



Cities After AI

Collective intelligence, the AI dividend and the future of living cities

Pavel Luksha, Founder, Gaia Evolution Foundation, and Co-founder, Living Cities Earth

Alexander Maneev, Chief Technology Officer, Global Mayors Academy

Lev Gordon, Founder, Living Cities Earth

Draft v0.1 for discussion · October 2026



Cities After AI

Collective intelligence, the AI dividend and the future of living cities

VERSION	Draft v0.1 for discussion, October 2026
AUTHORS	Pavel Luksha , Founder, Gaia Evolution Foundation, and Co-founder, Living Cities Earth Alexander Maneev , Chief Technology Officer, Global Mayors Academy Lev Gordon , Founder, Living Cities Earth
STATUS	Draft for discussion, not for citation. References are being added.
READ ONLINE	hub.aiforcities.earth/cities-after-ai

Contents

EXECUTIVE SUMMARY	4
KEY MESSAGES FOR THE GLOBAL DIALOGUE	6
OPENING: THE CITY WE DID NOT EXPECT	7
1. AI IS CHANGING WHAT CITIES ARE, NOT ONLY WHAT THEY DO	9
2. A CITY IS A LIVING SYSTEM WITHIN A LIVING PLANET	16
3. WHO PERCEIVES, DECIDES AND ACTS IN THE CITY IS CHANGING	21
4. THE GREATEST OPPORTUNITY IS A MORE COLLECTIVELY INTELLIGENT CITY	27
5. MORE POWERFUL TECHNOLOGY CAN PRODUCE LESS CAPABLE CITIES	34
6. HOW CITIES USE THE AI DIVIDEND WILL SHAPE THEIR FUTURES	44
7. CITIES MUST LEARN TO SHAPE THE TRANSITION TOGETHER	51
CONCLUSION: AN INVITATION TO BUILD THRIVING CITIES AFTER AI	61
GLOSSARY	65

EXECUTIVE SUMMARY

Artificial intelligence is changing not only what cities can do, but what cities are. It is reshaping how people work, learn and decide, how businesses and institutions operate, and how infrastructure behaves. Most of these changes will not arrive through a single city strategy. They will arrive through thousands of separate decisions made by city departments, companies, schools, utilities and households.

This paper argues that the most important question is not how quickly a city adopts AI, but what kind of city emerges as a result. It sets out seven connected propositions.

1. AI is a systemic force, not another urban technology. Cities therefore need to design from the conditions emerging **after** AI is fully deployed across city dimensions, rather than fitting AI into today's arrangements.
2. The subject of transformation is the whole city, understood as a **living system** within a living planet. City government has indispensable public responsibilities, but the transformation involves all six overlapping **voices of the city**: government and public institutions, business, civil society and communities, residents, young people, and AI developers, researchers and technology providers.
3. AI is changing who and what perceives, decides and acts in a city. AI systems are spreading into devices, organisations and infrastructure, and increasingly act on people's behalf. This creates new possibilities and new questions about authority and accountability.
4. The greatest opportunity is a **more collectively intelligent city**: one better able to hear more of its voices, develop shared understanding, make legitimate collective decisions, coordinate action and learn.
5. More powerful technology can produce less capable cities. One of the deepest risks is a **future by default**: a city reshaped by many separate decisions without sufficient collective understanding or deliberate choice about where it is heading. Alongside it come growing dependence, weakened human capabilities, the improvement of parts at the expense of the whole, concentrated benefits with displaced costs, and erosion of trust. **Long-term viability**, not technological performance, must become the measure of progress.
6. The **AI dividend** — time, attention, resources and new capabilities that AI frees up or creates — should translate into greater human capabilities, shared prosperity and stronger communities, rather

than being absorbed by heavier workloads or captured by a few powerful actors.

7. No proven model of the “city after AI” exists. Cities need a **new kind of leadership** and a shared capacity to experiment and learn, within cities and between them. Any city, large or small, can take this path, and none has to take it alone.

The paper concludes by introducing the *Lichtensteig Declaration on Living Cities and AI* as the first shared expression of these commitments, and by inviting cities, organisations and individuals into an open process of dialogue, experimentation and learning.

KEY MESSAGES FOR THE GLOBAL DIALOGUE

1

AI is changing what cities are, not only what they do. Cities should design from the conditions emerging after AI becomes a new normal.

2

The whole city is the subject of change, and it is part of a living planet. Success means people and the living systems that sustain them thriving together.

3

AI's greatest promise for cities is collective intelligence: cities that hear more of themselves, coordinate better and learn faster.

4

A city can become technologically smarter while becoming less capable of shaping its own future. The greatest danger is a future no one chose.

5

The AI dividend, meaning the time, attention, resources and new capabilities that AI frees up or creates, should strengthen people and communities, not simply disappear into cost cuts, heavier workloads or distant profits.

6

AI transformation is a leadership task for the whole city, not an IT project.

7

AI creates opportunities for cities of every size to leap ahead, but there is no proven path. Cities must experiment, learn and shape the transition together.

OPENING: THE CITY WE DID NOT EXPECT

Imagine the city you live in, in 2032. Its streets and buildings are familiar, but everyday life has changed. People rely on AI assistants to arrange purchases, find work, choose services and plan their journeys. Local businesses use AI to manage operations, develop products and reach customers. Schools, hospitals and public institutions have transformed in their own ways. Much of this is welcome: services are more accessible, new opportunities have opened, and people can do things that once required considerable time, money or expertise.

Yet the city has also changed in less obvious ways. A bookshop that served the neighbourhood for thirty years has closed, not because people stopped wanting it, but because their assistants rarely recommended it. Office buildings stand half-empty as companies reorganise work, changing the demand for housing and transport and the livelihoods available to residents. New economic activity has emerged, but much of the value it generates flows to companies elsewhere. And during last summer's heatwave, the hottest streets were in neighbourhoods whose residents rarely take part in consultations, even though the city's sensors had been recording the rising temperatures for years.

None of these changes required a decision about the future of the city as a whole. Each emerged from choices made by residents, businesses, institutions and technology providers pursuing their own objectives. Yet their combined consequences may prove more profound than anything in the city's development strategy.

Now imagine a different 2032. When the heat data begins to rise, the city brings it together with what residents of those neighbourhoods know about their streets, and with ideas from local businesses and young people. Within weeks, it decides together where shade, water and green space should go. The bookshop's owner, and others like her, have joined with the city to make sure local businesses can still be found. The time freed by AI in the city hall has been redirected toward precisely this kind of work.

AI could help make this second city possible. But it would take much more than widespread adoption of intelligent technologies. It would take a city that has become more capable of understanding, deciding and acting together.

This is the question at the heart of this paper: as AI transforms urban life, will cities become the sum of countless separate technological choices, or can they develop the collective intelligence to shape their own evolution?

Collective intelligence

The capacity of a city's people, communities, organisations and institutions to develop shared understanding, make and support legitimate decisions, coordinate action and learn together, achieving more than any could alone.

1.

AI IS CHANGING WHAT CITIES ARE, NOT ONLY WHAT THEY DO



AI is not simply a new tool for city services. It is changing how the whole city works: its economy, its institutions, its infrastructure, its relationships, and its place within the living world.

Not another technology project

For several decades, cities have absorbed new technologies as projects. A traffic control system, a permit portal or a housing database arrived with a budget, a department and a completion date. It is tempting to treat AI the same way, as one more project for the digital team. That framing misses what is happening.

AI changes how existing work is done, but it also changes the economy, the nature of work, the way knowledge is produced and decisions are made, the behaviour of infrastructure, and the roles of institutions. As analytical and decision-making capability becomes widely available and embedded in everyday systems, it changes who creates value, who holds knowledge, how decisions are made and how people and organisations interact. These are changes to the organisation and evolution of the city, not only to its functions. A city that manages AI as an IT project may discover in a few years that the ground has shifted under all of its life while only one department was watching.

AI decisions are already being made throughout the city

Many of the most consequential choices about AI affecting a city are not made by the city government at all. A logistics company changes its delivery routes. A school allows or forbids AI tools for homework. An employer automates its customer service. A property manager adopts algorithmic pricing. A family relies on an AI assistant to navigate health care or benefits. Each decision is made independently, often for good reasons, and each contributes to shaping the city's future.

This is why the question cannot be limited to “how should city hall use AI?” The more consequential question is how the city as a whole, with all its actors, will live with and shape AI.

AI arrives into pressures that are already entangled

AI is arriving at a moment when cities already face multiple pressures that reinforce one another: changing economies and labour markets, demographic shifts, strained public services, social fragmentation, climate disruption and mounting ecological risk. These do not queue up politely to be solved one at a time. Each worsens the others, and together

they produce challenges that no single institution or sector can address alone. Together, these interconnected and mutually reinforcing pressures constitute a polycrisis, in which ecological, social, economic and geopolitical disruptions compound one another and cannot be adequately addressed in isolation.

AI could help cities respond, but it will also transform the capabilities, institutions and relationships through which they respond. Research on government suggests the difficulty is not primarily technological. The World Governments Summit report *Rise of the Wise Government* (2026) concludes that the defining challenge is the accelerating pace of change, which is outstripping institutional capacity, workforce readiness and traditional models of administration. For cities, where these pressures land most concretely, that gap is the heart of the matter.

This is also why cities may become more important, not less, in an AI-rich world. Even as more work and interaction move into digital environments, urban life remains grounded in physical places, shared infrastructure, ecological relationships and institutions that must address collective needs and risks. Cities hold these interdependencies in close contact. They are where the consequences of technological change become visible first, and where new forms of collective response can be built.

Five shifts that change the nature of the city

AI is accelerating transformations already under way in urban life, while opening possibilities that go beyond earlier waves of digitalisation. Five shifts help explain why AI may change not only what cities do, but what they are.

1. *From established patterns of work and production to AI-shaped urban economies.* AI is changing the capabilities needed to produce goods and services, the division of labour between people and machines, and the organisation of firms and economic networks. Small businesses and individuals gain access to capabilities previously reserved for large organisations, while existing occupations, business models and local value chains face disruption. AI will reshape what cities specialise in, where economic value is created, and how much of that value stays within local communities.

