AI & SOCIETY
https://doi.org/10.1007/s00146-026-03285-2

REVIEW

Responsibility and planning amid growing urban complexity and AI proliferation: many hands, many things, many tasks

Daniele Chiffi[1] · **Stefano Moroni**[1]

Received: 27 April 2026 / Accepted: 20 July 2026
© The Author(s) 2026

Abstract
This paper focuses on issues pertaining to the risks and responsibilities associated with growing complexity and the spread of AI technologies in the urban context. Arguing that such risks are often systemic, it analyses three structural responsibility problems, each linked to a distinct form of dispersion. The *problem of many hands* concerns the difficulty of attributing responsibility due to the dispersion of agency among the multiple actors involved in the design, deployment, and use of technological systems. The *problem of many things* concerns the dispersion of causality because technological risks in the urban context emerge from complex interactions among technical, social, and material factors. We then introduce a novel responsibility-related problem, the *problem of many tasks*, which concerns the dispersion of scope associated with the multiple tasks performed by general-purpose AI systems. Taken together, these forms of dispersion complicate the allocation of responsibility for urban AI risks and motivate the adoption of specific planning strategies. Such strategies require a systemic, participatory, and trajectory-based approach to planning that clarifies responsibilities, ensures meaningful human control over AI, and proactively addresses urban risks rather than merely reacting to isolated hazards.

Keywords Responsibility · AI risks · Systemic risk · Planning

1 Introduction

Cities are complex systems in which many people coexist alongside a wide range of objects of different kinds—natural elements, artefacts and technologies (Chiffi and Moroni 2025). Because of this complexity, dealing with urban risks is very difficult. There may be natural hazards that affect cities (earthquakes, tsunamis, hurricanes, floods and volcanic eruptions), which may be exacerbated by the ways in which humans modify the built environment. For example, soil sealing and river canalisation may intensify flooding, while urban density may augment the human consequences of seismic events. Similarly, coastal development and deforestation may worsen storm surges and landslides, respectively. In addition, there may be risks arising from interactions between people (e.g. crime and terrorism), as well as risks arising from interactions between people and things (e.g. car accidents, traffic exposure risks, noise pollution and building collapses). Furthermore, cities are increasingly permeated by emerging technological tools that interact more and more with citizens (Navarro et al. 2017; Chen et al. 2025). Very often, these technologies are enhanced or directly supported by AI systems (Addad and Al-Taani 2024; Almulhim 2025). Unlike earlier technologies, these AI-based systems exhibit varying degrees of autonomy, adaptability, and interactivity, which can bring significant benefits, but may also introduce new risks or amplify existing ones. For instance, AI-driven urban "digital twins" can improve forecasting for infrastructure and land use, but they also carry the risk that overconfidence in models may amplify planning errors, causing new risks in the real city (Guikema and Flage 2025; Moroni 2025).

Risks that propagate through interconnections and trigger cascading effects across the various elements of a system are called "systemic risks". The concept of systemic risk gained prominence during the 2007–2008 financial crisis, which did not affect only individual firms or banks, but impacted on the entire economic system (Curran 2020). More recently, systemic risk has emerged as a key notion in

✉ Daniele Chiffi
daniele.chiffi@polimi.it

Stefano Moroni
stefano.moroni@polimi.it

[1] Politecnico di Milano, Milan, Italy

Published online: 05 September 2026

risk analysis, particularly in the context of technological hazards and their regulation (Danielsson et al. 2022; Sugawara 2025). Examples of systemic risks include the failure of a national power grid, the collapse of major financial institutions, global supply chain disruptions, and pandemics, all of which can trigger wider economic and social crises. More generally, the term "systemic risk" denotes the possibility that a disruption in some portions of a system may propagate through interconnections, thus causing large-scale failures across interconnected elements, particularly within complex systems, such as cities. In urban contexts, systemic risks are especially important because the dense and heterogeneous sociotechnical networks of cities allow disruptions to cascade rapidly across infrastructures, populations, and services.

To provide an example, in September 2003, "a tree fell on an important power line in Switzerland. The effects rippled down to Southern Italy, causing around €1.2 billion worth of damage. About 56 million people mainly in Italy, but also in Switzerland were without power for several hours".[1] In this case, the cascading effect occurred both internally within interconnected power grid systems and externally, with significant impacts on various other sectors.[2] Such risks may arise through combinations of events, sometimes heterogeneous in nature, whose consequences are partially generated or amplified by human activity, and mediated by economic structures, technological infrastructures, social behaviour, and policy decisions (Renn 2021). The nature of systemic risk requires analysis to move beyond a model of isolated hazards and instead consider the *relational* and *structural* conditions under which risks propagate, intensify, and become potentially collectively uncontrollable. For instance, when a power grid fails, the effects can cascade into water treatment shutdowns, traffic signal failures, and communication breakdowns. They thus create compounding disruptions across multiple essential services within a city.

In this paper, we adopt this systemic perspective on urban risks to critically investigate issues of responsibility attribution and responsibility gaps due to risks that develop within urban systems.

The paper is primarily theoretical even if it may have suitable practical implications. It is based on an extensive, interdisciplinary literature review. Section 2 frames the issue of responsibility in general terms. Section 3 distinguishes three different types of responsibility problems and the related gaps. Section 4 explores how planning theory and practice may address the problems highlighted, especially those associated with AI in urban contexts. Section 5 concludes the discussion.

2 Responsibility and the city

As Braham and van Hees (2018) observe, attributing *responsibility* is a fundamental means by which social, political, and economic affairs are regulated. The basic idea behind holding someone responsible is that people are appropriate recipients of adverse or beneficial treatment, thereby indicating that they should adjust their conduct or maintain it as it is. From the standpoint of justice, it is evident that the system for attributing responsibility must be properly calibrated. Put simply, a misallocation of responsibility is not only likely to be regarded as unjust, but may also convey misleading signals about which forms of behaviour are appropriate (Braham and van Hees 2018, p. 1).

