconnected vectors of AI-driven transformation are converging to reshape the smart city: (1) the evolution of public-sector chatbots and agentic systems that redefine the citizen–administration interface, moving from informational tools toward cognitive government; (2) the emergence of autonomous electric mobility—robotaxis, on-demand transit, and autonomous logistics—that is fundamentally altering urban spatial structure, connectivity, and cost; and (3) the deployment of intelligent robotics and urban infrastructure systems—from city brain platforms to maintenance robots and drones—that automate the physical management of the urban environment.

Crucially, we contend that these three vectors should not be understood in isolation. Drawing on the analogy of the Industrial Revolution in England, where steam power, mechanized textiles, iron production, and coal extraction cross-fertilized through geographic proximity and iterative recombination [12,13], we propose that the convergence of agentic governance, autonomous mobility, and urban robotics generates a process of *cumulative recursive hybridization* within cities. Autonomous vehicles produce urban data streams that feed intelligent traffic management systems; agentic administrations accelerate the licensing and regulation of new mobility services; robotic maintenance sustains the infrastructure on which both governance platforms and autonomous fleets depend. This cross-fertilization is inherently local: it requires the co-location of talent, regulatory authority, experimentation capacity, and institutional willingness to iterate—assets that are embedded in specific urban ecosystems rather than distributed evenly across national territories [14,15].

The paper thus advances two central propositions. First, building on the *mirror hypothesis* [10]—that organizations can only realize the potential of the technologies they adopt if they reconfigure themselves to reflect their possibilities—we argue that the city is the natural unit of analysis for AI-driven transformation, because it is at the municipal level that regulatory frameworks, service delivery, infrastructure management, and innovation ecosystems intersect. Second, we propose that cities that move early to integrate these three vectors will benefit from compounding returns analogous to those observed in historical innovation clusters, while those that hesitate—protecting incumbent structures or waiting for national directives—risk being relegated to adopting models designed, tested, and optimized elsewhere [16].

The remainder of the paper is organized as follows. Section 2 presents the conceptual framework, integrating the mirror hypothesis with theories of urban innovation ecosystems, intersectional recombination, and systems thinking to explain why cities are the natural locus of AI-driven transformation. Section 3 examines the first vector: the evolution of public-sector conversational AI, proposing a four-level maturity model and drawing on comparative international cases. Section 4 addresses the second vector: autonomous electric mobility and its implications for urban spatial structure, cost, availability, and policy. Section 5 explores the third vector: robotics and

intelligent infrastructure in the urban environment. Section 6 develops the cross-cutting argument—the city as locus of transformation—analyzing the dynamics of cumulative recursive hybridization, the local embedding of innovation ecosystems, and the conditions that separate pioneering cities from lagging ones. Section 7 concludes with implications for urban policy and directions for future research.

2. Conceptual Framework: Institutional Symmetry, Urban Innovation, and the Dynamics of Recombination

Understanding how agentic AI transforms cities requires more than a technology-adoption lens. It demands a conceptual framework that accounts for (a) the organizational conditions under which institutions capture value from new technologies, (b) the spatial dynamics that make cities uniquely fertile environments for innovation, and (c) the systemic mechanisms through which diverse technological domains cross-fertilize and compound. This section integrates three bodies of theory—the mirror hypothesis of institutional symmetry, the literature on urban innovation ecosystems, and the logic of intersectional recombination—to construct the analytical foundation that will guide the subsequent empirical sections.

2.1. The Mirroring Hypothesis Extended: From Communication Structures to Agentic Ecosystems

Our conceptual point of departure is Conway's foundational observation that "organizations which design systems are constrained to produce designs which are copies of the communication structures of these organizations" [17]. Conway's Law, as it came to be known, established a powerful principle: the architecture of a technical system reflects the social structure that produced it. Colfer and Baldwin [18] subsequently formalized this regularity as the *mirroring hypothesis*, demonstrating through a systematic review of empirical studies across industries that the alignment between organizational structure and technical architecture is not merely a constraint but an *equilibrium condition*—when the two are in symmetry, coordination costs fall and performance improves; when they diverge, friction and underperformance follow. Convergent findings in the technology-adoption literature reinforce this pattern: the "productivity paradox" of the 1980s–1990s showed that enterprise computing investments yielded returns only when firms restructured workflows and incentive systems in parallel [19,20]; Milgrom and Roberts theorized this as organizational complementarity [21]; and Brynjolfsson and colleagues demonstrated empirically that IT productivity requires co-investment in decentralized decision-making and workforce reorganization [22].

However, we argue that both Conway's original formulation and Colfer and Baldwin's empirical formalization, while foundational, describe a *static* version of mirroring—a tendency for existing structures to be reflected in existing designs. In this paper, we extend the mirroring hypothesis in two directions that are essential for understanding the impact of agentic AI on cities.

The first extension is dynamic and strategic. Organizations and ecosystems do not merely mirror their current communication structures in the systems they produce. When a genuinely new technology opens a previously unavailable space of possible configurations—new ways of coordinating, new divisions of labor, new modes of decision-making—organizations *explore* this space and tend to converge, through competitive pressure and institutional learning, toward the strategic structures that best exploit the technology's possibilities [23,24]. The mirroring, in this view, is not backward-looking (reflecting what exists) but forward-looking (converging toward what is now possible). The introduction of the internet did not simply mirror pre-existing corporate hierarchies into websites; it opened a space of networked, platform-based, and disintermediated organizational forms that firms progressively discovered and adopted—those that found the best configurations thrived, and those that clung to pre-internet structures were displaced [25]. The productivity paradox was, in this light, a period of exploration: the lag between technology availability and productivity gains reflected the time required for organizations to search,

experiment, and converge on the organizational architectures that mirrored the technology's actual possibilities rather than their inherited structures.

The second extension is ontological, and it is specific to the agentic AI era. In all prior technological transitions, the agents doing the coordinating, exploring, and decision-making within organizations were exclusively human. The communication structures that Conway described, the organizational units that Colfer and Baldwin mapped to technical modules, the complementary practices that Milgrom and Roberts theorized—all of these involved human actors coordinating with other human actors. Agentic AI fundamentally alters this premise. For the first time, the participating agents in an organizational ecosystem include not only humans but also artificial agents capable of autonomous reasoning, decision-making, and action [26,27]. An agentic chatbot that processes a citizen's permit application end-to-end is not a tool used by a human agent—it *is* an agent, with a defined role, a scope of authority, and interaction protocols with other agents (human and artificial) in the system. A fleet-management AI that coordinates robotaxis across a city is a coordination agent operating within the ecosystem's architecture. An urban intelligence platform that synthesizes data from autonomous vehicles, robotic sensors, and administrative systems to trigger maintenance workflows is an orchestrating agent that mediates the interactions of multiple subsystems.

This means that the organizational architecture that must mirror the technology is no longer composed solely of human roles, human communication channels, and human decision rights. It is a *hybrid architecture* of human and AI agents, where the roles, boundaries, interaction protocols, and coordination mechanisms must be designed for both kinds of participants [10,28]. The mirroring hypothesis, extended to the agentic era, thus states: *organizations and ecosystems will tend to converge toward the configuration of human and AI agents, coordination mechanisms, and decision architectures that best exploits the possibilities opened by agentic technologies—and the degree to which they achieve this alignment will determine the value they capture.*

We further extend this logic beyond the boundary of the individual organization to the level of the *city as an ecosystem*. A city is, in essential respects, an organizational architecture: it comprises governance structures, regulatory frameworks, service delivery systems, infrastructure networks, labor markets, and institutional cultures, all of which are interdependent and jointly shape the city's capacity to absorb, adapt, and benefit from new technologies. In the agentic AI era, a city's ecosystem includes not only its human actors—public servants, entrepreneurs, researchers, citizens—but also the growing population of AI agents that mediate, execute, and coordinate urban functions: administrative chatbots that handle citizen interactions, autonomous vehicles that navigate city streets, robotic systems that maintain infrastructure, and intelligence platforms that synthesize urban data. The city that achieves mirroring is the one whose institutional architecture—its regulatory agility, data-sharing frameworks, coordination protocols, and workforce capabilities—evolves in symmetry with this hybrid human-AI ecosystem. When this symmetry is present, the city captures compounding value from the recursive interaction of its three transformation vectors. When it is absent, the city falls into an *asymmetry trap*: powerful agentic technologies constrained by organizational structures designed for an exclusively human coordination logic.

For public administrations specifically, this extended mirroring hypothesis carries particular weight. Government agencies operate under constraints—legislative mandates, civil-service rules, procurement frameworks, accountability regimes—that make organizational reconfiguration slower and more politically fraught than in the private sector [8,29]. These constraints were designed for organizations composed entirely of human agents, and they do not easily accommodate AI agents that need defined roles, authority boundaries, audit trails, and accountability mechanisms of their own. This structural rigidity explains why most public-sector chatbot deployments worldwide remain at the informational level (Level 1), merely layering conversational interfaces onto unchanged back-end systems, rather than progressing toward the transactional and agentic architectures (Levels 3–4) that would require deep institutional redesign—including redesigning the organization to incorporate AI agents as legitimate participants in service delivery and decision-making [10,25]. The extended mirroring hypothesis thus predicts that the cities achieving the greatest impact from agentic

AI will be those that pursue not merely technological deployment but the co-evolution of their entire institutional ecosystem—human roles, AI agent roles, coordination architectures, and governance frameworks—toward the configurations that best exploit the possibilities of the agentic era. Figure 1 synthesizes the structure of the extended mirroring hypothesis, showing how Conway's foundational observation and Colfer and Baldwin's formalization are extended in the dynamic/strategic and ontological directions.