2. *From predominantly human-mediated urban life to new relationships between people and AI.* Increasingly capable AI systems are becoming part of how residents learn, work, shop, access services, seek advice and deal with institutions. This creates new opportunities for accessibility, learning and connection, and it changes how people interact with one

another. The consequences reach past convenience into social trust, cultural life, community relationships and the human capabilities that everyday urban experience exercises or allows to fade.

3. *From largely passive infrastructure to adaptive urban environments.* Buildings, mobility systems, energy networks and logistics operations are becoming connected and responsive. AI enables them to anticipate demand, adjust to changing conditions and coordinate with growing autonomy. Infrastructure is no longer only the fixed foundation on which urban activity takes place; it increasingly shapes that activity through continuous sensing, decision and action. This is not a distant prospect. In **Seoul** in 2026, city leaders and technology executives gathered specifically to address physical AI, including quadruped robots, humanoid machines and autonomous systems operating on real city streets.

4. *From inherited spatial patterns to new geographies of urban life.* Cities developed around particular relationships between employment, housing, mobility, commerce and public space. AI-enabled changes to work, production and services alter those relationships, with consequences for where economic activity locates, how buildings are used, what transport is needed and how neighbourhoods develop. AI could reinforce existing concentrations or enable more distributed activity; different cities will see different patterns. What it creates is a new range of possibilities for urban spatial organisation, not a single new urban form.

5. *From limited ecological visibility to a new relationship with the living world.* This shift is not an environmental appendix to the others. It is the ground beneath them. Cities depend on wider systems of energy, water, materials, soil and biodiversity, and the health of those systems sets the boundary within which everything else in this list succeeds or fails.

AI changes this relationship in both directions at once. It offers genuinely new capacities to understand ecological relationships, anticipate environmental risks and coordinate responses across urban and ecological boundaries, bringing into view conditions that cities have long affected without seeing. At the same time, AI itself consumes electricity, water, land and materials on a scale that is beginning to reshape energy systems. The **International Energy Agency** projects that data centre electricity consumption will more than double by 2030, to around 945 terawatt-hours, with AI the most significant driver of that increase. The Agency's 2026 report indicates that data centre electricity demand surged 17% in 2025 and AI-linked demand is expected to triple by 2030. The Agency also notes that data centres are growing larger and tend to cluster, often around cities. These costs frequently fall on places other than those enjoying the benefits, and communities have noticed: in **the United States**, membership of groups opposing AI data centres grew 8-fold within eight months, passing 640,000 by August 2026.

A city that treats AI as a question of services and efficiency will see neither side of this. It will miss what AI could help it understand about the living systems it depends on, and it will miss what its own AI use is costing those systems elsewhere. Both belong on the same page, from the beginning.

Taken together, these shifts describe something more fundamental than the adoption of new tools. Urban economies, relationships, infrastructure, spatial organisation and ecological interactions are evolving together, at different speeds, with consequences that cannot be understood by examining any one of them alone.

Two further shifts deserve separate attention: who and what can perceive and act within the city (Section 3), and how the city can develop the collective intelligence to understand, decide and act together (Section 4).

The city as a moving object

Cities have planned through relatively stable cycles of investment, regulation and public decision-making. AI accelerates change in economic activity, information flows and technological systems, while multiplyi-

Agency

The capacity to initiate or carry out actions that influence outcomes. Agency does not, by itself, confer legitimate authority to decide or act on behalf of others.

Living system

An evolving, interdependent whole whose components and relationships shape its development, adaptation and viability. In cities, this includes people, communities, culture, economic activity, institutions, infrastructure and ecological processes.

Collective intelligence

The capacity of a city's people, communities, organisations and institutions to develop shared understanding, make and support legitimate decisions, coordinate action and learn together, achieving more than any could alone.

ng the number of actors able to influence urban development. A plan written to last five years now meets systems that change from week to week.

The city is becoming a **moving object**. Its leaders must steer a system that is itself continuously in motion, and they must do so while moving themselves. This calls for a different practice: sensing change early rather than discovering it in annual reviews, adjusting often rather than waiting for the next planning cycle, and staying clear about which commitments must hold steady while everything else moves. The challenge is not to make city government faster. It is to develop the capacity to understand and shape an urban system whose evolution is continuous and distributed.

We have been here before, partly

Cities have lived through a previous wave of urban technology. The smart-city movement of the 2010s brought real benefits: better traffic management, more efficient utilities, more accessible public information. It also taught hard lessons. Where technology led and public purpose followed, results often disappointed, and some projects struggled to establish the trust and legitimacy they needed.

Toronto's Quayside development is the best-known example. Alphabet's Sidewalk Labs proposed an ambitious technology-enabled neighbourhood, and the proposal raised serious questions about data governance, privacy and public accountability. The company withdrew in 2020, after a project that had faced sustained criticism over surveillance and privacy, though it gave economic uncertainty and the state of the property market as its reasons. The withdrawal cannot be attributed to the controversy alone, but the episode showed how hard it is to reconcile technological ambition, commercial interest and legitimate public expectations.

The lesson is not that technology fails cities. It is that a mobilisation, with its pilots, contracts and dashboards, is not the same as building lasting capability. The current wave of AI touches far more of urban life than the smart-city wave did, which makes that lesson more important now, not less.

Designing from "after AI"

The familiar question most AI related discussions ask is: how do we introduce AI into today's city? This paper asks a different one: **what conditions are emerging on the other side of this transition, and what kind of city can thrive in them?**

This paper asks a different one: what conditions are emerging on the other side of this transition, and what kind of city can thrive in them?

After AI is the name we give to this perspective. It means looking beyond the phase of adopting new tools, toward the conditions and possibilities that emerge as AI becomes part of everyday urban life. It does not mean that technological change will be finished, and it certainly does not mean that cities end. Cities are where people, institutions, technology and the living world meet at close range, and where the consequences of choices are felt first. Their importance is likely to grow.

Designing from after AI means reconsidering how a city might work, how people and institutions might relate to one another, what new economic and social possibilities become available, and what forms of urban life are worth creating. Rather than fitting AI into inherited arrangements, cities can use this moment to question those arrangements and build better ones.

Inaction is not a neutral option. A city that waits will still be transformed, largely by decisions taken beyond the reach of its own institutions. Neither is uncritical acceleration. A city that adopts as fast as it can optimises itself into a future it never chose.

A THOUGHT FOR LEADERS:

Which decisions about AI are already being made across your city, by whom, and with what consequences for the future you want to create?

2.

A CITY IS A LIVING SYSTEM WITHIN A LIVING PLANET



The subject of transformation is the whole city, not city government alone. A city is a living system of people, communities, institutions, infrastructure and nature, whose capacity to thrive depends on the health of the larger living systems it belongs to.

The city is more than its administration

City government has indispensable responsibilities for public goods, collective risks, rights and decisions that require legitimate public authority. But no single institution can fully understand or control a city. Its knowledge, capabilities and resources are distributed across residents, businesses, communities, public institutions and the relationships between them. The city is shaped as much by their interactions as by formal plans and policies.

The ambition, therefore, is not simply an "AI-enabled municipality," with smarter services and unchanged habits. It is a city that can thrive after AI, with public institutions and the wider urban community developing new capacities together.

Seeing the city as a living system

We understand the city as a living system: an evolving whole in which people, organisations, cultures, infrastructure and ecological processes interact, generating patterns of life and development that no single actor designs or controls. Its vitality depends not only on the performance of individual parts, but also on the quality of their relationships, their capacity to adapt and their ability to sustain the conditions on which urban life depends.

This perspective builds on a tradition of understanding cities as complex, evolving wholes, including Marilyn Hamilton's *Integral City* approach and research on urban social-ecological systems.

Seeing the city this way reveals what sector-by-sector management tends to miss. Each department, company and community can optimise its own activities while the city as a whole becomes less capable. Traffic may move faster while streets become less welcoming to people. Services may become more efficient while relationships between institutions and communities weaken. Conversely, improvements in one area can strengthen many others when they build connections, trust and shared capabilities.

AI can influence these relationships in both directions. Its significance therefore cannot be judged only by how well individual applications perform, but also by how they affect the development of the wider urban

system.

The voices of the city

To understand who must be involved in shaping cities after AI, we use the idea of the *voices of the city*: distinct perspectives that need to be brought into dialogue. Building on the Integral City tradition and adapting it to the AI transition, we distinguish six of them:

- **City government and public institutions** carry public mandates and responsibilities for collective goods, rights and legitimate decisions.
- **Business** creates economic value, employment, products and services, and helps shape the city's opportunities and livelihoods.
- **Civil society and communities** organize collective action, contribute local knowledge and can help bring perspectives that public institutions often overlook.
- **Residents** bring their lived experience, needs, aspirations and knowledge of everyday urban life, whether or not they participate in formal civic processes.
- **Young people** have a particular stake in the city's long-term future and bring perspectives on learning, work, culture and emerging possibilities.
- **AI developers, researchers and technology providers** increasingly shape the capabilities and systems available across the city and carry responsibilities for their design and consequences.

These voices are perspectives, not fixed categories or equivalent sources of authority. A shop owner is also a resident and perhaps a parent; young people may be entrepreneurs, community organisers or researchers. The distinction helps reveal whose knowledge and interests are being considered and whose may be missing. It does not by itself guarantee representation or participation.

The list also raises questions about forms of agency and representation that do not fit neatly into conventional stakeholder categories. The natural world is already part of urban life, yet its interests and conditions are often inadequately represented in decisions. AI systems, meanwhile, are becoming increasingly capable of acting within the city. Neither technological sensing nor technological agency automatically confers legitimate authority to speak or decide on anyone else's behalf. We return to these distinctions in Section 3.

Cities after AI

A perspective on the future taking shape as AI becomes embedded in urban life, looking beyond technology adoption toward how cities, their people and institutions can develop under these new conditions.

Agency

The capacity to initiate or carry out actions that influence outcomes. Agency does not, by itself, confer legitimate authority to decide or act on behalf of others.

Cities are nested within larger living systems

Cities do not exist independently of the wider living world. People develop within communities and cities; cities depend on regional economies, landscapes, watersheds and ecological networks; and all are part of the living planet. These relationships extend across administrative boundaries and influence one another in both directions.

Cities depend on flows of water, energy, food and materials, while their patterns of production, consumption, construction and waste affect ecosystems far beyond their boundaries. At the same time, urban communities can help restore the living systems on which they depend, through healthier waterways, regenerated landscapes, biodiversity, green spaces and more sustainable relationships with surrounding regions.