In fact, whenever risks occur, especially within sociotechnical systems, we typically seek to ensure that responsibility for them can be attributed to someone (causally, legally and morally: Bleher and Braun 2022). In this setting, the prospect that risks may exist without any clear attribution of responsibility is itself a further risk, since it may generate or exacerbate harmful outcomes (Nyholm 2023). Responsibility, and in particular moral responsibility, is commonly understood in two main senses: *passive* and *active*. Passive responsibility is retrospective. It is concerned with assessing agents' conduct after an adverse event and with explaining why a particular action or outcome has occurred. Active responsibility, by contrast, is prospective: it concerns guiding present action so as to prevent undesirable future consequences, and it is captured by the question of what ought now to be done (Bovens 1998; van de Poel and Verbeek 2006; Pesch 2015). In regard to emerging technologies, including AI, recent discussions have highlighted various responsibility gaps that affect both passive and active forms of responsibility (Lüthi et al. 2023; Xia 2025). In the case of AI systems, these gaps are often linked not only to the learning capacities of those systems, but also, and more fundamentally, to their opacity, complexity, and unpredictability (Nyholm 2023; Santoni de Sio and Mecacci 2021). In urban contexts, these responsibility gaps become especially problematic because AI systems are embedded in sociotechnical infrastructures that affect large populations and crucial public services and critical systems (Zanotti et al. 2024).

Cities exhibit multiple risks and interdependencies in which potential harms are often widespread, difficult to contain, and may affect certain groups more severely (Curci and Chiffi 2024). Urban AI often operates unnoticed, even though it may *support* or *directly make* numerous background urban decisions. Especially in the latter case,

[1] See https://cordis.europa.eu/article/id/126674-the-cost-of-blackouts-in-europe (accessed July 2026).

[2] On cascading effects in the case of electric blackouts, see also Dobson et al. (2007), Dobson and Newman (2017), Haes Alhelou et al. (2019) and Busby et al. (2021).

attributing responsibility for risks arising from decisions made by AI is particularly challenging, since responsibility is always structured as a chain that must ultimately terminate with human agents.

3 Problems and gaps in responsibility

In this section, we discuss three problems related to responsibility gaps: (i) the "problem of many hands", (ii) the "problem of many things", and (iii) what we suggest calling the "problem of many tasks". This last is a new problem connected to certain types of AI systems operating in complex sociotechnical systems. We analyse the significance of these problems in the context of the risks affecting contemporary cities. The three problems are characterised by three different types of "dispersion": dispersion of human *agency* in the first case, dispersion of *causality* in the second, and dispersion of *scope* in the third. Although these three problems often impose theoretical limits on the allocation of responsibility, our analysis seeks to understand how such problems can, at least partially, be concretely mitigated from the perspective of planning theory and practice.

3.1 The problem of "many hands": dispersion of (human) agency

A major mechanism often invoked to explain responsibility gaps is the so-called "problem of many hands" (Thompson 1980, 2014). This notion concerns situations in which numerous individuals contribute to an activity, frequently within complex sociotechnical contexts, thereby making it difficult to identify where responsibility for a given negative outcome lies and to establish who has been responsible for what (Bovens 1998; van de Poel et al. 2012; van de Poel and Royakkers 2023). This involves not only many individuals within the same group or organisation (public or private), but also many individuals working in different public and private agencies, organisations, and associations. Many engineering disasters exhibit this structural difficulty in assigning responsibility to specific agents. This may happen, for instance, when a bridge collapses, so those involved in the potential attribution of responsibility are designers, contractor, subcontractors, material suppliers, inspectors, municipal technical offices, and maintenance staff. In certain cases, it may be difficult or impossible to attribute individual responsibility because the disaster may have been caused and can be explained not just by the behaviour of a single agent, but by the interaction among the behaviours of different agents. Any single agent may even behave correctly, but the joint interaction of the agents may determine the negative outcome. Consider also the example of urban traffic congestion: it may be caused by the behaviour of motorists, pedestrians, or by those who decide to carry out roadworks at a particular time of year, to mention only some of the agents that may be responsible for the problem. Taken in isolation, no motorist can be considered fully responsible for urban traffic congestion, and his/her behaviour may be negligible, but the collective behaviour of all the motorists in a city in the late afternoon, as they leave their workplaces, may be what creates heavy traffic in a city.

3.2 The problem of "many things": dispersion of causality

The *multiplicity of* (*human*) *agents* may pose problems in the allocation of responsibility, thus producing responsibility gaps. However, responsibility gaps may also be exacerbated by the *multiplicity of interconnected structural and technological* elements. As highlighted by Coeckelbergh (2020, p. 2057):

> "But there are not only many (human) hands; there are also what one could call 'many things': many different technologies. In AI process and history, various software is involved but also more literally various things, material technological artefacts: things that are relevant since they causally contribute to the technological action, and that may have some degree of agency. There are many interconnected elements".

The "many things" problem directs attention not just to the multiplicity of human actors involved in a decision potentially leading to a disaster, but also to the multiplicity of elements that are connected and that may interact in an unexpected way (Lima et al. 2021; Svetlova 2022). For instance, an autonomous vehicle operating in mixed traffic may fail to give way to a pedestrian crossing the road at an informal (i.e., unmarked) location because its sensors misinterpret the situation, resulting in an accident. Similarly, in automated energy grids or urban water management systems, failures may stem not from any single device or human actor, but from the complex interactions among multiple sensors, controllers, algorithms, and human operators. The assessment of interactions among multiple technological artefacts ("many things"), such as sensors, machine-learning systems, and localisation systems, is, therefore, crucial to gain a better understanding of their contribution to causality in conjunction with human actions.

Taken together, the problems of "many hands" and "many things" suggest that disasters are not merely the result of individual failures; rather, they are manifestations of broader dynamics.

3.3 The problem of "many tasks": dispersion of scope

While the problem of "many hands" and the problem of "many things" may also originate from non-AI technologies, we argue that, in the context of most AI systems, an additional problem may emerge. We call this problem the "problem of many tasks", which is associated with General-Purpose AI (GPAI). Unlike traditional fixed-purpose systems that are designed with a specific purpose and are limited to the tasks for which they have been trained, GPAI systems can perform, or be adapted to perform, a range of distinct tasks, including some for which they have been neither intentionally nor specifically trained (Gutierrez et al. 2023). Examples of GPAI systems are *Large Language Models* (LLMs) or *Multimodal Foundation Models* (MFMs), which can be deployed in different fields and systems, such as finance, education, health care, governance, and urban analytics.

