

# AI, Data and Ontology: The Epistemic Turn in Urban Planning<sup>1</sup>

**Stefano Borgo**

Istituto di Scienze e Tecnologie della Cognizione, CNR

[stefano.borgo@cnr.it](mailto:stefano.borgo@cnr.it)

ORCID ID: 0000-0001-6001-2765

**Camilla Perrone**

Università degli studi di Firenze

[camilla.perrone@unifi.it](mailto:camilla.perrone@unifi.it)

ORCID ID: 0000-0001-5829-0416

Received: October 2025 / Accepted: December 2025 | © 2025 Author(s).

This article is published with Creative Commons license CC BY-SA 4.0 Firenze University Press.

DOI: 10.36253/contest-16796

## Abstract

In recent years, cities have undergone a dramatic evolution as they have become increasingly connected, data-rich environments. Urban artificial intelligence (AI) has emerged as a transformative force, empowering city planners, policymakers, and researchers to address complex challenges ranging from resource management to public safety. However, the sheer abundance of data in contemporary urban settings also presents significant ontological challenges, which become even more crucial when undertaking urban planning. These challenges revolve around the conceptualization, classification, and integration of heterogeneous data sources into coherent models capable of supporting intelligent decision-making.

In this article, we examine these ontological issues, discuss existing frameworks that aim to unify fragmented information, and explore the practical implications for urban AI applications. The thesis is that ontologies—structured and formal representations of knowledge—offer a powerful tool to address the challenges outlined above, while serving as a blueprint for defining, categorizing, and interrelating the entities present in urban environments and putting them to work in urban planning.

**Keywords:** Artificial Intelligence; Applied Ontology; Data; City; Complexity

---

<sup>1</sup> This paper expands the work in “I dati, pur essendo quello che ci sembrano, non cessano mai di essere quello che sono”, by S. Borgo, published in “Dialoghi sull’intelligenza Artificiale” Atti dei Seminari del CNR-IRISS, CNR Edizioni 2025, a cura di N. Rampazzo, ISBN: 978-88-8080-766-7.

## 1. Introduction

We all rely on conceptual schemata to understand the world around us. These schemata are influenced, if not determined, by a series of factors: from our biology and cognitive abilities to our social experiences and cultural environment. Schemata offer a selective and, therefore, simplified interpretation of what surrounds us and of how it changes over time. They are useful both for commonsense things, such as to establish that when our cat goes behind the sofa, it does not cease to exist even if we cannot see it, and for complex relations, such as to determine that when prices rise, the money we have in our pocket allows us to buy fewer items.

These schemata are not static. The world around us changes, our knowledge increases, we might be interested in different things or in a different understanding of what happens; there are many reasons that push us to revise our conceptual schemata from time to time, and to adapt them or add new ones. The developments that have lasting consequences, such as the introduction of new tools, from electronic devices to medical treatments, are real challenges for our conceptual schemata. Some innovations, like the introduction of modern travel systems (trains, aeroplanes) and transmission systems (smartphones, satellites), even alter our way of thinking about the world and the meaning of notions that we treat as "natural" and "common sense", such as distance (in the former case) and communication (in the latter).

It would be interesting to analyse the impact of developments in Artificial Intelligence (AI) over the past 15 years on how we understand the world: the new tools made possible by approaches like Machine Learning (ML) and Large Language Models (LLMs), are changing our conceptual schemata. As Romele argues, "these new technologies are not only radically transforming our interactions with the world or our modes of production and consumption but also our worldview" (Romele, 2023).

This article examines one aspect of the complex relationship between us and the world, or rather us in the world. It aims to demonstrate that the innovative tools provided by AI, such as ChatGPT<sup>2</sup> and DeepSeek<sup>3</sup>, revitalise old problems that our conceptual schemata still struggle to address. Essentially, the relationship that AI tools establish with data (in terms of their generation, use and meaning) has not changed from the past. Today, much like in earlier times, we create tools with an opportunistic view of data, and this has consequences—consequences that are more impactful today, and this impact is a problem that calls for attention.

Here we invite the reader to pay attention to data, their use and their meaning. If the innovations brought by AI are significant, as they are, they also generate unjustified, if not unrealistic, expectations. These expectations are only in part attempts to foresee an uncertain future and to anticipate the changes that new technology brings to our way of interacting with the physical world, and to how our social systems evolve. In some cases, which are our focus in this paper, these unjustified expectations are due to the fact that AI tools and us are connected to reality via different media. AI tools rely on data, and the type and quality of these data are not questioned as they should. Therefore, while much technical work is being put into the analysis of algorithms' quality and transparency—an area in which the scientific community is rightly dedicating substantial energy—our focus here will be on *the problem of using data for what they seem to be without reflecting enough about what they actually are*.