These interconnected systems function at different scales and through different organizing principles. What enables an individual to flourish is not identical to what makes a community, a city or an ecosystem viable. The task is not to impose the same rules at every level, but to develop forms of cooperation and stewardship that respect these differences while strengthening their mutual capacity to thrive.

Why we need a human- and planet-centred approach to AI

Understanding cities as living systems within a living planet calls for a fundamental shift in how we approach AI: **from judging technologies primarily by their performance or benefits to individual organisations, toward asking whether they strengthen the flourishing of people, cities and the living systems that sustain them.** Human and planetary well-being cannot be treated as separate objectives. The test of any arrangement of people and AI is whether it helps both to thrive.

This does not eliminate difficult choices. Cities must constantly negotiate tensions between economic development, housing, infrastructure, social needs and ecological health. But these are not choices between human prosperity and an unrelated natural world. They are decisions about how people can flourish within the living systems that make their lives possible, now and over generations.

AI offers new possibilities for understanding these interdependencies and developing better responses. Its contribution to urban progress should ultimately be judged by whether it strengthens human capabilities

Living planet

The interconnected ecological and life-supporting systems of Earth, of which cities are an integral part and which their activities both depend on and influence.

Stewardship

The responsible, ongoing care of shared resources, capabilities and living systems, with attention to their long-term health and the well-being of present and future generations.

es, the city's collective capacity to act, and the health and regenerative potential of the living systems of which it is part.

THOUGHT FOR LEADERS:

Is your city's AI agenda focused primarily on improving individual services, or is it also helping strengthen the human, institutional and ecological relationships on which the city's future depends?

3.

WHO PERCEIVES, DECIDES AND ACTS IN THE CITY IS CHANGING



AI is making intelligence and new capabilities more widely available across the city, while increasingly enabling technological systems to act on behalf of people and organisations or operate with greater autonomy. This changes the distribution of agency in urban life and raises new questions about authority, accountability and the relationships between human and artificial intelligence.

Intelligence is spreading throughout the city

Until recently, sophisticated analytical and decision-support capabilities required considerable resources, specialised expertise and technological infrastructure. AI is making many of these capabilities accessible to a much wider range of actors: individuals, small businesses, schools, clinics, community organisations and public institutions. A family bakery can use AI to analyse customer demand and plan its stock; a small community organisation can translate its communications, reach more people and better coordinate its activities.

This is potentially a profound change in the distribution of capabilities. Individuals and small organisations may become able to undertake activities that once required much larger institutions, opening new possibilities for enterprise, learning, creativity and participation in urban life. At the same time, access to these capabilities remains unequal, and their consequences depend on who controls the technologies and the conditions under which they are used.

But possessing greater analytical capability is not the same as having greater agency. AI can help someone understand a situation and make a decision; increasingly, it can also carry out actions on their behalf. This distinction matters because the capacity to act, rather than simply to know, changes how outcomes are produced across the city.

From tools to agents

One of the most consequential developments is the emergence of AI agents: software systems capable of pursuing tasks and undertaking actions with some degree of autonomy. A resident's assistant may negotiate a service appointment, a small business's agent may coordinate orders with suppliers, and a municipal agent may handle requests across several departments. As such systems become more capable, they may increasingly interact with one another, performing tasks that previously required direct human coordination.

AI agent

An AI system capable of pursuing tasks and taking actions with some degree of autonomy, often on behalf of a person or organisation.

A related development is taking place in infrastructure. Energy systems, transport networks and other urban environments increasingly combine sensors, analytical models and automated controls. AI could extend their ability to anticipate conditions, coordinate responses and operate with greater autonomy. Unlike a personal assistant, however, such a system may not act on behalf of a single individual. Its actions may affect entire communities, requiring different arrangements for oversight and responsibility.

These developments could substantially increase the capacity of people and organisations to act. Residents could navigate complex services more easily, small enterprises could coordinate with larger partners on more equal terms, and institutions could respond to changing needs with greater speed and flexibility. Open technical protocols are also emerging to support interactions between AI agents and their access to tools and information, including A2A and MCP.

Yet technological interoperability is not the same as legitimate cooperation. When a resident's agent interacts with a commercial platform or a municipal system, whose instructions does it follow? Who may authorise a commitment, question an outcome or reverse an action? When an autonomous infrastructure system affects public interests, who has the authority to determine its objectives and limits?

These questions become more important as AI systems acquire greater autonomy. Some may eventually develop something like a voice of their own in urban life, generating proposals or positions that cannot be understood simply as the execution of a predefined instruction. Whether or how such contributions should be considered remains an open question. But the distinction between capability and legitimate authority must remain clear: **the ability to perform tasks, express positions or negotiate with other systems does not itself confer the right to represent people or make collective decisions.** Such authority must rest on legitimate human and institutional arrangements.

Interoperability

The ability of technological systems from different providers to exchange information and work together, without requiring all participants to use a single platform.

But the distinction between capability and legitimate authority must remain clear: the ability to perform tasks, express positions or negotiate with other systems does not itself confer the right to represent people or make collective decisions.

Working hypotheses for agents in the city

We do not yet know which approaches to governing AI agents will work best in the complexity of urban life. However, experience with AI systems and emerging agent technologies suggests several principles worth testing.

The first is **loyalty to those on whose behalf an agent acts**. A personal or organisational agent should serve the legitimate interests and intentions of those who authorise it, rather than quietly prioritising the commercial interests of its provider. Its authority should be explicit, limited and revocable, while remaining subject to applicable law and the rights of others.

The second is **answerability**. AI systems should operate within clear arrangements of responsibility, with identifiable people or organisations accountable for their deployment and effects. For agents acting on someone's behalf, this could mean identifying whom they represent, acting only within their mandate, seeking clarification when permission is uncertain, keeping appropriate records of consequential actions, and returning high-stakes, irreversible or contested decisions to accountable human decision-makers. People and organisations should also be able to change providers without unnecessarily losing control of their data and relevant knowledge.

The technical mechanisms for exchanging information between agents are developing rapidly, but the conventions that make their interactions trustworthy are less mature. Email offers a useful analogy: technical standards allow messages to travel between different providers, but they do not establish whether a sender is authorised to make a commitment or who is responsible for its consequences. Similarly, interoperable agents will need shared understandings of identity, consent, authority and accountability.

These are proposals for practical experimentation, not a settled framework or universal standard. Their usefulness, limitations and applicability to different urban contexts will have to be established through real-world experience.

Sensing is not the same as speaking for

AI can also make previously overlooked aspects of urban life more visible, including the condition of natural systems. Environmental sensors, models and analytical tools can help cities understand air quality, urban heat, soil conditions, water systems and biodiversity, and explore the consequences of different interventions.

But understanding the condition of a river is not the same as legitimately representing its interests. A technological system may provide valuable ecological information without possessing the authority to decide what should be done with it or to speak on behalf of the living system concerned.

Existing institutional precedents demonstrate the importance of this distinction. In 2017, **New Zealand** recognised the Whanganui River as a legal person through the Te Awa Tupua legislation. The law established Te Pou Tupua, an office comprising two people nominated through arrangements involving the Whanganui iwi and the Crown, to act as the human face of the river. This is not an AI application. It illustrates how ecological representation can be grounded in cultural relationships, legal recognition and institutional responsibility.

AI may significantly improve the knowledge available to such arrangements, but it cannot establish their legitimacy simply by collecting data or generating recommendations. The possibility of giving natural systems a more meaningful place in urban decision-making therefore calls for the development of institutions and relationships alongside new technological capabilities.

The city becomes faster-changing

As intelligence and the capacity to act become more widely distributed, the possibilities for urban life expand. People and smaller organisations may gain new autonomy, technological systems may coordinate activities previously separated by institutional boundaries, and new kinds of cooperation may become feasible.

Yet a city containing many capable people, organisations and AI agents is not necessarily a collectively intelligent city. Its actors may still pursue incompatible objectives, operate with fragmented knowledge or lack the relationships and institutions needed to address shared problems. Conversely, technological systems can become more autonomous without the people affected gaining greater influence over the decisions that shape their lives.

The critical question is therefore how these expanding capabilities can strengthen the city's ability to understand shared challenges, make legitimate collective decisions, coordinate action and learn. That is the opportunity explored in the next section.

Living system

An evolving, interdependent whole whose components and relationships shape its development, adaptation and viability. In cities, this includes people, communities, culture, economic activity, institutions, infrastructure and ecological processes.

THOUGHT FOR LEADERS:

Where could AI give your city's residents and organisations greater capacity to act, and what new responsibilities arise when technological systems begin acting on their behalf?

4.

THE GREATEST OPPORTUNITY IS A MORE COLLECTIVELY INTELLIGENT CITY



AI's greatest promise for cities is not simply making individual organisations smarter, but helping the city as a whole become better at understanding shared challenges, making decisions, coordinating action and learning from experience.

What collective intelligence is, and what it is not

By collective intelligence we mean the capacity of a city's people, communities, organisations and institutions to bring together their different forms of knowledge and capability, develop shared understanding, make better decisions and act together on challenges that none could adequately address alone. It is a capacity of the wider urban system, not a technology, platform or centralized decision-making mechanism.

A collectively intelligent city is not one that delegates its thinking to a central AI, nor one in which every resident possesses a digital assistant. It is a city whose people and institutions become more capable of understanding and acting together, supported by technologies that expand their ability to connect knowledge, explore possibilities and coordinate activities.

This is also the central proposition of the AI for Cities Global Summit: **AI becomes an amplifier of the city's intelligence**. The greatest opportunity lies not in substituting artificial intelligence for human or institutional judgement, but in enabling new forms of collective intelligence that bring human, technological and ecological knowledge into more productive relationships.

Collective intelligence

The capacity of a city's people, communities, organisations and institutions to develop shared understanding, make and support legitimate decisions, coordinate action and learn together, achieving more than any could alone.

This is also the central proposition of the AI for Cities Global Summit: AI becomes an amplifier of the city's intelligence.

Four connected capacities

Collective intelligence develops through four capacities that reinforce one another.

Hear more. AI can help a city recognise experiences, concerns and knowledge that existing institutions often miss: those of marginalised communities, young people, newcomers and people who rarely attend public meetings. It can also make the condition of natural systems more visible through environmental monitoring and analysis. An intelligent city is one that can hear more of itself, while recognising that collecting information does not automatically ensure that different perspectives are fairly represented.