In the AI context, a task can be understood as comprising at least the following three elements: a *problem*, a *goal* and an *action* (Gutierrez et al. 2023). The first component is primarily (but not exclusively) descriptive, the setting of a goal is typically normative and actionability constitutes a performative dimension connecting the problem to the purpose. A task changes whenever at least one of these three components changes. With respect to the allocation of responsibility, the difficulty posed by the "many tasks" problem arising from GPAI systems lies primarily in the second, principally normative component. Unlike fixed-purpose AI, GPAI systems do not necessarily have a single intended purpose. They can perform a broad range of tasks, including those that emerge only over time. In many cases, these purposes cannot be fully specified and are, therefore, not subject to adequate judgement or evaluation, whether prospective or retrospective. This gives rise to potential gaps in the allocation of active or passive responsibility.

The application of general-purpose AI systems in different contexts may result in new risks due to the unexpected interaction between AI and the multiplicity of all the other elements of the different sociotechnical systems in which these technologies are deployed. For instance, with fixed-purpose AI systems, surveillance and profiling in cities are limited to specific tasks (e.g., facial recognition, detecting crowd gatherings) and are, therefore, more readily bounded and subject to regulation. By contrast, GPAI can integrate information deriving from different contexts, connect previously separate data sources, and scale surveillance in ways that amplify risks and affect larger populations. Because GPAI systems operate across multiple tasks and domains, responsibility becomes more diffuse: harms may emerge in contexts far removed from the original deployment, making it harder to assign accountability to any individual or institution. For example, a GPAI system deployed in a city may simultaneously analyse traffic flows, energy consumption, public health data, and social media signals. While each task may appear safe in isolation, their combination or the system's autonomous generation of new tasks may have unintended consequences (such as discriminatory routing or privacy violations), illustrating the scalability and pervasiveness of risks unique to GPAI.

3.4 Relations among the three problems

The three problems discussed above are conceptually different even if related because they identify specific dimensions along which responsibility may become dispersed. The problem of "many hands" concerns *who acts*, and, therefore, the dispersion of agency. The problem of "many things" concerns *how outcomes are produced*, and, therefore, the dispersion of causality across human and technological components. The problem of "many tasks" concerns *what a system does*, or may come to do, over time and, therefore, involves the dispersion of scope.

These three problems, however, should not be treated as wholly separate. They are better understood as progressively more complicated versions of a general dynamic. The problem of "many things" typically presupposes and extends the problem of "many hands", since technological artefacts are usually adopted by human agents within sociotechnical systems. The problem of "many tasks", in turn, may involve the other two problems, while adding a temporal dimension. For example, a GPAI system deployed in a city may be used by a multiplicity of actors, thereby raising the problem of "many hands". It may also be embedded in a dense network of sensors, data infrastructures and other digital artefacts, thereby raising the problem of "many things". Over time, however, its scope may drift beyond the purposes for which it was originally configured, thereby generating the problem of "many tasks".

In concrete urban contexts, these three forms of dispersion rarely occur in isolation. Rather, they tend to interact with one another. Dispersion of scope increases the contexts in which human agents and artefacts interact, thereby widening the range of actors potentially involved and thickening the causal networks through which risks may propagate. Conversely, the multiplicity of "hands" and "things" surrounding urban AI makes task drift harder to detect and trace. Consequently, responsibility gaps are most severe where the three forms of dispersion overlap. This is also why the three problems cannot be adequately addressed one at a time. As we argue in the next section, addressing these

problems calls for an integrated planning perspective capable of tackling them together.[3]

4 Planning from an integrated perspective that incorporates AI risks

Savaya and Alfasi (2023, p. 517) note that "in the current planning system, some aspects of planning are negotiated case by case, but responsibility is neglected, and trial-and-error is substantially limited because of top-down land-use regulation". This problem becomes even more disruptive as complexity increases in urban systems and AI becomes widely deployed (Lami and Moroni 2026).

The preceding sections have shown that risks associated with urban contexts are best understood as *systemic*: they emerge from interactions among multiple human actors, technological artefacts, and social arrangements, and they often propagate across domains in ways that are difficult to foresee or contain (Ottenburger and Ufer 2023). These characteristics complicate the attribution of responsibility, giving rise to responsibility gaps associated with the problems of *many hands* and *many things*, and—more specifically in the context of GPAI—the problem of *many tasks*. In this section, we explore the implications of this analysis for urban planning, understood not merely as a technical activity of spatial coordination, but as a distinctive form of governance (Cozzolino and Moroni 2022) concerned with shaping the *conditions* under which urban risks emerge, evolve, and are addressed.

Traditional planning approaches to risk have often focused on identifiable hazards, discrete technologies, or clearly bounded projects. While such approaches remain fruitful in certain cases, they are increasingly insufficient in the face of complex—and frequently AI-enabled—urban systems where risks arise from interdependencies and feedback effects rather than from isolated failures. From a systemic perspective, planning is concerned not only with individual technologies or infrastructures, but also with the *relational configurations* that connect them to social practices, organisational routines, and other technical systems.

This shift has important implications for how planning engages with risk. Rather than aiming solely at the prediction and prevention of specific adverse events, planning must attend to the *structural conditions* under which disruptions can cascade across urban systems. Planning, therefore, plays a crucial role in mediating systemic risk by shaping patterns of interconnection, dependency, and redundancy within a city. This orientation aligns planning more closely with forms of active responsibility insofar as it seeks to guide present decisions in light of uncertain, but potentially far-reaching future consequences (Moroni and Chiffi 2021, 2022).

4.1 Planning and the problem of "many hands"

As we have seen, the problem of "many hands" highlights how responsibility for adverse outcomes may be dispersed among multiple actors involved in complex urban socio-technical arrangements. Urban planning is closely bound up with such arrangements because it typically influences the actions of public agencies, private developers, infrastructure providers, and also technology firms. While planning cannot directly eliminate the dispersion of agency, it can intervene in the way responsibility is structured *institutionally* (given that current institutions frequently remain tied to past choices, path dependency itself renders the consideration of alternative options more costly; see Tennekes et al. 2014).