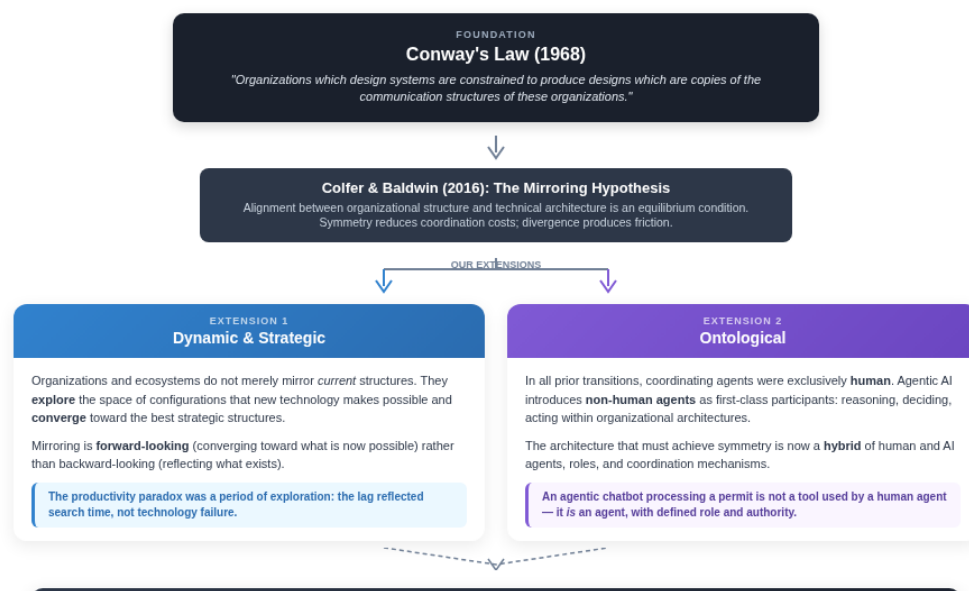


Figure 1. The extended mirroring hypothesis: from communication structures to agentic ecosystems. Conway's Law and Colfer and Baldwin's formalization are extended in two directions—dynamic/strategic (organizations converge toward best configurations new technology makes possible) and ontological (AI agents join humans as first-class participants in coordination architectures)—yielding a unified framework for the agentic era applied at the city level.

2.2. Cities as Innovation Ecosystems: Embedded Knowledge, Talent, and the Local Roots of Transformation

Why should the city, rather than the nation-state or the firm, be the primary unit of analysis for AI-driven transformation? The answer lies in a long tradition of scholarship on the geography of innovation, which consistently demonstrates that breakthrough innovation is not spatially neutral—it clusters in specific places, and those places are overwhelmingly cities [14,24,25].

Jane Jacobs argued half a century ago that cities are engines of economic life precisely because they bring together diverse activities in dense proximity, enabling the kind of unexpected cross-pollination that planned economies and isolated institutions cannot replicate [30]. Her insight anticipated what evolutionary economists would later formalize as *related variety*: the principle that innovation thrives not in environments of pure specialization (where everyone does the same thing) nor in environments of pure diversity (where there is no common language), but in ecosystems where distinct but complementary knowledge domains overlap and recombine [31,32].

This dynamic is sustained by several mechanisms that are intrinsically urban. First, *embedded knowledge*: the tacit know-how, institutional memory, and relational capital that accumulate in a place over time and cannot be easily transferred or replicated elsewhere [33]. The reason Silicon Valley dominates venture-backed technology, or that Shenzhen leads in hardware prototyping, is not simply a matter of policy incentives—it reflects decades of accumulated expertise, supplier networks, shared mental models, and trust architectures that are woven into the social fabric of the city [34]. Second, *talent density and circulation*: cities attract, develop, and retain specialized human capital, and the physical proximity of skilled workers in complementary domains accelerates the informal knowledge exchange—the corridor conversation, the accidental encounter at a conference, the engineer who

moves from a robotics startup to a transit authority—that seeds novel combinations [15,31]. Third, *fast supply chains and institutional proximity*: in a city, a startup developing autonomous delivery vehicles operates within reach of the municipal authority that issues permits, the university lab that tests sensor arrays, the logistics company that provides pilot routes, and the venture fund that supplies capital. This institutional proximity compresses the iteration cycle from idea to prototype to deployment in ways that remote or dispersed arrangements cannot match [37].

A critical and often underappreciated dimension of this argument is that knowledge and talent do not exist as abstract, portable assets—they are constituted through participation in ongoing projects and practices. Nonaka and Takeuchi's foundational work on organizational knowledge creation demonstrated that the most valuable forms of expertise—tacit knowledge—are generated through *socialization* and *externalization* within active work processes, not through formal training or documentation alone [38]. Lave and Wenger extended this insight through their concept of *communities of practice*, showing that expertise develops through *legitimate peripheral participation*: newcomers become skilled practitioners not by studying a body of knowledge but by progressively engaging in real, situated activity alongside experienced colleagues [39]. More recently, project-based learning theories have shown that in fast-moving technological domains, competence is inseparable from the projects in which it is exercised—engineers, data scientists, and urban planners develop frontier capabilities only when there are live deployments to work on, real problems to iterate against, and functioning teams embedded in ongoing operations [40,41].

The implication for cities is profound. A city that cancels its autonomous-bus pilot, shelve its chatbot initiative, or fails to launch robotics experiments does not merely lose time—it loses the very substrate in which talent forms. The engineers who would have learned to calibrate LiDAR in local traffic conditions, the public servants who would have developed expertise in AI-mediated service design, the urban planners who would have gained intuition for robotics-infrastructure integration—none of these capacities develop in the abstract. They develop in projects. And when the projects disappear, the talent either fails to form or migrates to cities where the projects exist. This creates a self-reinforcing dynamic: cities with active deployments attract and develop talent, which enables more ambitious deployments, which further deepens the talent pool—while cities without projects enter a vicious cycle of capacity erosion [15,43].

For the AI-driven transformation of cities, these dynamics are especially potent because the three vectors we examine—governance, mobility, and robotics—are all *place-based*: they are deployed in, regulated by, and experienced within specific urban territories. Unlike cloud software, which can be adopted from anywhere, an autonomous robotaxi fleet depends on local road infrastructure, local traffic regulation, local mapping data, and local citizen acceptance. An agentic public administration depends on local regulatory authority, local data registries, and local institutional culture. Urban robotics depends on local infrastructure conditions, local labor markets, and local procurement processes. The city is not merely the setting for these innovations—it is a constitutive element of their development.

2.3. The Medici Effect and Intersectional Innovation in the Urban Context

Frans Johansson's concept of the *Medici Effect* offers a complementary lens [43]. Johansson argued that the most transformative innovations in history have occurred at the *intersection* of disciplines, cultures, and domains—just as the Medici family's patronage in Renaissance Florence brought together sculptors, scientists, poets, financiers, and architects, creating an environment in which ideas from one domain sparked breakthroughs in another. The key insight is that innovation at intersections is qualitatively different from innovation within a single domain: it is less predictable, more combinatorial, and often more disruptive, because it escapes the path dependencies and mental models that constrain domain-specific thinking [43,44].

Applied to the smart city, the Medici Effect suggests that the most significant transformations will not come from advancing any single vector in isolation—a better chatbot, a faster robotaxi, a more efficient street-cleaning robot—but from the *intersections* among them. Consider: when an

agentic public administration can process autonomous-vehicle licensing in real time, it removes a bottleneck that currently delays deployment by months or years. When an autonomous logistics network shares data with a city brain platform, it enables predictive infrastructure maintenance that neither system could achieve alone. When robotic street maintenance operates on schedules optimized by the same AI that manages traffic flow, both systems become more efficient. These intersectional innovations are not planned from the top down; they emerge from the proximity, density, and diversity of the urban ecosystem—from the fact that the people building governance chatbots, the engineers deploying robotaxis, and the teams developing maintenance drones work in the same city, attend the same events, share the same infrastructure, and face the same regulatory environment.

This is why we argue that cities function as *innovation intersections* in the Johansson sense: they are the places where the Medici Effect operates at scale, because they concentrate the diversity of domains, the density of talent, and the institutional infrastructure needed for intersectional recombination to occur.

2.4. Systems Thinking and Cumulative Recursive Hybridization

The final element of our conceptual framework draws on systems thinking—the recognition that complex systems exhibit emergent properties that cannot be predicted from the behavior of their individual components [45,46]. In a systems perspective, the smart city is not a collection of independent technological applications (a chatbot here, a robotaxi there, a drone over there) but an interconnected system in which changes in one domain propagate through feedback loops to reshape others.

We synthesize the preceding theoretical elements into the concept of *cumulative recursive hybridization*: a dynamic in which multiple technological domains, co-located within an urban ecosystem, interact through iterative cycles of recombination, each cycle building on the outputs of the previous one and generating compounding returns. The term draws on the historiography of the Industrial Revolution, where scholars have shown that Britain's transformation was not driven by any single invention—not the steam engine alone, nor the spinning jenny, nor coke-smelted iron—but by the recursive interaction among them, each improvement in one domain enabling and demanding improvements in the others, within the geographic clusters of Lancashire, the Midlands, and the Scottish Lowlands [12,13,47].

The analogy is direct. In a city that adopts all three vectors simultaneously, the following feedback loops emerge: (a) *data loops*: autonomous vehicles and urban robots generate continuous streams of urban data—traffic patterns, road conditions, air quality, pedestrian flows—that feed city brain platforms and improve governance decision-making; (b) *regulatory loops*: agentic administrations capable of real-time licensing and adaptive regulation accelerate the deployment and iteration of mobility and robotics services; (c) *infrastructure loops*: robotic maintenance and intelligent infrastructure management improve the physical environment on which autonomous mobility depends; and (d) *talent loops*: the presence of frontier deployments in multiple domains attracts and retains the specialized workforce that sustains the ecosystem. These loops are recursive—each cycle increases the system's capacity for the next—and cumulative—the gains compound over time, creating increasing returns that widen the gap between pioneering cities and lagging ones [48].

This framework predicts that the cities which will lead the agentic AI transformation are not necessarily those with the largest budgets or the most advanced national policies, but those that cultivate the conditions for cross-domain recombination: institutional openness to experimentation, regulatory agility, talent density in complementary fields, and a willingness to allow—and learn from—failure. It also predicts that the divergence between leading and lagging cities will accelerate, as the compounding nature of recursive hybridization creates path dependencies that are difficult to reverse once established.

Section 6 will return to these dynamics with empirical evidence, examining how specific cities have—or have not—created the conditions for cumulative recursive hybridization to take hold.