Understand better. AI can help connect knowledge scattered across institutions, communities and sectors, identify patterns and relationships, and explore how different problems affect one another. This can strengthen the city's ability to understand complex challenges such as housing affordability, public health, urban heat or changing labour markets, where no single actor possesses a complete picture. The opportunity is not merely better analysis, but the development of shared understanding across those who need to respond.

Decide, coordinate and act together. Most outcomes that matter in a city, including housing, mobility, safety, care, learning, climate resilience and local prosperity, depend on the actions of many different organisations and communities. AI can support dialogue, help compare alternative courses of action and improve coordination across institutional boundaries. Decisions requiring public authority must remain with legitimate institutions, while businesses, communities and residents retain their own capacities and responsibilities for action. Collective intelligence becomes consequential when shared understanding can mobilise these distributed capabilities toward common purposes.

Learn and adapt faster. Cities can become better at observing the consequences of their decisions, identifying unintended effects and revising approaches in response to experience. AI can help connect implementation data with public feedback, environmental changes and other signals that are usually considered separately. This allows cities to learn not only whether a particular intervention worked, but how it affected the wider urban system and whether its benefits can be sustained.

These capacities form a connected process. More information has limited value if it does not improve understanding. Better understanding achieves little if relevant actors cannot make decisions and mobilise resources. And action cannot reliably improve without feedback and learning. AI can strengthen each capacity, but the connections between them depend on relationships, institutions, trust and legitimate authority.

What it already looks like

Elements of this approach are already emerging in cities around the world, although they differ considerably in their use of AI and in how far they connect public participation with decisions and action.

In **Barcelona**, the open-source participation platform *Decidim* has helped residents propose, debate and track decisions about their city's development since 2016. Its first participatory budgeting process brought nearly 40,000 residents into voting and allocated €30 million to

76 selected neighbourhood projects in 2021. Decidim is not itself an AI system, but it demonstrates something essential to collective intelligence: a digital platform becomes meaningful when it is connected to public deliberation, institutional commitments and decisions about real resources.

Bogotá has brought AI more directly into this process. In 2024, its conversational assistant Chatico enabled residents to contribute to the city's 2024–2028 Development Plan through familiar messaging channels. The city recorded approximately 147,000 participation entries from more than 41,000 people; 82% of contributors reported that they had not previously participated in developing such a plan. The process combined AI-supported communication with other forms of public consultation. Chatico has also supported participatory budgeting, helping residents take part in local investment choices. Its significance lies in making participation more accessible and connecting it to established public decision-making processes.

In **Bowling Green and Warren County, Kentucky**, the 2025 What Could BG Be? initiative invited residents to contribute to a long-term vision for their community. Nearly 8,000 people participated over 33 days, registering more than one million responses to ideas about the region's future. Google's Jigsaw Sensemaker, using Gemini, helped organise the contributions into themes and identify areas of agreement and disagreement. This allowed a relatively small community to undertake public sensemaking at a scale that would have been difficult through conventional consultation alone. 70% of participants said they felt more confident their voices would be heard, and 83% reported better understanding of other perspectives. The results are informing the BG2050 planning process, although the longer-term effects on decisions and implementation remain to be established.

A different form of collective intelligence can be seen in **Jakarta**, where PetaBencana has combined community knowledge, digital communication and automated information processing to support disaster response. Residents report flooding and other hazards through familiar social media channels; the platform helps organise and map this information, which can then be used alongside official data by residents, emergency responders and public authorities. Rather than treating people merely as recipients of warnings, the system makes their direct observations part of a shared understanding of unfolding events. It illustrates how collective intelligence can extend beyond consultation into coordinated action, especially when time is critical.

These cases do not demonstrate that cities have already developed comprehensive collective intelligence. Barcelona illustrates the institutional foundations of participation and follow-through; Bogotá shows how AI

can broaden access to civic processes; Bowling Green demonstrates AI-assisted sensemaking at scale; and Jakarta shows how community-generated knowledge can support real-time action. Together, they reveal components of a larger possibility that is only beginning to emerge.

An illustration: one question heard by the whole city

Imagine a neighbourhood deciding what to do with an empty plot of land.

Residents share ideas through public meetings, online consultations and, if they choose, AI assistants that help them formulate and communicate their views. A community organisation works to ensure that older residents, recent arrivals and others who might otherwise be excluded can participate. Young people propose alternative designs, local businesses explore how they might contribute, and environmental data reveals how different uses of the land could affect urban heat, water flows and biodiversity.

AI helps connect these contributions, explore possible scenarios and make disagreements visible. The city publishes an account of what was proposed and how different considerations were weighed, so that residents can question the interpretation before a decision is taken. Legitimate decision-makers then choose a course of action, while community groups, businesses and public institutions may each take responsibility for parts of its implementation.

The process does not end when construction is complete. The city and its communities continue observing how the space is used, whether it improves local well-being and ecological conditions, and what should be changed. What begins as a planning decision becomes an opportunity for continuing collective learning.

None of this requires speculative future technology. AI can expand what is possible, but the underlying achievement would be the city's ability to connect knowledge, participation, legitimate decisions, coordinated action and feedback into a functioning process.

The limits that keep it honest

More sensing does not automatically produce better understanding, and wider participation does not guarantee better decisions. AI-generated summaries may overlook minority perspectives, apparent consensus may conceal genuine disagreement, and sophisticated models may give misleading confidence about uncertain outcomes.

Collective intelligence depends on the quality of relationships and institutions as much as on technological capability. People need opportunities to question how information is interpreted, challenge proposed decisions and understand how their contributions influence outcomes. Some disagreements cannot and should not be eliminated by an algorithm; they require public deliberation and legitimate choices among competing interests.

Nor should every interaction be technologically mediated. Direct encounters, community relationships, professional judgement and the experience of working together are themselves sources of collective intelligence. AI should help strengthen these capacities, not gradually displace them.

The test cases are hard problems

The greatest value of collective intelligence may emerge in addressing challenges that no organisation can solve alone: housing affordability, youth opportunity, public health, climate adaptation, the future of work, neighbourhood regeneration and the restoration of damaged ecosystems.

Consider urban heat. Environmental monitoring can identify where temperatures are rising, but deciding what to do requires knowledge of residents' daily lives, health vulnerabilities, land ownership, transport patterns, infrastructure and local economic conditions. Public authorities, communities, businesses and environmental organisations must act in different ways, often across several years. AI could help them build a shared picture, test interventions and monitor outcomes, but the result depends on their capacity to work together.

This is the larger promise: enabling cities to address complex problems not only with better information, but through new forms of cooperation and collective action that were previously difficult to organise or sustain.

Collective intelligence *between* cities

The same principle applies beyond the boundaries of individual cities. Urban challenges are often shared, but relevant knowledge and experience remain dispersed. Cities repeatedly develop solutions, encounter limitations and learn lessons that others could benefit from, while lacking effective ways to combine these experiences.

Regeneration

Restoring and renewing the conditions that enable living systems to thrive, going beyond reducing harm toward strengthening ecological and social health.

International city networks are already demonstrating the value of cooperation. The **Cities Coalition for Digital Rights**, for example, has helped cities exchange approaches to digital governance and AI. A 2026 comparative study of 161 urban AI initiatives found that participating cities showed greater regulatory depth and more extensive AI policy activity, consistent with the role of international networks in strengthening municipal capabilities through shared learning.

The opportunity extends beyond exchanging good practices. AI could help city networks compare experiments, identify the circumstances in which different approaches succeed or fail, connect specialists and local knowledge across places, and develop shared capabilities for addressing challenges that exceed the resources of individual cities.

This is particularly important for cities that cannot independently sustain large research, technology and strategy teams. Collective intelligence between cities could give them access to knowledge, expertise and learning opportunities that would otherwise remain out of reach.

THOUGHT FOR LEADERS:

Where does your city already possess the knowledge needed to address an important challenge, but lack the relationships, institutions or coordination to turn that knowledge into collective action? How could AI help close that gap?

5.

MORE POWERFUL TECHNOLOGY CAN PRODUCE LESS CAPABLE CITIES



A city can become more technologically intelligent while losing the human, institutional and collective capacities that make it thrive. Technological progress does not automatically translate into urban progress, and some of the most consequential risks arise from the interaction of individually successful technologies across the wider city.

The familiar risks are real, but they are not the deepest

AI raises serious concerns about privacy, bias and discrimination, security, reliability, manipulation and accountability. These require safeguards, transparency, meaningful rights to challenge decisions and effective remedies. Responsible AI frameworks have developed important approaches to these questions, but ensuring that individual systems operate safely and fairly does not guarantee that their combined effects will benefit the city.

The larger challenge concerns the ways technological transformation can alter the economy, institutions, relationships and living systems of the city, sometimes without any individual decision or application appearing particularly problematic. Several distinct risks deserve attention.

Living system

An evolving, interdependent whole whose components and relationships shape its development, adaptation and viability. In cities, this includes people, communities, culture, economic activity, institutions, infrastructure and ecological processes.

The future by default

One of the deepest risks is that a city's direction is increasingly shaped without a corresponding collective choice about the future being created.

One of the deepest risks is that a city's direction is increasingly shaped without a corresponding collective choice about the future being created.

AI enters urban life through procurement contracts, private investment, platform adoption and countless decisions by businesses, institutions and individuals. These decisions are often reasonable within the objectives and responsibilities of those making them. Yet their cumulative

consequences can transform employment, economic opportunity, public space, social relationships and the distribution of power across the city.

We call this *a future by default*: a city whose development increasingly reflects the combined consequences of decisions made by different actors, without a corresponding capacity to understand and shape the emerging whole.

Future by default

A direction of urban development emerging from the cumulative effects of separate technological, commercial, institutional and individual decisions, without sufficient collective understanding or deliberate shaping of the future they are creating.

This is a distinct systemic problem, not simply another form of technological dependency or ineffective governance. A city could maintain competition between providers, protect individual rights and distribute economic benefits relatively fairly while still undergoing profound transformations that its people and institutions had never collectively considered. Choosing the future together does not mean imposing a comprehensive plan from above. It means developing the continuing capacity to recognise emerging changes, deliberate about desirable directions, coordinate action across the city and adjust course as the future unfolds.