From a planning perspective, one key implication for accountability frameworks is the need to make responsibility allocation clearer and more explicit within urban processes (too often, *diffused responsibility* becomes *nobody's responsibility*; see Floridi 2016). This obviously does not mean identifying a single agent responsible for every possible outcome; rather, it means clarifying roles, decision-making competences, and oversight obligations in advance. Planning procedures and instruments can thus function as sites where collective responsibility is defined, even in the absence of clear individual contributions to responsibility. By more explicitly embedding responsibility considerations in planning decisions, planners can help mitigate responsibility gaps that might otherwise emerge only retrospectively, after harm has occurred.

Here, a first important point concerns which institutional level of intervention—e.g. national, regional, provincial, or municipal—is responsible for what, and how these responsibilities are delineated (in many countries, and especially in relation to regional and urban issues, this is not always clear[4]). A second important point concerns which responsibilities (both active and passive) should be assigned to *elected* public authorities (for example, regional or municipal councils) and which responsibilities should instead pertain to *designated* agencies or committees (for example, planning units and departments), taking into account the

[3] This paper adopts primarily an institutional and organisational perspective on responsibility, focusing, especially in the next section, on the allocation of roles, competences, and governance arrangements in technologically mediated (including AI-mediated) urban decision-making. This does not, of course, imply that behavioural, interpretive, and distributed-cognitive dimensions of agency are irrelevant. However, we do not return to these aspects in what follows.

[4] On this, see Boelhouwer et al. (2006), Glaas et al. (2010), Landauer et al. (2015), Bae et al. (2016), Den Uyl and Russel (2018), Kösters et al. (2020), Morpurgo (2021) and Khaddour and Deng (2023).

need for coordination, accountability, and effective decision-making. A third point concerns which responsibilities should remain in *public* hands (both elected and designated bodies) and which should instead be explicitly assigned—and recognized as such—to *private* actors (Klein et al. 2017), clearly distinguishing among firms and companies, non-profit organisations, and individual professionals or ordinary citizens, and considering the implications of such assignments for transparency and public oversight (Table 1).

It is worth recalling here that during the COVID-19 pandemic, some countries (e.g. Italy and Spain) experienced serious problems due to the unclear attribution of responsibility in decision-making, for instance between the central government and regional administrations,[5] and, in some circumstances, also between the central state and large cities. Indeed, the pandemic highlighted the shortcomings of some national models of distribution of competencies and responsibilities across the different institutional levels (Nicola and Scaccia 2021; Alber et al. 2022).

4.2 Planning and the problem of "many things"

The problem of "many things" directs attention to the multiplicity of technological artefacts and infrastructures whose interactions with human agents contribute to (urban) risks. In this case, risks rarely originate from a single system in isolation; instead, they arise from the coupling of data infrastructures, algorithms, IoT sensors, physical assets, and human operators (Llorca Albareda 2025). Planning practices that treat technologies as discrete and independent elements may run the risk of overlooking these interaction effects.

A critical systemic planning approach, therefore, requires moving beyond component-based assessments and engaging in analysis of sociotechnical configurations as a whole. This also includes considering how AI systems interact with existing urban structures and infrastructures (Iapaolo and Lynch 2026), how failures in one domain may spread to others, and how design choices can either amplify or dampen cascading effects. In this respect, planning is concerned not only with regulating single technologies, but also with governing the *relations* among them. Such an approach reinforces the view that AI-related risks cannot be adequately addressed without attention to the broader systems in which these technologies are embedded.

Moreover, in cases involving the problem of "many things", it is important that planners promote *meaningful human control* over the system. This is especially important for autonomous technologies, but it also applies to assisted

Table 1 Coping with the problem of many hands

Point	Key questions (challenges)	Planning implications (Coping strategy)	Practical tools
Institutional level	Which level of government (national, regional, provincial, municipal) holds responsibility for specific urban outcomes, and how are these responsibilities delineated? This is often unclear in multi-level governance systems, with the consequent overlapping of duties or gaps	Conduct a formal mapping of responsibilities across institutional levels; create frameworks that clearly assign accountability for specific policy areas or urban projects; include mechanisms for coordination to reduce conflicts or gaps	Responsibility matrices, multi-level governance charts, inter-institutional agreements, formal coordination protocols
Elected authorities vs. designated agencies	Which responsibilities should be assigned to elected authorities (e.g., councils) as opposed to specialized agencies or committees (e.g., planning units, departments)? How to balance active (decision-making) vs. more passive (oversight) roles, while ensuring coordination and accountability?	Define decision-making scopes and oversight obligations for both elected bodies and agencies; assign active responsibilities (policy decisions, approvals) to elected councils and technical duties (implementation, monitoring) to designated agencies; establish clear coordination protocols and reporting lines	Organisational charts, delegation documents, terms of reference, workflow diagrams, reporting templates
Public vs. private actors	Which responsibilities should remain with public actors (elected or designated) and which can be explicitly delegated to private actors, including firms, non-profit organisations, professionals, or citizens? How to ensure the transparency, recognition, and supervision of assigned roles?	Explicitly assign responsibilities to private actors where appropriate, clearly documenting duties and expected outcomes; distinguish among types of actors (commercial, non-profit, professional); implement mechanisms for transparency, accountability, and public supervision; integrate these assignments into planning instruments or contracts	Contracts, memoranda of understanding, public–private partnership agreements, reporting dashboards

[5] See on this Ren (2020), Giovannini and Mosca (2021), Moroni et al. (2023), De Franco and Pacchi (2024), Vese (2023) and Salvati (2024).