## 2. Symbolic AI and Data-Driven AI

When discussing the meaning of data, it is useful to recall a classical distinction in AI: symbolic AI (the semantic-cognitive approach to AI) and sub-symbolic AI (the data-driven approach to AI) [Calegari et al., 2020].<sup>4</sup> Data-driven AI refers to artifacts that essentially use databases as inputs for computation (think of them as long lists of alphanumerical values). These databases collect values of physical quantities in a portion of the world. For example, a sensor installed on a road sends signals that increment a counter, such as indicating how many vehicles have passed on a certain road; a key pressed on a keyboard adds a letter to a sequence, such as building a linguistic expression. Simplifying the example, let us assume that data-driven AI uses the values from the sensor without considering the type of sensor used, its location, how it works, or why it has been put in that place. Similarly, it considers the sequence of typed letters without assuming that it was written to state a meaning. This approach in AI uses data as mere values, not as the result of some unique interaction between things in the world: that sensor and the precise surrounding environment, that device and that specific user. By taking into account how data are

---

<sup>2</sup> <https://chatgpt.com/>

<sup>3</sup> <https://www.deepseek.com/>

<sup>4</sup> The distinction is important but one should not think that there is homogeneity within either of them. For instance, the recent development of AI based on LLMs, called Generative AI, is one approach within the large family of sub-symbolic AI.

created, i.e., the characteristics of the interaction where data is produced, we can turn a sensor from a data generator to an information provider.

Symbolic AI is designed to take the latter position. Not only does it acquire the numerical value from the sensor but also assigns it a semantic meaning taking into account where it is situated, how it works and for what reasons it is used. In other words, symbolic AI aims to work with the meaning of data. There are many ways to work with meaning. It can be as simple as assigning labels to the data. In our example, symbolic AI could relate the number given by the sensor to the expression: "this is the number of times the sensor was activated during the day," or, considering the location and purpose of putting the sensor there, to the expression: "this is the number of vehicles that passed that road and triggered the sensor during the day." The two ways to describe the data carry different information. One is about the internal functioning of the sensor (it activated  $n$  times), the other is about the purpose for having it there (it registered the passage of  $n$  vehicles). Note that the latter is more interesting but also less reliable: the sensor might be activated or not due to unforeseen factors. After all, a sensor is a device and, as such, subject to malfunction, incorrect use or wrong setting, it may be badly designed or triggered by unforeseen causes like electromagnetic interference, people tampering with it, changes in ambient light, inertia after an activation, obstructions, bad communication connection and so on. Knowing how to interpret the data coming from the sensor provides us with significant advantages, but the sensor and its use must be reliable: the quality of the elaboration of an AI algorithm depends on the data we start with. The simple examples we have provided illustrate how complex it is to verify that the sensor interacts with its surroundings as expected and, consequently, that the meaning of the data it generates is correctly informing us about the world.

We have seen that symbolic AI tools are built to use data together with their interpretation.<sup>5</sup> For this reason, their algorithms use computational languages designed to preserve semantics, i.e., to provide not only computations but also how to interpret them in agreement with the interpretation of the input data. It follows that, when the initial information is reliable, what the AI algorithm processes (a statistic, an observation, a deduction) is equally reliable, and the decisions it makes are optimal relative to the available knowledge.

Conversely, data-driven AI tools are designed to process data without focusing on (and sometimes intentionally ignoring) contextual information about the data or how it was collected. This choice opens up techniques of data processing that are more flexible and may follow rules and procedures which cannot be explained from a cognitive perspective. On the one hand, we should clarify that the choice of focusing on pairs of data-meaning or on 'naked' data is not so crisp. On the other, moving away from semantics can be justified when there is a careful selection of the data to use as input (or when data is only partially available or data distribution, a statistical property, needs to be exploited). Suppose the road sensor was activated 5 times yesterday and 15 times today. The data-driven AI tool receives the figures (5 and 15), calculates the average (10), and delivers the following response: the average of the provided data is 10. This number is not connected to anything in the real world; from the tool's perspective, it could pertain to cars, thunderstorms, or football matches. It is crucial for us to comprehend what the response relates to by analysing the supplied data. The symbolic AI tool receives the data (5 and 15 along with their descriptions), calculates the average value (10), and delivers the following response: the sensor indicates an average of 10 daily vehicle crossings (or, there are an average of 10 daily sensor activations). Both tools gather the same data from the sensor: the first treats them as standalone numbers, while the second interprets them as statements about a specific part of the world. In the latter case, it is the presence of the label, utilised by the symbolic tool, that enables the number to be understood as a statement regarding vehicle crossings (or sensor activations) over a specified period of time.<sup>6</sup> Instead, in the first case it is the knowledge of the AI user, not provided to the AI tool, that enables to relate the output number to vehicle crossing.

Today, we utilise many dynamic and complex systems to manage urban traffic, telecommunications, and more. These systems can change rapidly and generate substantial amounts of data. The symbolic approach is complicated to develop and may lack the flexibility needed to handle highly dynamic systems. For this reason, together with computational, technological, and economic factors, the data-driven approach is, in certain cases, the only viable option at our disposal. There are situations, such as image classification and hand-written text

---

<sup>5</sup> What is meaning and how to distinguish data, information and knowledge are complex issue, even the use of the terminology can be confusing (Zins, 2007).