Parts improve while the whole weakens

Another risk arises when individual organisations or services optimise their performance without accounting for consequences elsewhere in the city. A transport system can become faster for its users while increasing congestion or reducing the quality of public space. A highly efficient commercial platform can improve convenience while undermining local businesses and community relationships. A city's economic growth can increase the pressures placed on its surrounding ecosystems.

An earlier generation of urban platform technologies offers a revealing precedent. A study by the **San Francisco County Transportation Authority** found that ride-hailing services such as Uber and Lyft accounted for approximately half of the increase in the city's traffic congestion between 2010 and 2016. These platforms created real convenience for individual passengers, yet their combined effect contributed substantially to congestion across the wider city. This was not a study of contemporary generative AI, but it illustrates how technologically coordinated activities can generate consequences that are not captured by measuring the performance of individual services.

AI may intensify such challenges as more decisions become automated, interconnected and responsive to changing conditions. Improving the performance of individual parts is valuable, but urban progress requires understanding what those improvements mean for the health and viability of the larger system.

Benefits concentrate while costs move elsewhere

AI can expand economic opportunity, lower barriers to entry and create new forms of value. But access to these benefits is unequal. Organisations with greater financial resources, expertise, data and market power may be better positioned to capture the gains, while smaller firms, workers and communities face difficult transitions. Cities themselves may differ considerably in their capacity to develop, negotiate and benefit from AI.

New forms of algorithmic coordination can also reinforce existing concentrations of market power. In **the United States**, the Department of Justice has pursued cases concerning the use of shared information and pricing algorithms by competing landlords, including a settlement with RealPage in 2025 (under the settlement RealPage can no longer use real-time confidential data for price recommendations). Such cases illustrate how technologies that improve commercial decision-making for individual firms can raise concerns about competition and housing affordability across urban markets.

The distribution of environmental costs raises another set of questions. AI depends on data centres, electricity, water, land and materials, often supplied in places far removed from those enjoying its benefits. The International Energy Agency's 2025 *Energy and AI* assessment projected that global data-centre electricity consumption could reach approximately 945 TWh by 2030 in its base scenario, more than double the 2024 level, with AI a major driver of additional demand.

The local implications are already significant in some places. According to **Ireland's** Central Statistics Office, data centres accounted for 23% of the country's metered electricity consumption in 2025, compared with 5% in 2015, with consumption rising from 6,973 GWh in 2024 to 7,663 GWh in 2025. This figure concerns all data centres, not AI alone, but demonstrates how concentrated digital infrastructure can become a major factor in energy planning.

The larger question is whether the creation of technological and economic value contributes to shared prosperity and ecological regeneration, or whether its gains and burdens become increasingly separated across communities, cities and regions.

Regeneration

Restoring and renewing the conditions that enable living systems to thrive, going beyond reducing harm toward strengthening ecological and social health.

Trust and legitimate decision-making can erode

As AI becomes involved in decisions affecting people's livelihoods, access to services and opportunities, the ability to understand, question and challenge those decisions becomes increasingly important. Opaque systems can weaken individual autonomy and institutional accountability, particularly when people have little practical ability to refuse their use or seek redress.

Rotterdam provides a concrete warning. The Dutch city introduced a machine-learning system to help identify possible welfare fraud, but suspended its use in 2021 following concerns identified through external review. A subsequent investigation by Lighthouse Reports and WIRED in 2023 found that the system's risk assessments disadvantaged vulnerable groups, including people with limited Dutch-language proficiency. Factors associated with social vulnerability could effectively become grounds for greater suspicion.

The case shows how an apparently technical effort to improve administrative efficiency can undermine fairness and the relationship between public institutions and the people they serve. Greater analytical capability is not necessarily better institutional judgement.

These concerns extend beyond public services. As agents and automated systems increasingly act across organizational boundaries, responsibility can become difficult to trace. Maintaining legitimate authority, explainable decisions and effective accountability is essential not only to protect individual rights, but also to preserve the trust on which collective action depends.

Dependence can grow faster than capability

AI can make individuals and institutions more effective while simultaneously increasing their dependence on technological providers. When essential services, institutional knowledge or operational capabilities become tied to proprietary systems, changing suppliers or recovering control can become difficult and expensive. Dependence may develop

not only through contracts and software, but also through the gradual loss of the skills and organizational knowledge needed to operate independently.

There is a related risk of weakening human judgement. In a 2025 multi-centre observational study in **Poland**, the adenoma detection rate during standard, non-AI-assisted colonoscopies fell from 28.4% to 22.4% after physicians had begun working with AI assistance. The findings do not establish that AI necessarily causes professional deskilling, but they raise a serious question about what happens when people stop regularly exercising capabilities that technology has begun to provide.

For cities, the concern extends beyond individual professional skills. Institutions can become more efficient at processing information while losing the experience needed to interpret complex situations, exercise judgement and respond to unfamiliar circumstances. Their ability to act without technological assistance, or to challenge its conclusions, may gradually diminish.

Human relationships can also be weakened by excessive technological mediation. Direct contact between residents and public servants, professional conversations, community meetings and civic deliberation do more than exchange information. They develop trust, empathy, shared understanding and the habits of cooperation. If AI increasingly replaces these interactions rather than supporting them, a city may become better at responding to individual requests while becoming less capable of acting together.

Why these risks must be understood systemically

These risks interact, but they are not interchangeable. Dependency can constrain a city's freedom to act; weakened skills and relationships can undermine human and institutional capability; fragmented optimisation can harm the wider city; concentrated benefits and displaced costs can increase inequalities; and failures of legitimacy can erode trust. Each can arise through different mechanisms, including decisions that were deliberately planned and technically well executed.

What connects them is a broader possibility: technological capabilities may develop without corresponding improvements in the capacities of people, institutions and communities, or sufficient attention to the larger living systems on which the city depends.

Addressing these risks requires the complementary capacities explored throughout this paper: accountable technological agency (Section 3), stronger collective intelligence (Section 4), deliberate use of the AI Dividend (Section 6), and new forms of urban leadership and collective learning (Section 7).

Agency

The capacity to initiate or carry out actions that influence outcomes. Agency does not, by itself, confer legitimate authority to decide or act on behalf of others.

Two paths for the city

These choices are not made once. They recur in procurement, business strategy, urban planning, service design, technological development and everyday decisions across the city. The following contrasts illustrate what may happen when systemic consequences receive little attention, and what becomes possible when they are deliberately considered.

DIMENSION	BY DEFAULT	BY COLLECTIVE DESIGN
Collective direction	Separate technological choices transform the city without corresponding collective deliberation	The city develops shared understanding and the capacity to shape emerging possibilities
Human capabilities and relationships	Skills, judgement and direct interaction weaken through excessive delegation	AI strengthens people's capacities, relationships and opportunities to cooperate
Systemic performance	Individual services improve while the wider city or surrounding ecosystems bear the costs	Improvements are assessed against their consequences for the larger living system: the health of the whole city comes first
Economic prosperity	Benefits concentrate among powerful actors while others face disruption	New value creation supports broad participation, local enterprise and shared prosperity
Ecological impact	Resource demands and environmental costs remain outside local calculations	Technological development accounts for ecological limits and contributes to regeneration
Autonomy and control	Dependence on a few platforms reduces the freedom of people and institutions to choose	Open systems, portability and institutional capability preserve meaningful autonomy
Rights and collective voice	Opaque decisions and unequal representation weaken trust and legitimacy	Decisions remain accountable, contestable and informed by diverse perspectives
The AI Dividend	Gains are absorbed by rising workloads, cost reductions or concentrated profits	Released and newly created capacity is directed toward human development and collective flourishing

These are contrasting possibilities, not inevitable outcomes or a complete model of urban transformation. Cities may advance in some dimensions while encountering serious difficulties in others. The purpose of the comparison is to make visible the choices that technological adoption alone cannot resolve.

The measure of progress: long-term viability

If the scale and speed of AI adoption are not adequate measures of progress, what should take their place?

We propose **long-term viability**: the capacity of people, communities, institutions and the living systems of which they are part to sustain themselves, adapt, develop and flourish over time.

Long-term viability

The capacity of people, communities, institutions and ecological systems to sustain themselves, adapt, develop and flourish over time and across generations.

AI can help optimise a growing range of activities, but optimisation cannot establish which objectives deserve priority or which conditions of life must be protected. A technically successful intervention may be socially damaging, institutionally unsustainable or ecologically destructive. Long-term viability requires examining the consequences of technological choices across interconnected systems and over different timescales, rather than judging performance through a narrow set of immediate outcomes.

This does not eliminate trade-offs. Cities will continue to face difficult choices between competing needs, economic opportunities, environmental pressures and the interests of different communities and generations. What changes is the basis on which those choices are made: their consequences for the wider living system must be understood, openly considered and reviewed.

Some human activities also deserve protection from excessive automation. Public deliberation, human care, professional judgement and meaningful civic encounter are not merely inefficient processes awaiting technological replacement. They help create the very capacities that allow people and communities to flourish. AI may support these activities, but decisions about whether and how to automate them must consider what human capabilities and relationships could be lost.

The ultimate test is therefore more demanding than whether human and machine intelligence become more capable together. Their combined development must contribute to the flourishing of people and communities while sustaining and regenerating the larger living systems on which urban life depends.

THOUGHT FOR LEADERS:

Beyond the performance of individual AI projects, how would you know whether technological transformation is strengthening or weakening your city's collective capabilities, autonomy and long-term viability?

6.

HOW CITIES USE THE AI DIVIDEND WILL SHAPE THEIR FUTURES



AI can release resources and create new capabilities across urban life. Whether these benefits become a foundation for human development, shared prosperity and greater collective capacity depends on how they are created, distributed and used.

What the AI dividend is

We define the **AI Dividend** as the time, attention, resources and new capabilities that AI frees up or creates for people and organisations. Part of this dividend comes from doing existing work more efficiently. AI can reduce the effort required to process documents, analyse information, respond to requests or coordinate routine activities, releasing human time and institutional resources for other purposes.

Another part comes from capabilities that were previously too expensive, complex or difficult to provide. AI can help a small business develop new products, enable personalised learning at a scale previously unavailable, or make it possible for communities and institutions to analyse challenges that once exceeded their resources. These gains are not simply savings to be redistributed. They can expand what people and organisations are able to accomplish.