Table 2 Coping with the problem of many things

Function of human (meaningful) control	Planning implication (coping strategy)	Practical mechanisms
Actively supervise system behaviour in real time	Ensure humans can monitor technologies continuously to detect anomalies or unintended behaviour	Real-time dashboards, monitoring interfaces, alert systems, sensor integration
Intervene or override decisions when necessary	Design systems so that humans can pause, adjust, or override automated processes to prevent harm	Human-in-the-loop controls, override buttons, emergency shutdown protocols
Contest or appeal against outcomes	Enable mechanisms for questioning, reviewing, or appealing against decisions made by autonomous or semi-autonomous systems	Complaint procedures, appeal boards, audit trails, decision logs, transparency portals
Continuously modify or redesign systems	Ensure adaptive governance whereby planners and operators can update system design to prevent cascading risks or errors	Iterative design processes, system audits, post-implementation reviews, regulatory update cycles

or semi-automated systems.[6] As Holford (2022, p. 834) writes:

> "Determining what is relevant requires human expertise on the one hand, as well as the authority to do so on the other – in other words, meaningful control. As such, responsibility can be placed squarely on the shoulders of human decision-makers, to which we simultaneously confer them the means to meet these responsibilities by providing them the full meaningful control they require – that is, human meaningful control. Partial measures, such as a distributed system of responsibility, must not be used as a surrogate to full human responsibility on the part of the human decision-maker".[7]

Planners can help establish meaningful human control over autonomous or partially autonomous technologies—such as self-driving cars, automated infrastucture monitoring systems, or AI-based decision-making tools—by appropriately shaping the legal, institutional, and technical environments in which these technologies operate. This ensures that humans can (i) actively supervise system behaviour in real time, for example by means of monitoring dashboards or alert systems, (ii) intervene or override decisions when necessary, such as stopping a robot during an unsafe maintenance operation, (iii) contest or appeal against outcomes, for instance by challenging an AI-driven loan decision, and (iv) continuously modify or redesign the systems to prevent harmful consequences, for instance by updating algorithms to eliminate biases or vulnerabilities. Achieving this requires establishing concrete mechanisms such as regulatory oversight, independent audit procedures, accountability frameworks, and ethical review boards, so that human authority and responsibility remain central even in complex sociotechnical contexts (Table 2).

4.3 Planning and the problem of "many tasks"

The problem of many tasks associated with GPAI introduces a qualitatively new source of systemic risk for urban planning. Unlike the problems of many hands and many things, which primarily concern the dispersion of *agency* and *causality*, respectively, among actors and artefacts, the many tasks problem concerns the indeterminacy of tasks over time and the related dispersion of *scope*. GPAI systems are not merely multifunctional; they can acquire new applications and roles over time. Their tasks are not completely fixed at the moment of design, procurement, or planning approval, but may develop and transform, sometimes producing functions unforeseen by developers or planners. The many tasks problem, therefore, arises not from changing (human) uses of AI, but from the intrinsic capacity of GPAI systems to support an expanding and evolving range of tasks beyond those initially envisaged.

This characteristic directly challenges an important aspect of certain planning practices: that technologies and infrastructures can be primarily evaluated and governed on the basis of their *intended function*. Environmental impact assessments and public procurement procedures typically presuppose a relatively stable correspondence between, for instance, a technology and the task it is meant to perform. In the case of GPAI, however, this correspondence becomes unstable. Rather than simply supporting multiple predefined functions, GPAI systems are capable of supporting new and

[6] Consider the Ethiopian Airlines Flight 302 disaster in 2019. The aircraft was equipped with the Maneuvering Characteristics Augmentation System (MCAS) designed to prevent stalling by automatically lowering the nose when a dangerously steep angle is detected. Because of faulty sensor data, the MCAS was activated. In this case, the system repeatedly overrode pilot inputs on the basis of erroneous sensor data, thereby amplifying rather than mitigating the developing emergency (Coeckelbergh 2020).

[7] For the debate on the notion of "meaningful human control", see e.g. Santoni de Sio and Van den Hoven (2018), Veluwenkamp (2022), Cavalcante Siebert et al. (2023), Christen et al. (2023), Davidovic (2023), Santoni de Sio et al. (2023), Robbins (2024) and Calvert (2025).

previously undefined tasks, with each new task potentially influencing the nature and distribution of risks within the urban system as a whole.

From a planning perspective, the novelty of the problem of many tasks lies in the fact that risk no longer originates primarily from a technology's intended use, but from how its uses proliferate and overlap across various domains. Task expansion is not an accidental feature, but a distinctive property of GPAI systems. Importantly, this expansion should not necessarily be understood as a deviation from the original purpose, but as an open-ended evolution of the system's possible roles within urban contexts. As a result, risks may progressively emerge and become systemic without any single planning decision clearly marking the point at which responsibility should be reassessed.

This raises a structural difficulty for responsibility attribution. If risks arise from the accumulating effects of successive transformations in the system's role across time and sectors, responsibility cannot be meaningfully allocated at the moment of initial deployment. Instead, responsibility gaps emerge from what might be described as *temporal displacement*: decisions that are individually reasonable within their original contexts jointly produce outcomes that were neither intended nor foreseen. In this sense, the problem of many tasks reveals a *temporal* dimension of responsibility gaps that is not fully captured by existing accounts focused on multiplicity of actors or artefacts alone. This temporal diffusion of tasks challenges planners, developers, and municipal authorities alike to determine who should be accountable for emergent risks.

For planning, this implies that governing GPAI requires a shift from *task-based authorisation* to what may be called *trajectory-based governance* (compare with, e.g., Colombo 2026, Osmond 2026). Precisely because the relevant tasks cannot be exhaustively specified ex ante, governance must focus on the trajectory through which a GPAI system progressively acquires new roles and tasks over time. Rather than merely asking whether a specific AI application is acceptable for a given task, planners must consider how the system may evolve across tasks and domains over time, and how such evolution will interact with other urban infrastructures and decision-making processes. Here, the notion of trajectory should refer not to a predetermined path of technological development, but to the evolving relationship between the GPAI system and the urban sociotechnical environment in which it is embedded. Trajectories should, therefore, be understood not simply as sequences of technical updates, but as progressive transformations in the role that a GPAI system occupies within the urban sociotechnical environment. Consequently, the object of planning shifts from evaluating isolated AI *applications* to governing *transitions* between successive configurations of the GPAI system's role within the urban environment. Planning thus becomes concerned with governing transitions in the evolving sociotechnical role of GPAI systems rather than regulating discrete and independent uses. For example, a GPAI system in a smart city might initially optimise traffic flows, but later incorporate environmental or public safety data, producing new tasks and effects that were not foreseeable at deployment. Trajectory-based governance would require planners to continuously monitor and govern the progressive reconfiguration of the system's role in relation to various evolving domains.