3. AI and Public Administration: From Chatbots to Cognitive Government

The first vector of urban transformation concerns the interface between public administrations and citizens. Over the past decade, conversational systems in the public sector have evolved from rule-based bots and dynamic FAQs into sophisticated assistants powered by natural language processing and, increasingly, by generative AI [49,50]. This trajectory signals more than a technological refresh: it entails organizational, cultural, and social change in how administrations interface with citizens and execute core processes. Framed through the extended mirroring hypothesis developed in Section 2, the progression from informational chatbots to agentic governance systems represents a gradual—and still largely incomplete—alignment between the possibilities opened by AI technology and the organizational architectures of the public sector.

3.1. A Functional Maturity Model for Public-Sector Conversational AI

To anchor the analysis, we propose a four-level maturity model that maps the capabilities, organizational implications, and citizen value of public-sector conversational systems [10,49]. The model is cumulative: each level builds on the capabilities of the preceding one, but the organizational and institutional demands increase substantially at each stage.

Level 1 – Guided informational. At this foundational level, conversational systems reduce search frictions and standardize responses, typically complementing municipal websites without requiring complex back-end integrations. The system provides structured FAQs, basic triage, and links to forms. The organizational demand is minimal: content governance, editorial standards, and analytics for topic/intent distribution. The value proposition centers on 24/7 availability and consistency, deflecting repetitive contacts from telephone and counter channels. This is the modal implementation worldwide—the majority of public-sector chatbots currently operate at this level [49,50].

Level 2 – Contextual guidance (NLP/RAG). At this level, assistants move beyond generic answers to understand user intent, retrieve relevant regulation or documentation, and provide contextually adapted responses. Architecturally, this requires generative models operating over retrieval-augmented generation (RAG) pipelines, drawing on curated corpora of policy, regulation, and procedures with source-linked provenance [51]. The organizational demand rises considerably: knowledge curation pipelines, policy/legal review workflows, document versioning, and prompt/response evaluation become necessary. The gain is a marked reduction in ambiguity and error rates, and a more personalized citizen experience.

Level 3 – Assisted transactions. The qualitative leap occurs here: chatbots initiate or complete specific procedural steps—booking appointments, filing requests, processing payments, handling document submissions—integrating securely with transactional back-end systems [10,52]. This level requires API-first integration with line systems (registries, scheduling, payments), robust identity verification, explicit consent management, and auditable event logging. The organizational implications are profound: the chatbot is no longer a communication channel but a transaction operator, and its failures have direct consequences for citizens. Product-oriented operating models emerge—conversation design, service choreography, incident response, and security and privacy impact assessments become routine functions.

Level 4 – Cognitive–agentic (agents with data and action). At the apex of the model, conversational agents operate autonomously on behalf of citizens, orchestrating multi-system processes with traceability and safeguards [10,26]. The assistant does not merely respond to requests but proactively identifies eligibility, sends renewal warnings, coordinates across agencies, and executes multi-actor workflows under governed mandates. The architecture demands interoperable data spaces, policy-as-code for eligibility and safeguards, real-time analytics, and multimodal interfaces (text, voice, video, AR). This is, in the language of our extended mirroring hypothesis, the level at which AI agents become first-class participants in the organizational architecture of public

administration—not tools used by human agents, but agents in their own right, with defined roles, authority boundaries, and accountability mechanisms.

The transition across levels is not merely a function of better language models. Moving from Levels 1–2 to Levels 3–4 requires not only improved natural language understanding but also identity, consent, audit, and interoperability layers—institutional capabilities that translate model performance into public value with accountability [10]. Figure 2 summarizes the four levels, their architectural requirements, organizational demands, and representative international cases.



Figure 2. Four-level maturity model for public-sector conversational AI. Each level builds on the capabilities of the preceding one, but the organizational and institutional demands increase substantially. Representative city deployments illustrate the current international landscape at each level.

3.2. Comparative International Landscape

A comparative analysis of chatbot deployments across Europe, Latin America, and Asia reveals three empirical patterns that reinforce the maturity model and illuminate the dynamics of the extended mirroring hypothesis [10,49].

The dominance of Level 1. The majority of implementations worldwide remain at the guided informational level. Systems such as Línea Madrid (Spain), WienBot (Vienna), Bobbi (Berlin), TMBbot (Barcelona), Govbot (Japan), Divinha (Curitiba), and Jugalbandi (India) provide standardized information, reduce contact-center load, and offer basic navigation. They deliver quick wins in availability and consistency, but yield limited transformation when kept siloed from back-end systems. Their persistence at Level 1 reflects, in mirroring terms, an organizational architecture that has not yet reconfigured to match the possibilities of the technology.

The emergence of Level 2–3 as a contested frontier. A smaller but growing cohort of deployments has advanced to contextual guidance or assisted transactions. In Europe, Estonia's Bürokratt stands as the singular Level 4 reference—a platform-of-platforms architecture that integrates services across agencies and channels, including voice, toward a single national assistant. Clara (Madrid) and Noa (Île-de-France) operate at Level 2 with contextual retrieval capabilities. In Latin America, Buenos Aires' Boti is a leading Level 3 case on WhatsApp, handling over 300,000 interactions per month including appointment bookings, filings, and alerts; Mexico City's TEO focuses on anti-corruption reporting (Level 3); and Bogotá fields multiple systems at different levels (Chatico at Level 2, Rebeca at Level 1). In Asia, Dubai's Rammas supports billing and inquiries with multi-language capability and deep system integration (Level 3); Singapore's VICA centralizes multi-

agency assistance (Level 2); and South Korea's OneService Chatbot enables complaint filing and service reservations (Level 3) [10,49].

Regional patterns of strategic differentiation. Europe leads in the diversity of approaches, with cases spanning from informational to early agentic, and a notable concentration of experimentation in Spain (Las Rozas' Miguel, Ciudad Real's Prado with 80+ languages and 3,500 sources, the Open Administration of Catalonia's shared generative service). Latin America stands out for high adoption via WhatsApp—a channel strategy that meets citizens “where they are” and proves decisive for inclusion and scale. Asia drives large-scale, multi-service integration projects. Across all regions, the geographic concentration of advanced deployments in a few jurisdictions confirms that policy clarity, shared infrastructure, and institutional program management capacity are differentiators for moving from demonstration to institutionalization [10,49].

3.3. Three Scenarios for Implementation

Building on the maturity model and the comparative evidence, three implementation scenarios delineate the trajectories that local and regional administrations may follow. These are not mutually exclusive; a single jurisdiction may occupy different scenarios across policy domains and evolve asynchronously over time [10].

Conservative: incremental optimization. Chatbots act as guided informational interfaces that standardize answers and reduce search frictions without altering back-office structures. Maturity corresponds to Level 1 and selectively Level 2. Benefits include low cost and rapid deployment; limitations include minimal end-to-end resolution and potential citizen frustration when transactions are expected but unavailable. Illustrative cases include Vienna and Madrid's informational assistants.

Disruptive: assisted transactions and mass personalization. Chatbots become transaction operators that initiate or complete procedural steps, personalize interactions, and interoperate with multiple line systems in real time. Maturity is Level 3 with early Level 4 features. Cycle times shrink, staff effort shifts to higher-value work, and the citizen experience approaches that of private-sector digital services. The organizational prerequisites are substantial: robust data governance, cybersecurity, change management, and identity/consent layers. Representative cases include Boti (Buenos Aires), Rammas (Dubai), and elements of the Bürokratt platform [10,49].

Systemic: networked cognitive government. Conversational agents become components of an inter-organizational cognitive infrastructure where public agencies, private providers, and civil society co-produce seamless services. Maturity aligns with Level 4. Value arises from proactive, preventive administration—eligibility nudges, renewal warnings, multi-actor workflows—and ecosystem orchestration under stringent governance. The administration ceases to be a building that citizens must visit and becomes a network of agents that accompany them throughout their lives [49]. This scenario represents the fullest expression of the extended mirroring hypothesis: the organizational architecture of government has been redesigned to incorporate AI agents as legitimate participants, with coordination mechanisms, decision rights, and accountability structures that reflect the possibilities of the agentic era. National exemplars remain rare; Estonia's interoperable layer and Singapore's multi-ministry approach are the closest references.

3.4. Toward the Windowless Administration: The Agentic Governance Vision

The trajectory described above points toward a vision of public administration fundamentally different from the multi-channel bureaucracy that most citizens experience today. In this vision—what the Anteverti-Esade scenarios report terms the “windowless administration” [49]—the administration has no service windows, no websites in the traditional sense. Instead, citizens interact through conversational, multimodal interfaces that adapt to and anticipate their needs. The burden of navigating institutional complexity shifts from the resident to the system: interfaces become adaptive (text, voice, rich interactions), policy logic is encoded as executable rules that guide end-to-end processes, and a resilient orchestration layer routes each request to purpose-built agents that

consult authoritative sources, apply policy-as-code, and perform actions under role-based permissions [10].

This is not merely a user-interface improvement. It represents a fundamental reframing of government—from a channel-centric service provider to an agent-orchestrated system that learns from context, anticipates needs, and executes on citizens' behalf within clear democratic mandates. The operating layer of government changes: it becomes a fabric of human–AI teams that coordinate across agencies, data sources, and service partners with accountability and auditability [10].

Realizing this vision requires the co-evolution of technology and organizational architecture that the mirroring hypothesis predicts. Administrations that retain paper-era hierarchies while adding conversational front-ends will achieve, at best, a veneer of modernization without gains in responsiveness, coordination, or resilience. Those that redesign their structures—decision rights, roles, KPIs, coordination mechanisms—to incorporate AI agents as active nodes in traceable workflows will be the ones that deliver proactive, personalized services at scale. The practical test of progress, as the IJEBM analysis argues, is the movement from reactive case handling to preventive administration, measured through time-to-complete, first-contact resolution, and equity of access [10].

4. Autonomous Electric Mobility and the Reshaping of Urban Space

The second vector of urban transformation concerns how people and goods move through cities. Mobility is arguably the most consequential domain of AI-driven change, because cities are defined—physically, temporally, and economically—by their transport infrastructures [54,55]. The medieval city was the city of walking and the cart; the modern city is the city of the automobile. Each transport technology has not merely served the city but *shaped* it: determining its spatial extent, its density patterns, its connectivity, and its social geography. The convergence of electric propulsion and autonomous driving, mediated by AI, promises a transformation of comparable magnitude—one that will redefine the cost, availability, and spatial logic of urban mobility [54,56].