The dividend arises throughout the city: in public services, businesses, schools, hospitals, community organisations and households. It is one of the most tangible ways in which the technological capabilities described in Section 3 can contribute to the collective intelligence and capacity for action discussed in Section 4.

But the dividend is not automatic. AI introduces costs of its own, including technology, infrastructure, training, oversight and environmental resources. More importantly, greater productivity does not necessarily translate into additional time for people, better services or wider prosperity. The relevant question is what net benefits actually emerge, who receives them and what they make possible.

Where the dividend can go

The benefits of AI can take different directions, often simultaneously. Organisations may use productivity gains to reduce costs, improve services, expand their activities or increase profits. Workers may gain time for more rewarding tasks, or find that higher productivity expectations absorb the time released. Technology providers may capture a substantial share of the value created, while the people and organisations using their systems retain a smaller part.

AI Dividend

The time, attention, resources and new capabilities that AI frees up or creates, through both greater efficiency and activities that were previously difficult or impossible.

Collective intelligence

The capacity of a city's people, communities, organisations and institutions to develop shared understanding, make and support legitimate decisions, coordinate action and learn together, achieving more than any could alone.

None of these outcomes should be judged in isolation. Reducing costs can release resources for other priorities, and commercial returns can support further innovation. The problem arises when efficiency gains fail to strengthen the capabilities of the people and institutions involved, or when the distribution of benefits undermines wider prosperity and long-term viability.

Research on government illustrates the difference between adopting AI and developing new institutional capacity. The 2026 World Governments Summit report *Rise of the Wise Government* found that more than 85% of surveyed officials in leading AI-adopting governments reported using AI for routine optimisation tasks such as drafting and research. Yet only around 20% reported full-scale transformation of internal processes, and just 4% reported it in public-facing services. Among emerging areas of work, citizen engagement and trust-building were among the least frequently expected to grow.

These findings concern a limited group of national governments rather than cities, but they reveal an important pattern: AI adoption can advance rapidly in established administrative functions while the deeper transformation of institutional roles, relationships and capabilities remains limited. The risk is not that efficiency is unimportant. It is that institutions become faster at their existing tasks without developing the capacities needed to address the challenges they have historically struggled to manage.

The scale of potential gains is nevertheless substantial. In a 2025 **UK government** trial involving more than 20,000 civil servants, participants reported saving an average of 26 minutes per working day with AI assistance. Such estimates show how much capacity might become available, but they cannot establish by themselves whether the time is subsequently used for better services, reduced workloads, stronger institutions or budget savings. That requires a further choice, followed by evidence of what actually changed.

What reinvestment can look like

For public institutions, one of the most important opportunities is to redirect capacity from routine administration toward work that requires human judgement, relationships and coordination. The alternative 2032 imagined at the opening of this paper illustrates the opportunity: time freed by AI in city hall is redirected toward bringing residents' knowledge together with urban heat data and coordinating action on shade, water and green space.

Long-term viability

The capacity of people, communities, institutions and ecological systems to sustain themselves, adapt, develop and flourish over time and across generations.

Early local-government experience shows what this might mean. In **Swindon, England**, a 2024 pilot using AI-assisted documentation (Magic Notes, by Beam) in adult social care reported a 63% reduction in administrative time across the participating staff. Care Act conversations fell from 90 minutes to 35, and write-up time from four hours to 90 minutes. An independent evaluation of a similar application in **Kent** in 2025 found that the time needed to prepare assessment records fell by an average of 2.7 hours per assessment and saved around 7 hours per person per week, while practitioners reported better-quality conversations with the people they supported.

These are early findings, not proof of lasting improvements in care outcomes. But they illustrate a meaningful possibility: social workers could spend less time preparing records and more time understanding people's needs, exercising professional judgement and building relationships. The value lies not only in the paperwork completed more quickly, but in what practitioners may become able to do with the attention they recover.

Other forms of reinvestment can be equally consequential. Public institutions can strengthen preventive services, coordinate responses across organisations, improve oversight and accountability, or develop the capacity for long-term planning and strategic foresight. These are activities that often receive less attention than immediate service demands, even though they are essential to the city's future.

Prevention offers a particularly important direction. **Dutch municipalities**, for example, have developed early-warning arrangements that use information about payment arrears to reach households before financial difficulties become severe. These arrangements rely substantially on institutional coordination and rules-based information processing rather than advanced AI. They nevertheless illustrate the kind of proactive public-service capability that AI could help extend, provided the necessary legal safeguards, professional support and human relationships are in place.

Across the wider city, capacity can also be reinvested in lifelong learning, care, community participation, new forms of cooperation and the regeneration of natural systems. AI may strengthen the ability to monitor ecological conditions, anticipate risks and coordinate restoration, but achieving these outcomes requires people, resources and institutions capable of responding to the knowledge it provides.

The essential shift is from treating productivity as an end in itself to using technological gains to develop capabilities that matter for the long-term flourishing of the city.

Regeneration

Restoring and renewing the conditions that enable living systems to thrive, going beyond reducing harm toward strengthening ecological and social health.

The dividend is also about new prosperity

The economic possibilities extend well beyond reducing the costs of existing activities. AI can lower barriers to entrepreneurship, make specialised capabilities more accessible, enable new products and services, and change what kinds of economic activity can flourish locally.

Evidence of these developments is already emerging. In a 2025 **OECD** study covering more than 5,000 small and medium-sized enterprises across seven countries, 65% of firms using generative AI reported improved employee performance. Around 35% enabled them to scale up (35%), and 29% indicated it allowed them to compete with larger companies. These are self-reported outcomes, and their economic scale remains uncertain, but they suggest that AI can expand the productive possibilities available to smaller organisations rather than merely make existing operations cheaper.

For cities, the opportunity is particularly significant. Local enterprises may gain access to analysis, design, translation, software development and other capabilities that previously required larger teams or external providers. New combinations of technological expertise, local knowledge and entrepreneurial initiative could support new industries, services and livelihoods. AI may also make different forms of local production, care, education and cultural activity economically viable.

None of this guarantees shared prosperity. Established businesses may face disruption, some occupations will change substantially, and value created locally may be captured by platforms or companies based elsewhere. Cities therefore have an interest not simply in accelerating AI adoption, but in strengthening the conditions that allow residents, workers and enterprises to participate in creating and benefiting from new economic opportunities.

This includes education and skills development, access to technology, support for entrepreneurs, collaboration between businesses and research institutions, and the ability of smaller firms to compete and cooperate within larger markets. It also means recognising that economic success cannot be measured solely by the productivity of individual companies. A thriving local economy depends on livelihoods, capabilities, diversity of enterprise and the quality of the wider urban community.

Different roles, shared responsibility

The AI Dividend is created across the whole city, but no single actor controls it. Public institutions can decide how to use the time, resources and capabilities generated within their own operations. They can also influence the wider distribution of benefits through procurement, education, economic development and partnerships.

Businesses have different responsibilities and opportunities. Their productivity gains belong within their own economic decisions, but investing in employees' capabilities, local partnerships and new forms of value creation can also strengthen the markets and communities on which they depend. Community organisations, schools, universities and other institutions can use AI to expand their reach and capabilities, while helping ensure that groups with fewer resources are not excluded from the transition.

Technology providers have a particular role because their products shape both the possibilities available to others and the distribution of economic value. Profitable innovation is legitimate, but cities also need fair competition, portability, trustworthy systems and opportunities for local actors to create and retain value.

A useful practical proposal is to make the AI Dividend visible. Before introducing a significant AI system, a public institution could identify the benefits it expects, the costs involved, who is likely to gain and how any released capacity is intended to be used. After deployment, it could publish a short **reinvestment statement** explaining what was actually saved or created, where the benefits went and what changed for staff, residents and the wider community.

Such statements could also be adopted voluntarily by businesses and other organisations. The purpose would not be to impose a universal reporting obligation before the approach has been tested, but to develop a practical way of distinguishing promised efficiencies from real improvements in capability and well-being.

The dividend is one mechanism, not the only one

The AI Dividend is a major opportunity for turning technological advances into human and collective development. It provides time, resources and new capabilities that can support learning, prevention, care, economic innovation, cooperation and ecological regeneration.

But cities cannot achieve these outcomes through reinvestment alone. Relationships, public institutions, legitimate decision-making, leadership and ecological stewardship must also develop. A city may have abundant technological resources yet lack the trust or collective intelligence to use them well. Conversely, even modest technological gains can become consequential when they strengthen institutions and communities that know how to act together.

The strategic opportunity is therefore larger than distributing the benefits of AI fairly, important as that is. It is to **use the capacities AI makes available to expand what people, communities and cities can become capable of doing together**, while contributing to the health of the larger living systems on which they depend.

Stewardship

The responsible, ongoing care of shared resources, capabilities and living systems, with attention to their long-term health and the well-being of present and future generations.

It is to use the capacities AI makes available to expand what people, communities and cities can become capable of doing together, while contributing to the health of the larger living systems on which they depend.

Living system

An evolving, interdependent whole whose components and relationships shape its development, adaptation and viability. In cities, this includes people, communities, culture, economic activity, institutions, infrastructure and ecological processes.

THOUGHT FOR LEADERS:

What new capabilities could your city develop with the time, resources and opportunities created by AI — and how will you know whether those benefits are reaching the people and communities who need them?

7.

CITIES MUST LEARN TO SHAPE THE TRANSITION TOGETHER



No proven model of the city after AI exists. Cities therefore need new forms of leadership, the capacity to experiment and learn, and stronger cooperation within and between cities to shape their development in an increasingly intelligent and rapidly changing world.

A new kind of urban leadership

AI transformation cannot be delegated to the IT department or the chief technology officer alone. It reaches the entire urban economy, the capabilities and relationships of residents, the functioning of institutions and the ecological systems on which cities depend. The challenge is therefore not simply to introduce and govern a new technology, but to lead a wider transformation of urban life.

The need for stronger leadership is becoming increasingly apparent. UN-Habitat's *Global Assessment of Responsible AI in Cities* identified major gaps in the institutional, governance and technological capacities needed for responsible AI adoption. Its findings suggest that technological experimentation is advancing faster than many cities' ability to evaluate consequences and develop coherent approaches. At Bloomberg CityLab in **Madrid** in April 2026, ten mayors from cities including Bogotá, Boston, London, Nairobi, San Antonio and Tokyo launched the Mayors AI Forum, seeking to give cities a more active role in shaping how AI is developed and deployed, rather than merely purchasing technologies designed elsewhere.