This reorientation has important implications for planning institutions. It suggests the need for mechanisms that trigger reassessment not only when a technology is modified, but also when its role within the urban system changes. Such shifts in role, however, should not be understood narrowly as technical modifications of the GPAI system itself. Rather, reassessment should be triggered whenever the system's functioning within the urban sociotechnical network changes in ways that alter the distribution and intensity of potential impacts; such changes may include the integration of new categories of data, the extension of the system into new decision-making domains, a significant increase in the population affected by its outputs, or the establishment of new interactions with other AI-enabled or digital infrastructures. The planners' oversight must, therefore, be continuous and reflexive, able to respond to the emergence of new task configurations and their interrelated effects. This does not imply permanent intervention in every technological development, but rather the creation of *institutional capacities* to recognise when a trajectory has entered a qualitatively different phase requiring renewed evaluation. The challenge, therefore, lies not merely in making planning instruments more flexible, but in enabling planning institutions to recognise when changes in a GPAI system's trajectory warrant renewed planning judgement. In this respect, the many-tasks problem highlights a fundamental defect of overly detailed planning instruments and reinforces the need for innovative forms of planning that can accommodate open-ended technological futures. The issue here is not generically ad hoc "adaptive" planning instruments. Rather than being "indeterminately adaptive", planning instruments should be "conditionally open" to new and unpredictable circumstances. Such openness does not imply discretionary adaptability, but rather institutional arrangements that allow revision in response to unforeseen developments while maintaining clear conditions of accountability and democratic control. The aim is not to multiply specific directives and subsequently amend them case by case, but to refrain from fully pre-specifying outcomes, establishing only unequivocal framework conditions while preserving a wide yet clearly delimited domain for adaptive responses (Rauws et al. 2020).

In practice, this conditionally open approach can be supported not only by responsive frameworks and continuous oversight, but also by appropriate democratic practices and

participatory mechanisms that sustain collective scrutiny and accountability in the governance of evolving GPAI systems. Such practices and mechanisms can serve as an essential complement to trajectory-based governance, providing a distributed capacity for scrutiny and early detection of task drift and emergent functions before they accumulate into systemic risks. The point is not so much how new technologies can promote greater digital participation (a topic widely discussed, but of limited relevance here) as it is how genuine in-person democratic-participatory processes can function as proactive mechanisms for addressing the temporal displacement of responsibility inherent in the problem of many tasks. They can become robust and meaningful safeguards that complement planners' oversight and support the reflexive management of evolving technological trajectories in urban systems (Table 3).[8]

5 Conclusion

Urban systems expose people to a great variety of financial, physical (natural hazards), technological, and social risks. Since these risks are interconnected, they must be jointly addressed when designing and implementing effective risk-management plans (Renn et al. 2018). As a result, researchers and practitioners concerned with urban risks should move beyond their specific and often limited domains of interest (Cutter 2021). In this paper, we have focused on issues pertaining to the risks and responsibilities associated with growing complexity and the spread of AI technologies in the urban context.

Our analysis has shown that urban risks are best understood in many contexts as systemic: they emerge from the interactions among multiple human actors, technological artefacts, and social arrangements, often propagating across domains in ways that are difficult to predict or contain (Ottenburger and Ufer 2023). This systemic nature of urban risks complicates the attribution of responsibility, giving rise to gaps linked to the "many hands", "many things" and "many tasks" problems.

Therefore, planning cannot rely solely on traditional hazard-focused approaches or assessments of individual technologies. Instead, it must adopt a critically systemic perspective focusing on the relational configurations and institutional structures within which risks arise and evolve. Key strategies include clarifying roles and responsibilities among public and private actors, embedding meaningful human control over technological systems, and shifting

Table 3 Coping with the problem of many tasks

Issue	Description	Planning approaches	Practical mechanisms
Task indeterminacy and evolution	GPAI systems evolve functionally over time; their tasks are not fixed at deployment. Task drift may produce unforeseen functions and risk accumulation	Adopt trajectory-based governance: continuously monitor the evolving role of the GPAI system; define boundary conditions; maintain planning frameworks that remain conditionally open to unforeseen task evolution	Trajectory review procedures; integration pathway assessments; continuous contextual monitoring; collaborative boundary-setting processes
Temporal displacement of responsibility	Responsibility gaps emerge over time as cumulative effects of task drift produce unintended outcomes	Embed responsibility in ongoing governance structures, with publicly accountable reassessment criteria and oversight responsibilities	Periodic and condition-based reassessment procedures; rolling authorisation cycles; institutional memory mechanisms; review forums
Cumulative risk	Overlapping task expansions gradually produce systemic risks affecting multiple domains (e.g., traffic, environment, public safety)	Map cascading effects and interdependencies; plan safeguards while incorporating normative limits through selectively open planning frameworks	Scenario analysis; cross-sector review panels; democratic scrutiny mechanisms

[8] On this issue, see e.g. Ovadya (2023), Fontes et al. (2024), Helbing et al. (2024), Sanchez et al. (2025), Sieber et al. (2025) and Taylor et al. (2025).

from task-based regulation to a trajectory-based governance which monitors the evolving functions of AI over time. Planning instruments, rather than being indeterminately adaptive, should be conditionally "open" to new and unexpected circumstances, combining foresight with responsiveness. Furthermore, participatory and democratic processes are essential to this approach. Engaging community members in planning can help identify biases, ensure equitable outcomes, and monitor the evolution of AI systems in a city. Such practices transform planning from a reactive activity into a proactive form of governance able to shape the urban risk realm while embedding accountability and democratic safeguards.

Author contributions The authors have contributed equally to the article. The authors did not receive support from any organisation for this research.

Funding Open access funding provided by Politecnico di Milano within the CRUI-CARE Agreement.