4.1. How Mobility Defines the City: Space, Time, and Cost

Three parameters have historically determined the form and functioning of cities through their transport systems [54]. The first is *space*: urban physiognomy and connectivity are dictated by the dominant mobility technology of each era. Railway lines created linear urban corridors; the automobile enabled low-density suburban sprawl; metro systems shaped radial density patterns around stations. The second is *time*: urban life is organized around an implicit constraint of approximately one hour of commuting. A city grows to the extent that its inhabitants can traverse it within this limit, and the effective size of the city is thus a function of the speed and reliability of its transport. The third is *cost and availability*: access to transport at different times, in different zones, and at different price points determines who can participate in urban economic and social life—and who is excluded.

For over a century, the cost per kilometer of private automobile transport has remained remarkably stable at approximately €0.40 in constant prices [54]. Electric propulsion disrupts this equilibrium by dramatically reducing energy and maintenance costs. Autonomous driving disrupts it further by eliminating the driver—the single largest cost component of any for-hire transport service, accounting for approximately 60% of total ride-hailing costs [57a,58a]. The combination of these two disruptions opens the possibility of individual on-demand transport at costs approaching €0.10 per kilometer, and substantially less for shared or collective autonomous services [54,59a]. This is not a marginal cost reduction; it is a structural shift that could fundamentally alter the economics of urban mobility.

4.2. The Autonomous Mobility Revolution: From Pilot to Scale

What was a speculative prospect five years ago is now an operational reality in a growing number of cities. Waymo, Alphabet's autonomous driving unit, currently provides over 500,000 paid robotaxi rides per week across ten U.S. cities—up from 175,000 at the start of 2025, representing a 157% increase in twelve months [57]. The company operates a fleet of over 3,000 robotaxis and is targeting one million weekly rides by the end of 2026. Its geographic footprint has expanded from initial deployments in Phoenix, San Francisco, and Los Angeles to include Austin, Atlanta, Miami, Dallas, Houston, San Antonio, and Orlando, with planned launches in Nashville, Las Vegas, San Diego, Detroit, Washington D.C., Seattle, and Denver [57,58]. Critically, Waymo has announced plans to begin operations in London in 2026 and has deployed test vehicles in Tokyo—marking the first moves toward international expansion beyond the United States [58].

In China, the scale of deployment is equally striking. Baidu's Apollo Go service operates over 1,000 robotaxis across 22 cities, including Beijing, Shanghai, Wuhan, and Shenzhen, completing over 17 million cumulative orders [59]. In Wuhan, where the largest single-city deployment operates, Apollo Go has achieved per-vehicle profitability—a milestone that demonstrates commercial viability at scale [59]. Weekly ride volumes reached 250,000 by late 2025, comparable to Waymo's volumes at that time [59]. Moreover, Baidu has announced agreements with Lyft to expand into Europe, with vehicles manufactured by Jiangling Motors expected to operate in Germany and Britain starting in 2026 [59]. Chinese competitors including Pony.ai, WeRide, and AutoX are pursuing parallel international strategies, with deployments or regulatory approvals in the UAE, Singapore, and South Korea [60].

The revolution extends well beyond robotaxis. **Autonomous buses** are entering regular public transit operations in Chinese cities. In August 2025, WeRide and Shenzhen Bus Group launched Shenzhen's first Level 4 fully driverless robobus line—the B888 route in Luohu District—connecting Luohu Port to MixC Market over a 6.6 km route, with vehicles equipped with over 20 sensors providing 360-degree perception [61]. This is not an isolated pilot: Beijing, Xiong'an New Area, Guangzhou, Changsha, Wuxi, Zhengzhou, Chongqing, and Hainan have all introduced autonomous shuttles on public roads, with technology provided by WeRide, Baidu Apollo, QCraft, and UISEE [61,62]. WeRide is partnering with bus manufacturer Yutong for global rollout, signaling that autonomous public transit—not just individual ride-hailing—is entering its scaling phase.

Autonomous delivery vehicles represent the third dimension of the mobility transformation, and China is again the leading laboratory. By mid-2025, over 15,000 autonomous delivery vehicles (ADVs)—mid-sized electric vans with two to ten cubic meters of storage handling payloads of up to a ton—were operating on roads nationwide, with approvals in over 200 Chinese cities [63,64]. Platform giants Meituan, JD.com, and Cainiao (Alibaba) have embedded driverless vehicles within their fulfillment systems, using them to replace human couriers on predictable last-mile routes. In Beijing's Shunyi district, Meituan operates a hybrid model where autonomous vehicles transport parcels to transfer stations and couriers complete the final hundred meters [63]. Manufacturer Neolix, which holds China's first official permit for autonomous delivery on public roads, has received approximately 30,000 global orders and now produces over 1,000 vehicles per month [64]. Two manufacturers command close to 90% of the professional autonomous delivery market, indicating rapid consolidation and industrial maturity [64].

The significance of autonomous delivery for cities is substantial: it reduces the number of individual human-driven delivery trips (a major source of urban congestion and emissions), enables consolidated multi-drop routing, and extends reliable logistics to areas and hours currently underserved. As the "Ciutats i IA" essay observes, the mobility transformation encompasses not only the movement of people but the movement of goods—and both dimensions are being reshaped simultaneously by the same convergence of electric propulsion and autonomous navigation [54].

The technology is thus no longer confined to controlled test environments. Across robotaxis, autonomous buses, and delivery vehicles, it is scaling commercially, expanding geographically, and entering the phase of mass deployment that will trigger the urban transformations discussed below.

4.3. Urban Implications: Availability, Equity, and Spatial Restructuring

The most radical impact of autonomous electric mobility may not be cost reduction but *availability transformation* [54]. Conventional public transport—buses, metro, taxis—concentrates service in central zones, high-density corridors, and peak hours. Peripheral areas, low-density neighborhoods, and off-peak hours are systematically underserved. This creates a geography of mobility inequality: residents of well-connected central neighborhoods enjoy frequent, affordable transport, while those in peripheral areas face long waits, limited routes, and high costs.

Autonomous vehicles are indifferent to time of day, demand density, and geographic peripherality. A robotaxi responds at 3 a.m. as readily as at 8 a.m.; it serves a suburban residential street as willingly as a downtown boulevard [54]. On-demand autonomous transit—minibuses and shuttles that dynamically adjust routes based on real-time demand—can provide frequent, flexible service to areas where fixed-route transit is economically unviable [65]. The result is a potential connectivity revolution: living far from the city center need no longer mean being disconnected from urban opportunities. In low-traffic conditions, peripheral locations could be within ten minutes of multiple urban destinations [54].

The implications for urban spatial structure are profound. If high-quality, low-cost transport becomes universally available regardless of location and time, the pressure on central housing markets could ease, as previously peripheral zones become viable residential alternatives. The chronic housing affordability crisis afflicting most major cities is, in significant part, a mobility crisis: housing is unaffordable in central locations because those are the only locations with adequate transport [54,66]. Autonomous electric mobility does not solve the housing problem directly, but it fundamentally changes the spatial equation by making connectivity less dependent on proximity to fixed infrastructure.

For goods movement, the impact is analogous. Autonomous delivery vehicles—small vans and pods operating on optimized multi-drop routes—can consolidate shipments, reduce the number of individual delivery trips, and provide service at lower cost and with greater availability than human-driven alternatives [54,67]. The proliferation of micromobility services (shared bicycles, scooters) is already reshaping last-mile logistics; autonomous systems will extend this logic to heavier loads and longer distances.

4.4. Two Policy Scenarios: The Decisive Role of Municipal Government

The realization of these benefits is not automatic. It depends critically on the policy choices made by municipal governments—a point that connects directly to the extended mirroring hypothesis and the argument developed in Section 2 [54].

Scenario A: Proactive municipal regulation. Cities that proactively license autonomous mobility services—robotaxis, on-demand transit, autonomous logistics—while managing the transition for incumbent sectors (taxi drivers, delivery workers, conventional transit operators) can steer the technology toward a city with far fewer private cars and a rich diversity of autonomous services, public and private, for both passengers and freight. In this scenario, private car ownership declines because the alternative—on-demand, low-cost, always-available autonomous transport—is simply more convenient and more affordable. Street space currently devoted to parking can be reclaimed for public use. Traffic volumes may decrease as shared autonomous vehicles replace single-occupant private cars [54,65].

Scenario B: Regulatory paralysis and incumbent protection. Cities that refuse to license autonomous services—whether to protect existing taxi collectives, delivery operators, or conventional transit—risk producing the opposite outcome. In this scenario, private citizens acquire their own autonomous vehicles, which instead of parking circulate empty or reposition to remote lots, returning when summoned. The result is *more* vehicles in circulation, *more* congestion, and a city environment that is worse, not better, than the status quo [54]. This is a powerful example of how regulatory choices designed to protect a specific group can end up harming the broader urban community—and it illustrates the mirroring hypothesis at the municipal level: a regulatory

architecture designed for human-driven transport, applied unchanged to autonomous technology, produces perverse outcomes because it fails to mirror the new technology's possibilities.

The cities at the forefront—San Francisco, Austin, Wuhan, Shenzhen—have chosen some variant of Scenario A, establishing regulatory frameworks that accommodate autonomous vehicles while imposing safety and data-sharing requirements [57,59,65b]. European cities remain largely in a pre-decision phase, though London's openness to Waymo and Germany's early regulatory frameworks for autonomous driving on public roads signal movement [58,67]. The decisive factor is municipal agency: national legislation provides the legal framework, but it is cities that issue operating permits, designate service zones, manage road infrastructure, and set the terms of engagement with autonomous mobility providers.

4.5. *The International Landscape: A Widening Gap*

The current geography of autonomous mobility deployment reveals a rapidly widening gap between pioneering and lagging cities. The leading cities are overwhelmingly in the United States (San Francisco, Phoenix, Los Angeles, Austin) and China (Wuhan, Shenzhen, Beijing, Shanghai), with emerging operations in the UAE (Dubai, Abu Dhabi) and early-stage preparations in Europe and Japan [57–60]. This concentration is not coincidental: it reflects the interplay of regulatory openness, technology ecosystem density, and institutional willingness to experiment that the conceptual framework in Section 2 identifies as the conditions for innovation clustering.