But the transformation required goes deeper than stronger oversight of AI projects. **Urban leadership must increasingly focus on developing the city's collective intelligence and capacity to act:** helping people and institutions understand changing conditions, connect across organisational and community boundaries, navigate uncertainty, and mobilise their capabilities toward shared purposes.

Urban leadership must increasingly focus on developing the city's collective intelligence and capacity to act: helping people and institutions understand changing conditions, connect across organisational and community boundaries, navigate uncertainty, and mobilise their capabilities toward shared purposes.

Collective intelligence

The capacity of a city's people, communities, organisations and institutions to develop shared understanding, make and support legitimate decisions, coordinate action and learn together, achieving more than any could alone.

This means learning to steer a city that continues to change as decisions are being made: sensing emerging shifts early, adjusting course as conditions evolve, and holding steady on the commitments that protect long-term human and planetary flourishing.

Public institutions have a distinctive responsibility in this process. Markets, communities and technological networks can generate innovation and organise many forms of activity, but they cannot alone guarantee the provision of public goods, protection of rights or adequate responses to collective risks. City governments must retain the capacity to make legitimate decisions, manage competing interests and steward long-term public value. Their role becomes more important as AI expands the possibilities for action throughout the wider city.

At the same time, public authority alone cannot produce the collective intelligence the city needs. New forms of leadership must connect institutional responsibility with the knowledge, initiative and capabilities distributed across businesses, communities, universities, residents and technological systems. This requires stronger institutions, but also new relationships between institutions and the societies they serve.

Above all, leaders must develop the capacity to look beyond immediate performance and anticipate the longer-term possibilities emerging from technological change. The task is not simply to manage the transition, but to help the city understand what futures are becoming possible, which ones it wishes to pursue, and how to create the conditions for people and the larger living world to flourish.

Such leadership begins with mayors and senior political and civic leaders because the transformation is systemic. But it cannot end there. Business leaders, community organisers, educators, researchers and young people also have roles in developing the city's capacity to understand and shape its future.

Decisions close to people, and the city's own identity

There is no single desirable model of the city after AI. Cities differ in their histories, cultures, economies, institutional capacities, ecological conditions and aspirations. AI creates opportunities to develop these distinctive strengths rather than simply making urban life conform to the technological models of a few leading providers or cities.

Decisions should therefore be made as close as possible to the people and communities affected, where the relevant capacities and legitimate authority exist. At the same time, many urban challenges cross neighbourhood, institutional and municipal boundaries. Their resolution requires coordination and sometimes decisions at a wider level. The principle is not maximum decentralisation, but an appropriate distribution of responsibility that combines local initiative with the capacity to address shared challenges.

This also means recognising that a city's identity is more than its physical appearance or historical character. It includes its cultural relationships, forms of enterprise, community institutions and ways of organising collective life. Some of these traditions may be worth preserving; others may need to evolve. Designing a city after AI should give its people greater freedom to develop their own future, rather than requiring them to adapt to an externally defined technological model.

Any city can leap ahead

Major technological transitions can unsettle established advantages and create opportunities for cities to develop capabilities that previously required much greater resources. AI may allow smaller businesses, local institutions and communities to undertake activities that were once poss-

ible only for much larger organisations. It may also make new forms of urban coordination, service provision and economic development feasible.

This creates the possibility of *leapfrogging*: cities advancing through new technological and institutional arrangements without having to reproduce every stage of earlier development. Previous transitions, including the spread of mobile payments in parts of Africa and digital public infrastructure in Estonia, demonstrate that established technological pathways are not always necessary. They do not, however, establish that any particular city will leapfrog through AI.

The opportunity is open to cities of different sizes and stages of development. Large cities may be able to combine their research institutions, economic resources and diversity of experience in new ways. Smaller and mid-sized cities may benefit from closer relationships between actors, more accessible leadership and the possibility of experimenting across a manageable urban scale. Yet these advantages are neither universal nor automatic. Smaller cities may also have fewer financial and institutional resources, while large cities can be highly innovative and adaptable.

What matters is not a city's position in an AI adoption ranking, but whether it can develop the capabilities, relationships and institutions needed to respond creatively to changing conditions. The possibility of leapfrogging should encourage ambition and experimentation, not create another competition in which technological speed substitutes for meaningful progress.

Learning by doing, and judging by the right measures

Because no proven model exists, cities will need to discover workable approaches through practical experimentation. But experimentation must be more than a succession of technology pilots.

A pilot can demonstrate that an AI system performs a task accurately or efficiently without establishing that it creates lasting value for the city. An application may reduce administrative costs while weakening professional judgement, improve access to a service while excluding some residents, or generate economic benefits while increasing ecological pressures elsewhere.

Leapfrogging

Developing new capabilities or reaching more advanced outcomes through technological and institutional innovations without necessarily following established paths of development.

Experiments should therefore begin with a clear understanding of the problem being addressed and the human, institutional and ecological outcomes that matter. Their evaluation should consider technological performance alongside effects on human capability, relationships and trust, the quality of collective decisions, the distribution of benefits and costs, institutional autonomy and the health of the wider living system.

Real-world experimentation also requires the capacity to recognise failure and change course. An unsuccessful experiment can produce valuable knowledge if its results are examined openly. Stopping a system that does not meet its objectives, or revising an approach when unintended consequences become apparent, is part of responsible learning.

Cities are already developing institutional precedents for this approach. In **Helsinki**, the Testbed Helsinki initiative creates opportunities for businesses, researchers and public institutions to test innovations in real urban environments, including education, mobility, the built environment and sustainability. Its significance lies not in offering a validated model for AI transformation, but in demonstrating how a city can organise itself as a setting for practical experimentation and cooperation.

Over time, such experiments can generate more reliable knowledge about what works, under which conditions, for whom and with what consequences. Their findings can inform shared approaches, diagnostic tools and eventually standards where these prove useful. The direction should be from tested experience toward more reliable practice, rather than from a supposedly universal model.

Learning between cities

No city holds a complete answer to the challenges of the AI transition. Yet cities share many of the same questions about economic development, human capability, institutional change, technological dependency, public participation and ecological viability.

A distributed international learning process can help cities compare experiments across different urban contexts, examine why certain approaches succeed or fail, identify recurring patterns and adapt insights to different local conditions. This is more ambitious than collecting examples of successful AI applications. It means developing the ability to understand why certain approaches work in one context, where they encounter limitations and what other cities can learn from

Living system

An evolving, interdependent whole whose components and relationships shape its development, adaptation and viability. In cities, this includes people, communities, culture, economic activity, institutions, infrastructure and ecological processes.

them. Such cooperation would allow cities not merely to adopt practices developed elsewhere, but to contribute to a growing body of knowledge about how urban life can flourish in an AI-rich world.

Universities and research institutions can play a particularly important role in this process. They can help cities understand emerging technologies, study economic and social consequences, support experimentation and evaluation, develop the capabilities of public and civic leaders, and connect local experience with wider research and international learning.

The *Astana Declaration on Universities after AI*, adopted by university leaders in October 2026, similarly calls for shared frameworks, experimentation, comparative learning and the exchange of successful and unsuccessful practices. This creates an opportunity for universities and cities to work together on a common challenge: understanding how their institutions and wider communities can develop in an AI-rich world.

The partnership need not be limited to large research universities or major technology centres. Local colleges, vocational institutions, independent researchers, community organisations and international partners may provide the capabilities that a particular city lacks. What matters is creating a continuing relationship between inquiry, experimentation, learning and action.

The longer-term possibility is a new form of collective intelligence between cities: shared capacities for understanding systemic challenges and developing responses that no city could build alone. Such cooperation could be especially valuable for places that lack the resources to pursue extensive research, strategy and experimentation independently.

What this means for each voice

The transition will be experienced differently across the city's voices, and each has a distinctive contribution to make.

City government and public institutions must look beyond improving municipal services toward developing the intelligence and capacity of the whole city. Their responsibilities include protecting public goods, ensuring accountability, maintaining strategic autonomy, convening different actors and enabling legitimate collective choices. A practical starting point is to identify a shared urban challenge that cannot be addressed through departmental action alone and bring the relevant people and organisations into a process of joint understanding and experimentation.

Business has opportunities to create new products, services, livelihoods and forms of local prosperity. Its role is not limited to adopting AI for greater productivity. Businesses can help build the city's technological and economic capabilities, strengthen local partnerships, invest in their people and contribute to addressing collective challenges. Business associations and entrepreneurial networks can become important partners in shaping the city's economic transition.

Civil society and communities can use AI to extend their reach, strengthen their knowledge and organise action, while ensuring that technological development remains connected to lived experience. Their contribution is particularly important in identifying needs that formal institutions overlook, making participation accessible, protecting the integrity of community representation and helping translate collective decisions into meaningful local action.

Residents are not merely users of AI-enabled services or sources of data. They are participants in shaping urban life, bringing knowledge of their neighbourhoods, relationships, needs and aspirations that no technological system can fully reproduce. AI can help them access knowledge, develop capabilities and engage with institutions, but the opportunity must remain grounded in their autonomy, rights and freedom to choose how they participate.

Young people will experience the consequences of today's choices over the longest period and will also be among those creating the next generation of technologies, enterprises and institutions. They need opportunities not only to learn how to use AI, but to develop judgement, creativity, cooperation and the capacity to imagine and build desirable futures. Their meaningful participation in the city's transformation should be a continuing part of urban life, not an occasional consultation.

AI developers, researchers and technology providers have an opportunity to create technologies that strengthen the wider urban ecosystem rather than only individual applications. This means designing systems that are trustworthy, interoperable, accountable and responsive to local needs. Working with cities and communities can also help developers understand problems that cannot be solved through technical performance alone, while shared expectations across multiple cities may create opportunities for more reliable and accessible technological solutions.

These roles are different, but they are connected. A thriving city after AI will require new forms of cooperation between people and institutions whose capacities and responsibilities are complementary.

First moves for any city

There is no validated sequence of steps that every city should follow. Nevertheless, several starting points are worth exploring, adapting and testing.