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Addad M, Al-Taani S (2024) Leveraging AI for energy-efficient smart cities: architectural and urban planning solutions for sustainable growth. A comparative case study of Amman city and international examples. J Civ Eng Archit 18(12):620–632

Alber E, Arban E, Colasante P, Dirri A, Palermo F (2022) Facing the pandemic: Italy's functional "health federalism" and dysfunctional cooperation. In: Steytler N (ed) Comparative federalism and Covid-19. Routledge, London, pp 15–32

Almulhim AI (2025) Integrating artificial intelligence into smart infrastructure management for sustainable urban planning. Technologies 13(11):1–32

Bae Y, Joo YM, Won SY (2016) Decentralization and collaborative disaster governance: evidence from South Korea. Habitat Int 52:50–56

Bleher H, Braun M (2022) Diffused responsibility: attributions of responsibility in the use of AI-driven clinical decision support systems. AI Ethics 2(4):747–761

Boelhouwer PJ, Boumeester HJFM, Van der Heijden HMH (2006) Stagnation in Dutch housing production and suggestions for a way forward. J Hous Built Environ 21(3):299–314

Bovens M (1998) The quest for responsibility: accountability and citizenship in complex organisations. Cambridge University Press, Cambridge

Braham M, van Hees M (2018) Voids or fragmentation: moral responsibility for collective outcomes. Econ J 128(612):1–19

Busby JW, Baker K, Bazilian MD, Gilbert AQ, Grubert E, Rai V, Rhodes JD, Shidore S, Smith CA, Webber ME (2021) Cascading risks: understanding the 2021 winter blackout in Texas. Energy Res Soc Sci 77:1–10

Calvert SC (2025) Principles and framework for the operationalisation of meaningful human control over autonomous systems. Sci Eng Ethics 31(5):27

Cavalcante Siebert L, Lupetti ML, Aizenberg E, Beckers N, Zgonnikov A, Veluwenkamp H, Abbink D, Giaccardi E, Houben G-J, Jonker CM, van den Hoven J, Forster D, Lagendijk RL (2023) Meaningful human control: actionable properties for AI system development. AI Ethics 3(1):241–255

Chen Y, Tan J, Yang R (2025) Technological relatedness, technological complexity and the development of digital technology in Chinese cities. Appl Geogr 179:1–15

Chiffi D, Moroni S (2025) Uncertainty, innovation and the city: philosophical explorations. Edward Elgar, Cheltenham

Christen M, Burri T, Kandul S, Vörös P (2023) Who is controlling whom? Reframing "meaningful human control" of AI systems in security. Ethics Inf Technol 25(1):1–7

Coeckelbergh M (2020) Artificial intelligence, responsibility attribution, and a relational justification of explainability. Sci Eng Ethics 26(4):2051–2068

Colombo E (2026) The Sentinel System: a first operational implementation of external, trajectory-aware AI oversight. SSRN preprint. https://doi.org/10.2139/ssrn.6675918

Cozzolino S, Moroni S (2022) Structural preconditions for adaptive urban areas: framework rules, several property and the range of possible actions. Cities 130:1–7

Curci F, Chiffi D (2024) Fragility and antifragility in cities and regions. Edward Elgar Publishing, Cheltenham

Curran D (2020) Connecting risk: systemic risk from finance to the digital. Econ Soc 49(2):239–264

Cutter SL (2021) Urban risks and resilience. In: Shi W, Goodchild MF, Batty M, Kwan MP, Zhang A (eds) Urban informatics. Springer, Singapore, pp 197–211

Danielsson J, Macrae R, Uthemann A (2022) Artificial intelligence and systemic risk. J Bank Finance 140:106290

Davidovic J (2023) On the purpose of meaningful human control of AI. Front Big Data 5:1–6

De Franco A, Pacchi C (2024) The space of rules and rules of space: reflecting on local ordinances during the Covid-19 pandemic. City Territ Archit 11(1):1–12

Den Uyl RM, Russel DJ (2018) Climate adaptation in fragmented governance settings: the consequences of reform in public administration. Environ Polit 27(2):341–361

Dobson I, Carreras BA, Lynch VE, Newman DE (2007) Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization. Chaos (Woodbury NY) 17(2):1–14

Dobson I, Newman DE (2017) Cascading blackout overall structure and some implications for sampling and mitigation. Int J Electr Power Energy Syst 86:29–32

Floridi L (2016) Faultless responsibility: on the nature and allocation of moral responsibility for distributed moral actions. Philos Trans R Soc A Math Phys Eng Sci 374(2083):1–13