European cities, with notable exceptions, risk falling behind. Barcelona's experience is illustrative: an autonomous bus pilot was launched and then discontinued, and today the city has no operational autonomous mobility services [54]. As the "Ciutats i IA" analysis argues, if Barcelona—or any European city—wants to see how autonomous buses and robotaxis actually work at scale, it must currently look to China or California [54]. The gap is not merely technological; it is institutional and cultural, reflecting a regulatory and political environment that has been slower to accommodate autonomous mobility than its American and Asian counterparts. This gap, as the cumulative recursive hybridization framework predicts, will compound over time: cities with operational deployments develop data, expertise, regulatory know-how, and public acceptance that make subsequent iterations easier, while cities without them fall further behind with each passing year.

5. Robotics and Intelligent Urban Infrastructure

The third vector of urban transformation concerns the physical fabric of the city itself: how streets are cleaned, infrastructure is maintained, incidents are detected, and emergency responses are coordinated. While less visible to the general public than chatbots or robotaxis, the deployment of intelligent systems and robots in the management of urban space represents a profound shift—from reactive, labor-intensive maintenance to proactive, data-driven, and increasingly autonomous operation of the city as a physical system [54,68]. In the framework of the extended mirroring hypothesis, urban robotics and intelligent infrastructure constitute the domain where AI agents most directly take on roles previously reserved for human workers, fundamentally changing the coordination architecture of city operations.

5.1. *City Brain Systems: From Traffic Optimization to Urban Intelligence Platforms*

The most mature manifestation of AI in urban physical management is the "city brain"—an integrated platform that ingests data from thousands of sensors, cameras, and connected systems to monitor, analyze, and act upon urban conditions in real time. The concept was pioneered by Alibaba's ET City Brain, launched in Hangzhou in 2016, and has since been deployed in dozens of Chinese and Asian cities [69,70].

The results in Hangzhou are well documented: incident-detection accuracy exceeded 92%, average driving speeds increased by approximately 15%, daily commutes shortened by three minutes, and emergency response teams reached destinations seven minutes faster [69,70]. The city,

once ranked fifth among China's most congested, dropped to 57th on the national congestion index [69]. The system operates by integrating data from traffic lights, surveillance cameras, GPS signals from vehicles, and mobile phone location data, using AI to optimize signal timing, detect incidents, reroute traffic, and allocate emergency resources dynamically. City Brain has since expanded to Guangzhou (emergency service optimization), Suzhou (accident detection), and numerous other cities across China, as well as implementations in Kuala Lumpur and Macau [70,71].

What distinguishes city brain systems from conventional smart-city dashboards is their capacity for *autonomous action*, not merely monitoring. The system does not simply alert a human operator to a traffic jam; it adjusts signal timing to alleviate it. It does not merely display an accident report; it identifies the optimal ambulance route and dispatches resources. In mirroring hypothesis terms, city brain platforms represent an organizational architecture in which AI agents have been assigned operational decision rights within defined parameters—the system acts on behalf of the city within governed mandates, just as the agentic chatbots described in Section 3 act on behalf of citizens.

The evolution of city brain systems points toward comprehensive urban intelligence platforms that integrate not only traffic data but environmental monitoring (air quality, noise, temperature), infrastructure condition assessment (road surface, utility networks, building facades), public safety, and energy management. China's smart city AI is increasingly moving into environmental control, with systems that monitor and respond to pollution events, manage urban heat islands, and optimize energy distribution across city districts [72]. This integration creates the data substrate on which the other two vectors depend: autonomous vehicles require real-time traffic and road-condition data; agentic administration benefits from situational awareness of urban conditions for proactive service delivery.

5.2. Urban Service Robots: Street Cleaning, Maintenance, and Beyond

Below the scale of city-wide platforms, a growing fleet of specialized robots is taking on tasks that have traditionally required large numbers of human workers. The deployment is most advanced in Chinese cities, where fiscal constraints—municipal budgets are lean, and the imperative to maintain urban quality with limited resources is acute—have driven rapid adoption [54,68].

In Shenzhen's Shijing sub-district, 36 autonomous cleaning robots developed by Cowa Robot patrol an area of approximately 2.7 million square meters, sweeping streets, collecting waste, and operating continuously across day and night shifts [73]. Guangzhou has announced plans to increase its fleet of unmanned sanitation equipment to 1,000 units by 2026, and Hangzhou now explicitly requires the inclusion of unmanned equipment in new public sanitation tenders [73,74]. These are not experimental pilots—they represent institutionalized procurement decisions that embed autonomous robots into the routine operational architecture of city maintenance.

The logic extends beyond cleaning. Autonomous inspection robots monitor infrastructure conditions—bridges, tunnels, utility networks—detecting cracks, corrosion, and structural anomalies that human inspectors might miss or reach only at considerable cost and risk [68]. Robotic systems for vegetation management, road surface repair, and facade inspection are in various stages of deployment or advanced testing across Chinese, Japanese, and Korean cities [75].

5.3. Drones, Emergency Automation, and Urban Patrol

Drones add an aerial dimension to urban robotics. Their applications in cities span logistics (Walmart has completed over 20,000 drone deliveries across U.S. hubs and announced plans to expand coverage to 1.8 million additional households [76]), emergency response (drone-carried defibrillators, search-and-rescue in disaster scenarios), infrastructure inspection (power lines, rooftops, facades), and environmental monitoring (air quality sampling, flood mapping).

In the domain of urban security and patrol, Chinese cities have moved furthest. Chengdu deployed teams of robot police officers in June 2025, combining quadruped robots, wheeled robots, and humanoid robots to patrol city streets [77]. Hangzhou placed AI-powered traffic policing robots on active duty in December 2025, and in Wuhu, a humanoid officer designated Intelligent Police Unit

R001 oversees a busy junction, using cameras, speakers, and an AI system to detect cyclists and pedestrians in the wrong lane [77,78]. In Shenzhen, EngineAI's PM01 humanoid robots—standing 1.38 meters tall—patrol alongside human officers [77]. These deployments, while still limited in scale, signal a trajectory in which robotic agents share public space with citizens as visible components of urban governance infrastructure.

Shenzhen is being designed as China's first "robot-friendly" urban district, where robots will transition from closed training environments to open, on-street operation in neighborhood blocks [79]. This is a deliberate urban planning decision—the city is reconfiguring its physical and regulatory infrastructure to accommodate robotic agents as permanent inhabitants of public space, not merely as temporary experimental devices.

5.4. *The Integration Challenge: From Isolated Robots to Systemic Urban Intelligence*

The current state of urban robotics is characterized by fragmentation: cleaning robots, patrol robots, delivery drones, and city brain platforms typically operate as independent systems, managed by different agencies or contractors, with limited data sharing or coordination. The transformative potential lies in integration—and this is where the cumulative recursive hybridization framework developed in Section 2 becomes directly operative.

Consider a fully integrated scenario: city brain platforms ingest data from autonomous vehicles, delivery robots, cleaning machines, patrol drones, and infrastructure sensors to build a real-time model of urban conditions. Cleaning robots are dispatched to areas identified as high-priority based on pedestrian flow data from the mobility network. Infrastructure maintenance robots are routed to locations flagged by the sensors embedded in autonomous vehicles' road-surface detection systems. Emergency drones are pre-positioned based on predictive models fed by the city brain's incident-detection algorithms. Patrol robots share security data with traffic management systems to coordinate responses to accidents. Each system feeds data to the others; each becomes more effective because the others exist.

This integrated scenario does not yet exist anywhere in full. But the components are operational, and the cities that are building them—Shenzhen, Hangzhou, Guangzhou, Seoul—are creating the conditions for integration to emerge. The critical enabler is not any single technology but the *ecosystem density* that permits cross-system data flows, shared standards, and coordinated governance. This is, once again, a function of municipal agency: it is the city government that sets data-sharing protocols, defines interoperability standards, manages procurement to ensure compatibility, and creates the institutional architecture within which autonomous systems from different vendors and domains can interact.

For European and American cities, where urban robotics deployments remain more limited and more fragmented, the risk is not merely that they will lack individual robotic capabilities but that they will miss the systemic integration effects that arise when multiple autonomous systems operate within the same urban ecosystem. A cleaning robot in isolation is a labor-saving device. A cleaning robot connected to a city brain platform, sharing data with autonomous vehicles and coordinated with infrastructure inspection drones, is a node in an intelligent urban system—and the value of the node is a function of the network it belongs to.

6. The City as Locus of AI-Driven Transformation

The preceding sections have examined three vectors of AI-driven urban transformation—agentic governance, autonomous mobility, and urban robotics—each with its own maturity trajectory, international landscape, and institutional prerequisites. This section brings the argument together. Drawing on the conceptual framework developed in Section 2 and the empirical evidence assembled in Sections 3–5, we argue that the city is not merely the setting where these transformations unfold but the *generative locus* where they interact, recombine, and compound. The key to understanding which cities will lead the agentic AI era lies not in any single vector but in the

dynamics of cumulative recursive hybridization across all three—and in the local conditions that enable or inhibit that hybridization.

6.1. Cumulative Recursive Hybridization in Practice: How the Three Vectors Interact

The Industrial Revolution did not occur because steam engines improved, or because spinning jennies were invented, or because coke-smelted iron became available. It occurred because these innovations *interacted* within specific geographic clusters—Lancashire’s cotton towns, Birmingham’s metal trades, the Scottish Lowlands’ engineering workshops—each improvement in one domain enabling and demanding improvements in the others, in recursive cycles that compounded over decades [12,13,47]. The process was not planned; it emerged from the density, proximity, and diversity of the local ecosystem.