1. **Make AI a city-wide leadership question.** Bring senior public and civic leaders together with representatives of business, education, civil society and other relevant actors to examine how AI is changing the city's economy, services, relationships and opportunities.
2. **Develop a shared understanding of the transition.** Identify where AI capabilities are emerging, who stands to benefit, where significant risks or dependencies may arise, and which existing strengths could support new forms of prosperity and collective action.
3. **Choose a meaningful challenge for collective experimentation.** Begin with a problem that matters to the city and cannot be solved by one institution alone. Explore whether AI can strengthen the combination of knowledge, human judgement, legitimate decisions and coordinated action needed to address it.
4. **Make benefits and responsibilities visible.** For significant applications, identify expected benefits, costs, accountability arrangements and implications for human capabilities. Track what actually happens, including how the AI Dividend is used, and protect the city's ability to review, modify or leave technological arrangements.
5. **Strengthen the human and institutional foundations.** Invest in the skills, relationships, public participation and forms of cooperation that technological systems cannot provide by themselves. Decide where human judgement, care and deliberation should remain central.
6. **Build partnerships for experimentation and learning.** Work with universities, research organisations, businesses, communities and other cities to test ideas, evaluate results and share lessons, including unsuccessful approaches.

These are not steps toward installing a predefined model. They are ways to begin developing the capacity to understand and shape urban transformation through practice.

Towards a shared learning architecture for many cities

Most cities cannot independently maintain all the technological, research, institutional and strategic capabilities that the AI transition may require. Nor should they have to reproduce the same work separately. Shared resources, comparative learning, leadership development and practical support could allow cities to build on one another's experience while developing approaches suited to their own circumstances.

The AI for Thriving Cities initiative proposes to contribute to this process through an open international dialogue, shared principles, practical experimentation and the gradual development of resources that cities can use and adapt. One possible outcome is AI for Our City: a collaboratively developed set of approaches to help cities consider the transformation across their different voices and living systems. Its content and usefulness must emerge from experience rather than be assumed in advance.

The longer-term ambition is to make such learning accessible to thousands of cities, potentially reaching 10,000 and beyond. This is an aspiration for collective participation and capability development, not a claim that a universal blueprint already exists.

The opportunity is to build more than a repository of tools or good practices. Cities could become participants in a continuing international process of learning how human, artificial and collective intelligence can evolve together in ways that strengthen prosperity, freedom, cooperation and the living planet.

Living planet

The interconnected ecological and life-supporting systems of Earth, of which cities are an integral part and which their activities both depend on and influence.

THOUGHT FOR LEADERS:

What capacities would your city need to develop to shape its future in the AI age, and which partners — within your city and beyond it — could help you build them?

CONCLUSION: AN INVITATION TO BUILD THRIVING CITIES AFTER AI

AI creates opportunities for cities that extend far beyond more efficient services and smarter infrastructure. It could help people and institutions develop new capabilities, create more inclusive and prosperous economies, address challenges that have long resisted conventional approaches, and restore the living systems on which urban life depends. Yet none of these outcomes follows automatically from technological progress. They depend on the choices cities make and, increasingly, on their capacity to understand and shape their own development together.

No city can achieve this transformation alone, and no universal blueprint can yet tell us how to accomplish it. What cities can do is establish shared commitments, begin experimenting with new approaches, and learn from one another as practical experience develops.

The Lichtensteig Declaration on Living Cities and AI

As a starting point, we propose the **Lichtensteig Declaration on Living Cities and AI**, named, subject to the town's agreement, after the Swiss town where the Global Mayors Academy is based. The Declaration expresses ten principles for guiding the transformation:

1. **Choose the future together.**
2. **Strengthen the city's collective intelligence.**
3. **Learn to flourish within a living planet.**
4. **Grow human capability.**
5. **Reinvest the AI dividend, and share it widely.**
6. **Put long-term health before optimisation.**
7. **Decide close to people; honour local identity.**
8. **Keep the city in charge of its intelligence.**
9. **Make AI answerable to people.**
10. **Evolve together, expanding freedom.**

These principles express what we believe cities should strive for, not a prescribed model of how they must operate. They are intended to provide common ground across different cultures, institutional arrangements and stages of technological development, while leaving cities and their communities free to discover approaches that fit their circumstances.

Living system

An evolving, interdependent whole whose components and relationships shape its development, adaptation and viability. In cities, this includes people, communities, culture, economic activity, institutions, infrastructure and ecological processes.

The distinction between principles and implementation is essential. We do not yet know which institutional arrangements, technological solutions or forms of cooperation will best fulfil these commitments. Practical approaches must emerge through experimentation, evaluation and shared learning. Only then can more reliable frameworks and standards develop. Participants pursue the principles within their respective responsibilities, mandates and legal contexts.

How the process will work

The Declaration is intended to begin a process of collective inquiry and action, rather than conclude a discussion.

Its draft will open for public consideration on **31 October 2026, World Cities Day**, during the opening days of the AI for Cities Global Summit. Over the following weeks, representatives of the city's different voices will be invited to examine the principles, challenge their assumptions, identify what may be missing and explore how they might be applied in practice. We also propose a period of critical review and stress-testing, inviting participants to examine potential conflicts, unintended consequences and circumstances in which the principles may be difficult to fulfil.

The aim is to develop version 1.0 through an open process, with changes and substantive disagreements made visible, and to open it for signature on **12 December 2026**, at the Summit's conclusion, online and, subject to local agreement, in Lichtensteig.

Participating in this process does not require agreement on every practical question. There are three complementary ways to contribute:

- **endorse the shared direction and offer feedback;**
- **participate in collective learning and development of the ideas; or**
- **undertake experiments and practical initiatives, sharing evidence of what succeeds, what fails and why.**

There are precedents for developing AI principles through public deliberation, including the **Montréal Declaration for Responsible AI**, which emerged from an extensive process of consultation and co-construction. Our ambition is to extend this participatory approach to the transformation of cities as living systems, engaging residents, businesses, public institutions, communities, researchers and technology developers in a shared exploration of what human- and planet-centred AI could mean in urban life.

Human- and planet-centred AI

An approach to developing and using AI that prioritises human capabilities and flourishing alongside the health and regeneration of the living systems on which people and cities depend.

The process must also be consistent with the principles it advocates. We intend to explore their application in the summit's own AI-supported activities, including transparent accountability, informed participation and respect for people's control over their contributions. Further experiments with interested cities, potentially beginning in **Lichtensteig**, could help test how these commitments translate into actual urban practice. These are proposals to develop and evaluate, not established solutions.

Who we invite

We invite cities and city networks, businesses and entrepreneurs, civil society organisations and communities, universities and researchers, technology developers, foundations, international institutions and residents to take part.

Some will contribute ideas and critique. Others will bring experience, practical challenges, institutional capabilities or resources. Cities may develop their own experiments and share the results; universities may help evaluate them; technology developers may explore new ways of supporting human and collective intelligence; communities may help ensure that emerging approaches remain grounded in the lives and aspirations of the people they affect.

What connects these contributions is the possibility of building a shared capacity to learn about the future of cities while that future is taking shape.

The opportunity before us is not merely to make existing cities more efficient through AI. It is to discover how cities themselves can become more intelligent, capable, prosperous and life-enhancing — places where growing technological power is matched by growing human and collective capability, and where the flourishing of urban life contributes to the flourishing of the living planet.

GLOSSARY

Agency:

The capacity to initiate or carry out actions that influence outcomes. Agency does not, by itself, confer legitimate authority to decide or act on behalf of others.

AI agent:

An AI system capable of pursuing tasks and taking actions with some degree of autonomy, often on behalf of a person or organisation.

AI Dividend:

The time, attention, resources and new capabilities that AI frees up or creates, through both greater efficiency and activities that were previously difficult or impossible. How these benefits are realised, distributed and used determines their wider value.

Cities after AI:

A perspective on the future taking shape as AI becomes embedded in urban life, looking beyond technology adoption toward how cities, their people and institutions can develop under these new conditions.

Collective intelligence:

The capacity of a city's people, communities, organisations and institutions to develop shared understanding, make and support legitimate decisions, coordinate action and learn together, achieving more than any could alone.

Future by default:

A direction of urban development emerging from the cumulative effects of separate technological, commercial, institutional and individual decisions, without sufficient collective understanding or deliberate shaping of the future they are creating.

Human- and planet-centred AI:

An approach to developing and using AI that prioritises human capabilities and flourishing alongside the health and regeneration of the living systems on which people and cities depend.

Interoperability:

The ability of technological systems from different providers to exchange information and work together, without requiring all participants to use a single platform.

Leapfrogging:

Developing new capabilities or reaching more advanced outcomes through technological and institutional innovations without necessarily following established paths of development.

Living planet:

The interconnected ecological and life-supporting systems of Earth, of which cities are an integral part and which their activities both depend on and influence.

Living system:

An evolving, interdependent whole whose components and relationships shape its development, adaptation and viability. In cities, this includes people, communities, culture, economic activity, institutions, infrastructure and ecological processes.

Lock-in:

Dependence on a particular provider, platform, technology or institutional arrangement that makes changing course difficult or costly.

Long-term viability:

The capacity of people, communities, institutions and ecological systems to sustain themselves, adapt, develop and flourish over time and across generations.

Regeneration:

Restoring and renewing the conditions that enable living systems to thrive, going beyond reducing harm toward strengthening ecological and social health.

Stewardship:

The responsible, ongoing care of shared resources, capabilities and living systems, with attention to their long-term health and the well-being of present and future generations.

Voices of the city:

Six overlapping, non-exhaustive perspectives used to guide dialogue about urban transformation: city government and public institutions; business; civil society and communities; residents; young people; and AI developers, researchers and technology providers. They are not mutually exclusive groups or equivalent sources of decision-making authority.

About AI for Thriving Cities

1

BELIEVE · Lichtensteig AI Declaration

What we believe, on one page. Ten principles that people, organisations and cities can sign.

2

UNDERSTAND · Cities After AI

Why it matters. A white paper on how cities can shape the AI transition instead of drifting into a future nobody chose.

3

COMMIT · Living City AI Commitments

What we do, and how we show it. Practical commitments, a system dossier, risk levels, red lines and the evidence for each principle.

4

BEHAVE · Civic Agent Commons

How AI agents behave. Ten rules an agent can load as its instructions, and a card that says whom it represents.



THE HUB
hub.aiforcities.earth