Fontes C, Al Haddad C, Antoniou C, Lütge C (2024) Assessing the collective challenges of AI: an applied framework focused on people centeredness and public participation. Urban Gov 4(4):245–254
Giovannini A, Mosca L (2021) The year of Covid-19: Italy at a crossroads. Contemp Ital Polit 13(2):130–148
Glaas E, Jonsson A, Hjerpe M, Andersson-Sköld Y (2010) Managing climate change vulnerabilities: formal institutions and knowledge use as determinants of adaptive capacity at the local level in Sweden. Local Environ 15(6):525–539
Guikema S, Flage R (2025) Digital twins as a security risk? Risk Anal 45(2):269–273
Gutierrez CI, Aguirre A, Uuk R, Boine CC, Franklin M (2023) A proposal for a definition of general purpose artificial intelligence systems. Digit Soc 2(3):36
Haes Alhelou H, Hamedani-Golshan ME, Njenda TC, Siano P (2019) A survey on power system blackout and cascading events: research motivations and challenges. Energies 12(4):1–28
Helbing D, Mahajan S, Carpentras D, Menendez M, Pournaras E, Thurner S, Verma T, Arcaute E, Batty M, Bettencourt LM (2024) Co-creating the future: participatory cities and digital governance. Philos Trans R Soc A Math Phys Eng Sci 382(2285):1–10
Holford WD (2022) "Design-for-responsible" algorithmic decision-making systems: a question of ethical judgement and human meaningful control. AI Ethics 2(4):827–836
Iapaolo F, Lynch CR (2026) Rethinking urban AI: a critical posthumanist perspective on intelligence, autonomy, and agency. Urban Geogr 47(1):65–87
Khaddour LA, Deng W (2023) Multi-criteria sustainability risk management for post-war residential reconstruction: the case of Damascus. J Hous Built Environ 38(3):1939–1982
Klein J, Juhola S, Landauer M (2017) Local authorities and the engagement of private actors in climate change adaptation. Environ Plan C Polit Space 35(6):1055–1074
Kösters M, Bichai F, Schwartz K (2020) Institutional inertia: challenges in urban water management on the path towards a water-sensitive Surabaya, Indonesia. Int J Water Resour Dev 36(1):50–68
Lami IM, Moroni S (2026) Urban decision-making is not so much a matter of intelligence, artificial or otherwise. A discussion of AI starting from cost-benefit analysis. Cities 168:1–8
Landauer M, Juhola S, Söderholm M (2015) Inter-relationships between adaptation and mitigation: a systematic literature review. Clim Change 131(4):505–517
Lima G, Cha M, Jeon C, Park KS (2021) The conflict between people's urge to punish AI and legal systems. Front Robot AI 8:1–9
Llorca Albareda J (2025) Uncovering the gap: challenging the agential nature of AI responsibility problems. AI Ethics 5:3857–3870.
Lüthi N, Matt C, Myrach T, Junglas I (2023) Augmented intelligence, augmented responsibility? Bus Inf Syst Eng 65(4):391–401
Moroni S (2025) Insurmountable limitations of city-scale digital twins? On urban knowledge and planning. Comput Urban Sci 5(1):1–12
Moroni S, Chiffi D (2021) Complexity and uncertainty: implications for urban planning. In: Portugali J (ed) Handbook on cities and complexity. Edward Elgar, Cheltenham, pp 319–330
Moroni S, Chiffi D (2022) Uncertainty and planning: cities, technologies and public decision-making. Perspect Sci 30(2):237–259
Moroni S, De Franco A, Pacchi C, Chiffi D, Curci F (2023) Planning and meta-planning to cope with disruptive events: what can be learnt from the institutional response to the Covid-19 pandemic in Italy. City Territ Archit 10(1):1–13
Morpurgo D (2021) The limits of planning: avoidance, concealment, and refusal of religious diversity in Northeast Italy. Plan Theory Pract 22(1):72–89
Navarro JLA, Ruiz VRL, Peña DN (2017) The effect of ICT use and capability on knowledge-based cities. Cities 60:272–280
Nicola F, Scaccia G (2021) The Italian model to fight Covid-19: regional cooperation, regulatory inflation, and the cost of one-size-fits-all lockdown measures. Adm Law Rev 73(1):53–58
Nyholm S (2023) Responsibility gaps, value alignment, and meaningful human control over artificial intelligence. In: Placani A, Broadhead S (eds) Risk and responsibility in context. Routledge, London, pp 191–213
Osmond M (2026) Liability frameworks for agentic AI systems. In: ICML 2026 Workshop AIWILD. https://openreview.net/forum?id=SZhxWYzKb1.
Ottenburger SS, Ufer U (2023) Smart cities at risk: systemic risk drivers in the blind spot of long-term governance. Risk Anal 43(11):2158–2168
Ovadya A (2023) Reimagining democracy for AI. J Democr 34(4):162–170
Pesch U (2015) Engineers and active responsibility. Sci Eng Ethics 21(4):925–939
Rauws W, Cozzolino S, Moroni S (2020) Framework rules for self-organizing cities. Environ Plan B Urban Anal City Sci 47(2):195–202
Ren X (2020) Pandemic and lockdown: a territorial approach to Covid-19 in China, Italy and the United States. Eurasian Geogr Econ 61(4–5):423–434
Renn O (2021) New challenges for risk analysis: systemic risks. J Risk Res 24(1):127–133
Renn O, Klinke A, Schweizer PJ (2018) Risk governance: application to urban challenges. Int J Disaster Risk Sci 9(4):434–444
Robbins S (2024) The many meanings of meaningful human control. AI Ethics 4(4):1377–1388
Salvati E (2024) Fragmentation and intergovernmental conflict during the Covid-19 crisis: the complex relationship between national and regional governments in Italy. Reg Fed Stud 34(3):405–434
Sanchez TW, Brenman M, Ye X (2025) The ethical concerns of artificial intelligence in urban planning. J Am Plan Assoc 91(2):294–307
Santoni de Sio F, Mecacci G (2021) Four responsibility gaps with artificial intelligence: why they matter and how to address them. Philos Technol 34(4):1057–1084
Santoni de Sio F, Van den Hoven J (2018) Meaningful human control over autonomous systems: a philosophical account. Front Robot AI 5:1–4
Santoni de Sio F, Mecacci G, Calvert S, Heikoop D, Hagenzieker M, van Arem B (2023) Realising meaningful human control over automated driving systems: a multidisciplinary approach. Minds Mach 33(4):587–611
Savaya Y, Alfasi N (2023) Complexity, responsibility and care: an intertwined perspective on planning. Plan Theory Pract 24(4):511–527
Sieber R, Brandusescu A, Sangiambut S, Adu-Daako A (2025) What is civic participation in artificial intelligence? Environ Plan B Urban Anal City Sci 52(6):1388–1406
Sugawara SE (2025) Addressing "beyond control": nuclear safety goals in the age of risk governance. J Risk Res 28(3–4):231–246
Svetlova E (2022) AI ethics and systemic risks in finance. AI Ethics 2(4):713–725
Taylor RR, Murphy JW, Hoston WT, Senkaiahliyan S (2025) Democratizing AI in public administration: improving equity through maximum feasible participation. AI Soc 40(5):3653–3662
Tennekes J, Driessen PP, Van Rijswick HF, Van Bree L (2014) Out of the comfort zone: institutional context and the scope for legitimate climate adaptation policy. J Environ Plan Policy Manag 16(2):241–259
Thompson DF (1980) Moral responsibility of public officials: the problem of many hands. Am Polit Sci Rev 74(4):905–916

Thompson DF (2014) Responsibility for failures of government: the problem of many hands. Am Rev Public Adm 44(3):259–273
van de Poel I, Royakkers L (2023) Ethics, technology, and engineering: an introduction, 2nd edn. Wiley-Blackwell, Chichester
van de Poel I, Nihlén Fahlquist J, Doorn N, Zwart S, Royakkers L (2012) The problem of many hands: climate change as an example. Sci Eng Ethics 18(1):49–67
van de Poel I, Verbeek PP (2006) Editorial: ethics and engineering design. Sci Technol Hum Values 31(3):223–236
Veluwenkamp H (2022) Reasons for meaningful human control. Ethics Inf Technol 24(4):1–9
Vese D (2023) Managing the pandemic: the Italian strategy for fighting Covid-19 and the challenge of sharing administrative powers. Eur J Risk Regul 14(1):113–140
Xia YH (2025) Filling the responsibility gap: agency and responsibility in the technological age. Sci Eng Ethics 31(6):1–19
Zanotti G, Chiffi D, Schiaffonati V (2024) AI-related risk: an epistemological approach. Philos Technol 37(2):66

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.