The same dynamic is beginning to emerge in the cities at the forefront of the AI-driven transformation, although the process is still in its early stages and the full integration remains aspirational. The interactions among the three vectors can be mapped through four feedback loops, as theorized in Section 2.4 and illustrated in Figure 3:

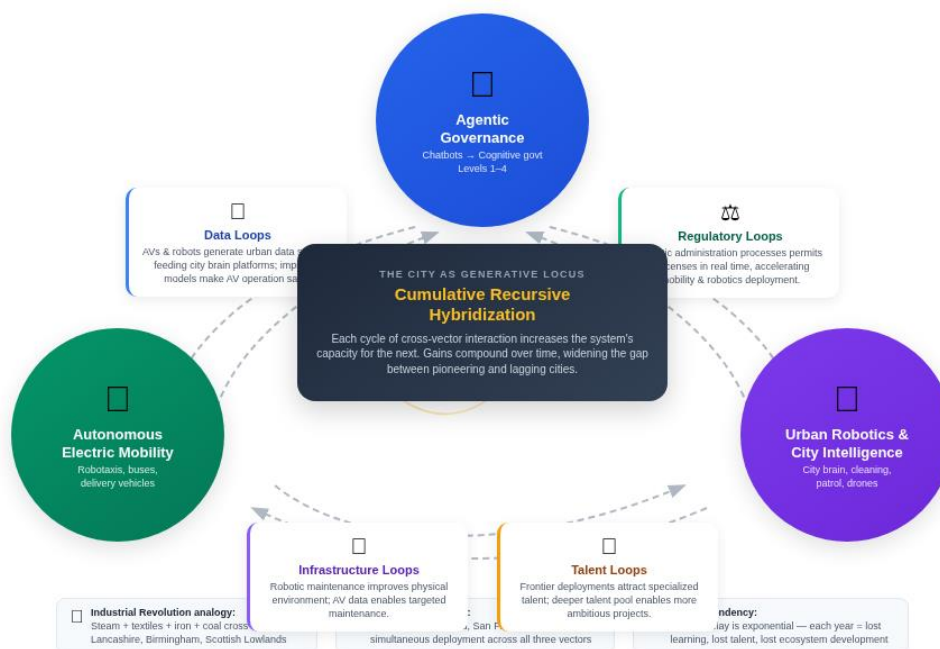


Figure 3. Cumulative recursive hybridization: four feedback loops connecting three vectors of urban AI transformation. The three vectors—agentic governance, autonomous electric mobility, and urban robotics—interact through data, regulatory, infrastructure, and talent loops within urban ecosystems. Each cycle increases the system’s capacity for the next, generating compounding returns that widen the gap between pioneering and lagging cities.

Data loops. Autonomous vehicles—robotaxis, buses, delivery vans—are mobile sensor platforms. As they navigate city streets, they continuously collect high-resolution data on road conditions, traffic patterns, pedestrian flows, air quality, and urban infrastructure state. This data feeds city brain platforms, improving their predictive models and operational decisions. In Hangzhou, the city brain’s traffic optimization depends on the density of data inputs; as autonomous vehicles proliferate, the volume, granularity, and freshness of data increase, making the system more effective [69,70]. Conversely, the city brain’s traffic optimization and incident-detection capabilities make autonomous vehicle operation safer and more efficient, closing the feedback loop. Delivery robots and cleaning machines contribute additional data streams—pavement conditions, waste accumulation patterns, micro-climate variations—that no single system would collect on its own.

Regulatory loops. Agentic governance systems capable of processing permits, licenses, and regulatory decisions in real time can dramatically accelerate the deployment of autonomous mobility and robotics. Today, a city that requires months of bureaucratic process to issue an operating license for a robotaxi fleet, or that lacks procedures for permitting cleaning robots on public sidewalks, creates bottlenecks that slow the entire ecosystem. An administration operating at Level 3 or 4 of the maturity model—with API-integrated transactional capabilities and agentic orchestration—could process such decisions in days or hours, adapting regulations dynamically as the technology evolves [10]. Shenzhen’s decision to design itself as a “robot-friendly” urban district [79] is an example of regulatory architecture that mirrors the possibilities of the technology: the city is not waiting for robots to arrive and then figuring out how to regulate them; it is proactively redesigning its institutional framework to accommodate robotic agents as permanent participants in urban life.

Infrastructure loops. Robotic maintenance systems—cleaning robots, inspection drones, road-surface monitoring—improve the physical infrastructure on which autonomous vehicles depend. Potholes, debris, and degraded road markings are among the most common causes of autonomous vehicle disengagement; a city that maintains its infrastructure proactively through robotic systems creates a more reliable operating environment for autonomous mobility [55,62]. In turn, the data generated by autonomous vehicles about infrastructure conditions enables more targeted and efficient robotic maintenance, creating a self-reinforcing cycle of improvement.

Talent loops. The most consequential feedback loop may be the least visible. As argued in Section 2.2, talent and knowledge are constituted through participation in active projects. A city that simultaneously deploys agentic governance systems, autonomous mobility services, and urban robotics creates a dense ecosystem of frontier projects across multiple AI domains. This attracts and develops specialized talent—AI engineers, data scientists, urban planners with robotics expertise, public servants skilled in AI-mediated service design—who would not develop these capabilities in a city without live deployments [15,39–42]. The talent, in turn, enables more ambitious projects, attracts investment, and deepens the ecosystem. San Francisco’s concentration of autonomous vehicle expertise is inseparable from the fact that Waymo, Cruise (now wound down), Zoox, and numerous startups operated there simultaneously, creating a labor market and knowledge commons that no single company could have generated alone [57,58]. Shenzhen’s emergence as a robotics capital reflects the same dynamic: the co-presence of DJI (drones), UBTECH (humanoid robots), Cowi (cleaning robots), WeRide (autonomous vehicles), and dozens of smaller firms creates an ecosystem in which engineers move between companies, ideas cross-pollinate, and the city’s collective capability compounds [73,77,79].

6.2. *Why Cities, Not Nations: The Local Embedding of Transformation*

A recurrent finding across all three vectors is that advanced deployments are concentrated in a small number of cities, not distributed evenly across national territories. Estonia’s Bürokratt is a national platform, but it is an exception; most governance chatbot innovation happens at the municipal level (Buenos Aires, Madrid, Dubai). Autonomous mobility is concentrated in specific cities (San Francisco, Wuhan, Shenzhen), not deployed uniformly across the United States or China. Urban robotics follows the same pattern (Shenzhen, Hangzhou, Guangzhou).

This concentration is not accidental. It reflects the fundamentally local nature of the conditions required for AI-driven urban transformation. National governments can provide legal frameworks, fund research, and set standards—but they cannot replicate the ecosystem density that makes transformation possible. The critical ingredients are:

Regulatory authority at the municipal level. It is the city that issues operating permits for robotaxis, sets zoning rules for robot-friendly districts, procures autonomous cleaning services, and decides whether its administration will deploy agentic chatbots. National legislation enables; municipal government acts [54,65].

Institutional proximity and fast iteration cycles. In a city like Shenzhen, the robotics startup developing a cleaning robot operates within kilometers of the municipal sanitation authority that will

procure it, the university lab that tests its sensors, the autonomous vehicle company whose road-condition data could optimize its routes, and the city brain platform that could coordinate its operations. This proximity compresses the cycle from idea to pilot to deployment to iteration in ways that geographically dispersed arrangements cannot match [37].

Embedded knowledge and communities of practice. As the theoretical framework in Section 2.2 established, expertise develops through participation in live projects. A city with active deployments across multiple AI domains develops communities of practice—networks of practitioners who share tacit knowledge, solve problems collaboratively, and build the institutional memory that makes subsequent deployments more effective [38,39]. This knowledge is embedded in the city’s social fabric: it travels when people change jobs within the same metropolitan area, but it does not easily transfer to distant cities that lack the project base in which to apply it.

The Medici Effect at urban scale. The intersectional innovations that arise when governance, mobility, and robotics professionals interact in the same urban ecosystem—the chance conversation between a chatbot developer and a robotaxi engineer that sparks an idea for real-time permit processing; the urban planner who realizes that cleaning robot data could inform housing policy—occur because diverse domains are co-located in dense proximity [43,44]. These innovations cannot be planned or mandated by national policy; they emerge from the ecosystem’s structure.

This analysis explains why the relevant unit of comparison for AI-driven urban transformation is not “the United States vs. China” or “Europe vs. Asia” but “San Francisco vs. Barcelona,” “Shenzhen vs. Berlin,” “Wuhan vs. Madrid.” The transformation is happening city by city, and the differences between leading and lagging cities within the same country (San Francisco vs. Detroit, Shenzhen vs. Chengdu) can be as large as the differences between countries.

6.3. From Experimentation to Systemic Transformation: What Distinguishes Leading Cities

Not every city that experiments achieves systemic transformation. Many cities have launched isolated pilots—an autonomous bus here, a chatbot there—without progressing to the integrated, cross-domain deployments that generate compounding returns. What distinguishes the cities that are advancing from those that are stalling?

The evidence from Sections 3–5 points to four differentiating factors:

Simultaneity across vectors. Cities that deploy across multiple vectors simultaneously—governance, mobility, and infrastructure—create the conditions for cross-fertilization. Shenzhen is deploying autonomous buses (WeRide), cleaning robots (Cowa), patrol robots (EngineAI), and building a robot-friendly urban district, all while its administration modernizes digital services. This simultaneity is not a coincidence; it reflects a municipal strategy of comprehensive AI adoption that creates the ecosystem density required for recursive hybridization. Cities that pursue one vector in isolation—a chatbot initiative here, an autonomous bus pilot there—miss the compounding effects.

Institutional willingness to redesign, not just adopt. The extended mirroring hypothesis predicts that value comes not from deploying technology onto existing structures but from co-evolving institutional architectures with technological capabilities. The cities that are pulling ahead—Estonia (national administration redesigned around Bürokratt), Shenzhen (urban districts redesigned for robotic agents), Wuhan (regulatory framework redesigned for autonomous vehicles)—have accepted that the institutional architecture itself must change. Cities that deploy AI while preserving inherited bureaucratic structures, incumbent protections, and legacy procurement processes remain in the asymmetry trap.

Tolerance for iteration and failure. The “Ciutats i IA” essay makes the point forcefully: there are no established best practices to copy in AI-driven urban transformation. The field is exploratory, and the path forward requires experimentation, evaluation, and continuous adjustment [54]. Cities that accept failure as a component of learning—that launch pilots, evaluate results, iterate, and scale what works—develop adaptive capacity. Cities that demand guaranteed outcomes before acting, or that cancel pilots at the first sign of difficulty (as Barcelona did with its autonomous bus), erode their own capacity to learn and improve.

Public-sector entrepreneurship. In every leading case, municipal government has played an active, entrepreneurial role—not merely as regulator or procurer but as ecosystem orchestrator. Hangzhou’s city government co-developed City Brain with Alibaba. Shenzhen’s government is designing robot-friendly districts. Buenos Aires’ government deployed Boti on WhatsApp as a deliberate channel strategy. These are acts of institutional entrepreneurship: public officials taking initiative, accepting risk, and shaping the trajectory of technological adoption in their cities. The role of the public sector is not to step aside and let technology companies operate freely, nor to regulate restrictively and slow adoption, but to actively co-create the conditions for transformation—setting standards, facilitating integration, managing transitions for affected workers and communities, and ensuring that the benefits are broadly shared.

6.4. *The Divergence Ahead: Path Dependency and the Widening Gap*

The cumulative recursive hybridization framework predicts that the gap between leading and lagging cities will widen, not narrow, over time. This prediction rests on the logic of path dependency and increasing returns that characterizes innovation ecosystems [16,48].

Cities that have already deployed across multiple vectors are accumulating advantages that are difficult to replicate: operational data that improves AI systems, regulatory know-how that accelerates subsequent deployments, talent pools that deepen with each new project, public acceptance that builds through positive experience, and institutional memory that reduces the cost and risk of future innovation. Each cycle of deployment, evaluation, and iteration makes the next cycle easier and more productive. The result is a compounding trajectory in which early movers accelerate while latecomers face not only the challenge of catching up technologically but the far more daunting challenge of building the institutional, human, and relational capital that the leading cities have accumulated over years of active deployment.

For cities that have not yet moved decisively, the window of opportunity is narrowing. The historical analogy is instructive: during the Industrial Revolution, the cities that industrialized first—Manchester, Birmingham, Glasgow—maintained their advantages for over a century, shaping national and global economic geography for generations [12,13]. Cities that missed the industrial transition—or that actively resisted it to protect pre-industrial interests—were relegated to peripheral status. The same dynamic is plausible for the agentic AI era: the cities that build the institutional architectures, talent ecosystems, and cross-domain integration capacities in this decade may establish advantages that persist for decades to come.

This does not mean that latecomers cannot succeed. But it does mean that the cost of delay is not linear—it is exponential, because each year of inaction represents not just lost time but lost learning, lost talent formation, and lost ecosystem development. Cities that aspire to participate in the AI-driven transformation must act now—not with isolated pilots or cautious studies, but with the kind of comprehensive, simultaneous, institutionally transformative commitment that the leading cities have already made. The question, as the “Ciutats i IA” essay poses it, is whether a city wants to be one that tries, leads, and shapes its own future—or one that watches from the sidelines as others build the future and then adopts, on others’ terms, the models designed elsewhere [54].

7. Conclusions and Future Research Directions

This paper has argued that the agentic AI era is reshaping cities through three interconnected vectors—the transformation of public administration toward cognitive government, the emergence of autonomous electric mobility, and the deployment of robotics and intelligent infrastructure in the urban environment—and that the city, not the nation or the firm, is the natural unit of analysis for understanding this transformation.

Three principal contributions emerge from the analysis.

First, we have extended the mirroring hypothesis from its original domain of firm-product architecture (Conway [17], Colfer and Baldwin [18]) in two directions. The *dynamic extension* holds that organizations and ecosystems do not merely mirror their current structures but explore and

converge toward the best strategic configurations that a new technology makes possible. The *ontological extension* holds that agentic AI fundamentally changes the nature of the participating agents: for the first time, the organizational architectures that must achieve symmetry with technology include not only human roles but also AI agents with defined decision rights, authority boundaries, and coordination protocols. The city's institutional ecosystem—governance structures, regulatory frameworks, talent markets, infrastructure management—must evolve to mirror the possibilities of this hybrid human-AI coordination architecture. Where it does, compounding value follows; where it does not, the city falls into an asymmetry trap.

Second, we have proposed the concept of *cumulative recursive hybridization* to explain why the convergence of the three vectors within specific urban ecosystems generates compounding returns analogous to those observed during the Industrial Revolution. The mechanism operates through four feedback loops—data, regulatory, infrastructure, and talent—each of which is recursive (each cycle increases the system's capacity for the next) and cumulative (gains compound over time). The concept draws on the Medici Effect (Johansson [43]), systems thinking, and the historiography of industrial clusters (Allen [12], Mokyr [13]) to explain why transformation is inherently local and why cross-domain interaction, not single-vector advancement, is the source of the most consequential innovations.

Third, the comparative analysis across more than twenty governance chatbot deployments, the rapidly scaling autonomous mobility ecosystems of the United States and China, and the emerging urban robotics landscape has identified the conditions that distinguish leading cities from lagging ones: simultaneity of deployment across vectors, institutional willingness to redesign (not merely adopt), tolerance for iteration and failure, and public-sector entrepreneurship as ecosystem orchestration. These conditions are fundamentally local—they depend on municipal regulatory authority, talent ecosystem density, institutional proximity, and the embedded knowledge that accumulates only through active projects.

The implications for urban policy are direct. Cities that wish to participate in the AI-driven transformation cannot afford to wait for national directives, pursue isolated pilots, or protect incumbent structures at the expense of systemic innovation. They must act simultaneously across governance, mobility, and infrastructure; they must redesign institutional architectures to incorporate AI agents as legitimate participants; and they must accept the iterative, experimental character of a transformation for which no established playbook exists. The role of municipal government is not passive regulation but active ecosystem orchestration—co-creating, with firms, universities, and citizens, the conditions under which cumulative recursive hybridization can take hold.

Several limitations of this study point toward avenues for future research. First, the conceptual framework of cumulative recursive hybridization, while grounded in historical analogy and contemporary evidence, has not been tested through formal modeling or econometric analysis; developing quantitative measures of cross-vector interaction effects and ecosystem density would strengthen the framework considerably. Second, the comparative analysis of governance chatbots, while extensive, draws primarily on publicly available documentation and reported capabilities rather than on systematic user-experience data or impact evaluations; longitudinal studies tracking the progression of specific deployments across maturity levels would provide more rigorous evidence of the mirroring hypothesis in action. Third, the paper has focused predominantly on cities in the United States, China, and Europe; the dynamics of AI-driven urban transformation in African, South Asian, and Southeast Asian cities—where institutional contexts, infrastructure basements, and demographic pressures differ markedly—deserve dedicated investigation. Fourth, while we have deliberately avoided centering the analysis on ethical frameworks, the governance challenges raised by hybrid human-AI coordination architectures in the public sphere—accountability, transparency, surveillance, labor displacement—require sustained interdisciplinary attention that goes beyond the scope of this paper.

The agentic AI era is not a distant prospect. It is already reshaping the cities that have chosen to engage with it. The central message of this paper is that the cities which will define the next era of urban civilization are those that understand the transformation as systemic, local, and institutional—not merely technological—and that act accordingly, with the ambition, simultaneity, and institutional courage that the moment demands.

References

1. Lipsey, R.G.; Carlaw, K.I.; Bekar, C.T. *Economic Transformations: General Purpose Technologies and Long-Term Economic Growth*; Oxford University Press: Oxford, UK, 2005.
2. McKinsey Global Institute. *The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value*; McKinsey & Company: New York, NY, USA, 2024.
3. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arber, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. *On the Opportunities and Risks of Foundation Models*. arXiv 2021, 2108.07258.
4. Noy, S.; Zhang, W. Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science* 2023, 381, 187–192.
5. Glaeser, E.L. *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier*; Penguin Press: New York, NY, USA, 2011.
6. Batty, M. *The New Science of Cities*; MIT Press: Cambridge, MA, USA, 2013.
7. Kitchin, R. *The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences*; SAGE: London, UK, 2014.
8. Dunleavy, P.; Margetts, H.; Bastow, S.; Tinkler, J. *Digital Era Governance: IT Corporations, the State, and E-Government*; Oxford University Press: Oxford, UK, 2006.
9. Mergel, I.; Edelman, N.; Haug, N. *Defining Digital Transformation: Results from Expert Interviews*. *Gov. Inf. Q.* 2019, 36, 101385.
10. Almirall, E. *Reinventing Public Governance: From Digital Governments to Agentic Governance*. *Int. J. Eng. Bus. Manag.* 2026, forthcoming.
11. Wirtz, B.W.; Langer, P.F.; Fenner, C. *Artificial Intelligence in the Public Sector—A Research Agenda*. *Int. J. Public Adm.* 2021, 44, 1103–1128.
12. Allen, R.C. *The British Industrial Revolution in Global Perspective*; Cambridge University Press: Cambridge, UK, 2009.
13. Mokyr, J. *The Lever of Riches: Technological Creativity and Economic Progress*; Oxford University Press: Oxford, UK, 1990.
14. Florida, R. *The Rise of the Creative Class: And How It's Transforming Work, Leisure, Community and Everyday Life*; Basic Books: New York, NY, USA, 2002.
15. Moretti, E. *The New Geography of Jobs*; Houghton Mifflin Harcourt: Boston, MA, USA, 2012.
16. Arthur, W.B. *Increasing Returns and Path Dependence in the Economy*; University of Michigan Press: Ann Arbor, MI, USA, 1994.
17. Conway, M.E. *How Do Committees Invent?* *Datamation* 1968, 14, 28–31.
18. Colfer, L.J.; Baldwin, C.Y. *The Mirroring Hypothesis: Theory, Evidence, and Exceptions*. *Ind. Corp. Change* 2016, 25, 709–738.
19. Solow, R.M. *We'd Better Watch Out*. *New York Times Book Review*, 12 July 1987, p. 36.
20. David, P.A. *The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox*. *Am. Econ. Rev.* 1990, 80, 355–361.
21. Milgrom, P.; Roberts, J. *Complementarities and Fit: Strategy, Structure, and Organizational Change in Manufacturing*. *J. Account. Econ.* 1995, 19, 179–208.
22. Brynjolfsson, E.; Hitt, L.M. *Beyond Computation: Information Technology, Organizational Transformation and Business Performance*. *J. Econ. Perspect.* 2000, 14, 23–48.
23. March, J.G. *Exploration and Exploitation in Organizational Learning*. *Organ. Sci.* 1991, 2, 71–87.
24. Nelson, R.R.; Winter, S.G. *An Evolutionary Theory of Economic Change*; Harvard University Press: Cambridge, MA, USA, 1982.

25. Brynjolfsson, E.; McAfee, A. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*; W.W. Norton: New York, NY, USA, 2014.
26. Xi, Z.; Chen, W.; Guo, X.; He, W.; Ding, Y.; Hong, B.; Zhang, M.; Wang, J.; Jin, S.; Zhou, E.; et al. The Rise and Potential of Large Language Model Based Agents: A Survey. *arXiv* 2023, 2309.07864.
27. Wang, L.; Ma, C.; Feng, X.; Zhang, Z.; Yang, H.; Zhang, J.; Chen, Z.; Tang, J.; Chen, X.; Lin, Y.; et al. A Survey on Large Language Model Based Autonomous Agents. *Front. Comput. Sci.* 2024, 18, 186345.
28. Dafoe, A.; Bachrach, Y.; Hadfield, G.; Horvitz, E.; Larson, K.; Graepel, T. Cooperative AI: Machines Must Learn to Find Common Ground. *Nature* 2021, 593, 33–36.
29. Pollitt, C.; Bouckaert, G. *Public Management Reform: A Comparative Analysis—Into the Age of Austerity*, 4th ed.; Oxford University Press: Oxford, UK, 2017.
30. Jacobs, J. *The Economy of Cities*; Vintage Books: New York, NY, USA, 1969.
31. Frenken, K.; Van Oort, F.; Verburg, T. Related Variety, Unrelated Variety and Regional Economic Growth. *Reg. Stud.* 2007, 41, 685–697.
32. Boschma, R. Proximity and Innovation: A Critical Assessment. *Reg. Stud.* 2005, 39, 61–74.
33. Gertler, M.S. Tacit Knowledge and the Economic Geography of Context, or The Undefinable Tacitness of Being (There). *J. Econ. Geogr.* 2003, 3, 75–99.
34. Saxenian, A. *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*; Harvard University Press: Cambridge, MA, USA, 1994.
35. Storper, M.; Venables, A.J. Buzz: Face-to-Face Contact and the Urban Economy. *J. Econ. Geogr.* 2004, 4, 351–370.
36. Marshall, A. *Principles of Economics*, 8th ed.; Macmillan: London, UK, 1920.
37. Feldman, M.P.; Audretsch, D.B. Innovation in Cities: Science-Based Diversity, Specialization and Localized Competition. *Eur. Econ. Rev.* 1999, 43, 409–429.
38. Nonaka, I.; Takeuchi, H. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*; Oxford University Press: Oxford, UK, 1995.
39. Lave, J.; Wenger, E. *Situated Learning: Legitimate Peripheral Participation*; Cambridge University Press: Cambridge, UK, 1991.
40. DeFillippi, R.J.; Arthur, M.B. The Boundaryless Career: A Competency-Based Perspective. *J. Organ. Behav.* 1994, 15, 307–324.
41. Grabher, G. Learning in Projects, Remembering in Networks? Communitarity, Sociality, and Connectivity in Project Ecologies. *Eur. Urban Reg. Stud.* 2004, 11, 103–123.
42. Glaeser, E.L.; Kerr, W.R.; Ponzetto, G.A.M. Clusters of Entrepreneurship. *J. Urban Econ.* 2010, 67, 150–168.
43. Johansson, F. *The Medici Effect: What Elephants and Epidemics Can Teach Us About Innovation*; Harvard Business School Press: Boston, MA, USA, 2004.
44. Hargadon, A. *How Breakthroughs Happen: The Surprising Truth About How Companies Innovate*; Harvard Business School Press: Boston, MA, USA, 2003.
45. Meadows, D.H. *Thinking in Systems: A Primer*; Chelsea Green Publishing: White River Junction, VT, USA, 2008.
46. Sterman, J.D. *Business Dynamics: Systems Thinking and Modeling for a Complex World*; McGraw-Hill: Boston, MA, USA, 2000.
47. Landes, D.S. *The Unbound Prometheus: Technological Change and Industrial Development in Western Europe from 1750 to the Present*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2003.
48. David, P.A. Clio and the Economics of QWERTY. *Am. Econ. Rev.* 1985, 75, 332–337.
49. Almirall, E.; Fernández, M.; Tudurí, M.; Fernández, B. *Artificial Intelligence and the Public Sector: Scenarios for 2030*; Anteverti and Esade: Barcelona, Spain, 2025.
50. Følstad, A.; Brandtzæg, P.B. Chatbots and the New World of HCI. *Interactions* 2017, 24, 38–42.
51. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-T.; Rocktäschel, T.; et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*; Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H., Eds.; Curran Associates: Red Hook, NY, USA, 2020; pp. 9459–9474.

52. Larsen, A.G.; Følstad, A. The Impact of Chatbots on Public Service Provision: A Qualitative Interview Study with Citizens and Public Service Providers. *Gov. Inf. Q.* 2024, 41, 101927.
53. van Noordt, C.; Misuraca, G. Artificial Intelligence for the Public Sector: Results of Landscaping the Use of AI in Government across the European Union. *Gov. Inf. Q.* 2022, 39, 101714.
54. Almirall, E. *Ciutats i IA (Cities and AI)*. Esade Center for Innovation in Cities: Barcelona, Spain, 2025.
55. Newman, P.; Kenworthy, J. *The End of Automobile Dependence: How Cities Are Moving Beyond Car-Based Planning*; Island Press: Washington, DC, USA, 2015.
56. Senge, P.M. *The Fifth Discipline: The Art and Practice of the Learning Organization*; Doubleday: New York, NY, USA, 1990.
57. Waymo. Waymo One Surpasses 500,000 Weekly Paid Rides. Waymo Blog, 2026. Available online: <https://waymo.com/blog/> (accessed on 25 March 2026).
58. Hawkins, A.J. Waymo Is Now Doing over 500,000 Paid Robotaxi Rides per Week. *The Verge*, 2026. Available online: <https://www.theverge.com/> (accessed on 25 March 2026).
59. Baidu. Apollo Go Autonomous Ride-Hailing Service Surpasses 20 Million Cumulative Rides. Baidu IR, 2026. Available online: <https://ir.baidu.com/> (accessed on 25 March 2026).
60. Pony.ai. Pony.ai Investor Presentation, Q4 2025. Available online: <https://pony.ai/> (accessed on 25 March 2026).
61. WeRide; Shenzhen Bus Group. WeRide and Shenzhen Bus Group Launch Shenzhen's First Fully Driverless Robobus Line. WeRide Press Release, August 2025. Available online: <https://ir.weride.ai/> (accessed on 25 March 2026).
62. QCraft. Autonomous Buses Operating in Chinese Cities: Status Report 2025. Available online: <https://www.qcraft.ai/> (accessed on 25 March 2026).
63. Meituan. Meituan Autonomous Delivery: Scaling Last-Mile Logistics with Autonomous Vehicles. Meituan Technical Blog, 2025. Available online: <https://tech.meituan.com/> (accessed on 25 March 2026).
64. Neolix. Neolix Autonomous Delivery Vehicle: 30,000 Global Orders and Counting. Company Report, 2025. Available online: <https://www.neolix.ai/> (accessed on 25 March 2026).
65. International Transport Forum. *The Shared-Use City: Managing the Curb*; OECD/ITF: Paris, France, 2018.
66. Glus, P.; Pagonis, T.; Czabanowski, R. Autonomous Vehicles and Urban Policy: Towards a Research Agenda. *Transp. Rev.* 2021, 41, 1–20.
67. Florida, R. *The Geography of Autonomous Vehicles*. CityLab/Bloomberg, 2024. Available online: <https://www.bloomberg.com/citylab/> (accessed on 25 March 2026).
68. Alibaba Cloud. City Brain: Now in 23 Cities in Asia. Alibaba Cloud Blog, 2020. Available online: <https://www.alibabacloud.com/blog/> (accessed on 25 March 2026).
69. Xue, L.; Zhang, J. City Brain: Practice and Thinking of Urban Governance with Artificial Intelligence. *Front. Eng. Manag.* 2022, 9, 207–217.
70. Pan, Y.; Tian, Y.; Liu, X.; Gu, D.; Hua, G. Urban Big Data and the Development of City Intelligence. *Engineering* 2016, 2, 171–178.
71. Chen, Y.; Zhang, D. Integration of Environmental Monitoring with Urban AI Platforms in China. *Environ. Sci. Pollut. Res.* 2023, 30, 54102–54118.
72. Cowa Robot. Cowa Deploys 36 Sanitation Robots in 2.7 Million Square Meter Area in Shenzhen. Yicai Global, 2025. Available online: <https://www.yicai.com/> (accessed on 25 March 2026).
73. Guangzhou Municipal Government. Guangzhou to Increase Unmanned Sanitation Fleet to 1,000 Units by 2025. Government Press Release, 2025.
74. Bogue, R. Robots for Infrastructure Inspection. *Ind. Robot* 2023, 50, 217–222.
75. Walmart. Walmart Drone Delivery Expands to 1.8 Million Additional Households. Walmart Corporate, 2025. Available online: <https://corporate.walmart.com/> (accessed on 25 March 2026).
76. South China Morning Post. Chengdu Deploys Robot Police Teams Combining Quadruped, Wheeled, and Humanoid Robots. SCMP, June 2025. Available online: <https://www.scmp.com/> (accessed on 25 March 2026).

77. Xinhua. Hangzhou Places AI-Powered Traffic Policing Robots on Active Duty. Xinhua News Agency, December 2025. Available online: <https://www.xinhuanet.com/> (accessed on 25 March 2026).
78. Shenzhen Municipal Government. Shenzhen Designed as China's First "Robot-Friendly" Urban District. Government Press Release, 2025.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.