<sup>6</sup> Exploiting even small parts of the meaning of data can have an important impact on AI systems (e.g., see <https://www.technologyreview.com/2024/09/25/1104465/a-tiny-new-open-source-ai-model-performs-as-well-as-powerful-big-ones/>)

processing, where AI symbolic tools do not even provide possible alternatives. For this reason today research aims to integrate these two approaches as much as possible to maximise their potential. The tools created in this manner are referred to as *hybrid systems*.

### 3. Data, AI and the World

Let us return to our theme: the relationship between AI and the external world. Every AI tool provides a solution to a specific problem by processing the data it is given. To greatly simplify, these data may have been collected over the years and kept in databases (e.g., a digital library of written books or a record of meteorological observations), supplied continuously by sensors distributed throughout the environment (e.g., the data coming in everyday from meteorological or traffic sensors), or directly inputted by users through tools such as computer keyboards or mobile phones.

We have observed that these data can be either labelled or unlabelled. If labelled, we must ensure that this classification is done correctly. If not—as occasionally we cannot label all the data or it may be too costly or complicated—we need to ensure that the data are suitable for the AI tool with which we intend to process them. Without this, the AI tool's output might be incorrect; that is, not a solution to the problem we are addressing. Essentially, regardless of whether the data are labelled, it is up to the user of the AI tool to decide whether to use the results based on the reliability of the initial data. Even if we are confident that the tool has processed the data impeccably (and in some cases, we can demonstrate this), it remains the user's responsibility to ascertain, at least in principle, whether the obtained results are appropriate. This holds true whether the data are collected in databases or generated at run-time by sensors or other input devices. When dealing with dynamic and complex systems, we must ensure that these systems are reliable. Similarly, when the data that these systems use are dynamic and complex, we must guarantee their reliability. But how can one ensure the reliability of data?

### 4. Applied Ontology for AI

Let us begin by acknowledging that the introduction of sensors and the accumulation of data are a means to transform the world, to change it from a physical complex system to a cyber-physical complex system, and from the environment in which we live to the environment as we conceive it. This change in complexity is a natural consequence of augmenting the world with sensors: by adding data-generators to the world, we build an information layer over it which becomes part of the world we live in, by using this information in dealing with the world, we turn the latter into a cyber-physical system and, finally, by choosing which sensors to use and where to put them, we put at the center not the world as it is but our conceptualisation of it. In other words, of the world we get what we wish to understand and what we choose to look at.

The choice of what to measure, of which sensors to use, of the locations where to put them, and of how to collect measurements influence our knowledge of the world and the meaning of the data. Data do not exist independently of us; they are by all means the result of our decisions and our capabilities to perceive, interpreting and modifying the world. For this reason, we can assert that to comprehend the data, we must first comprehend the conceptual schemata that guide our understanding of the world. In other words, we need to uncover the conceptual schemata employed when deciding what to measure, which devices to use, where to position them and so on.

We provide data to our AI tools because we believe that these data accurately reflect the world in the aspects that interest us. Thus, we consider some data reliable because, and only as long as, they align with our perception of how the world is. However, we must not forget that data result from a relationship between a device (the sensor) and a part of the world with which the device interacts. This relationship is likely more complex than we might have imagined with our conceptual schemata since, as mentioned, every schema is a simplification. Consequently, when systems and data become intricate and complex, it becomes challenging to understand them for what they actually are. Yet, we need to use them for what they are in order to give reliable input to our AI tools. In recent decades, part of the scientific community has sought solutions to this problem by examining data not as objects to be processed but as narratives about the world to be comprehended. This goal has prompted the integration of scientific methodologies and philosophical analyses giving rise to a unified research area known as *applied ontology*. Today, applied ontology systems, such as DOLCE [Borgo et al., 2022], are international standards<sup>7</sup> and serve as conceptual tools for a cohesive and integrated analysis of: our understanding of the world, the meanings we ascribe to the data we generate, and how we organise all this information. In this way, it has become clear that

<sup>7</sup> <https://www.iso.org/standard/78927.html>

enhancing our conceptual schemata contributes to the development of more reliable AI tools and ensures their proper use. Applied ontology also enhances the semantic and cognitive transparency of information science, and helps to make systems comprehensible and accessible to our analyses. The increasing interest in this type of research within the AI community bear witness to an important step towards constructing increasingly reliable tools and utilising them to their fullest potential. The importance of semantics can be expressed as a maxim that rephrases José Saramago writing: "Data, while being what they seem to us, never cease to be what they are...".<sup>8</sup> We have observed that data form a specific type of thing (or, better, statement), and that data serve to fuel AI tools. We have emphasised the importance of not only using data for our intended purposes but also of deeply understanding their meaning if we aim to use them correctly. This last point is particularly relevant when we intend to use AI tools to manage the environments we have built for ourselves and in which we spend most of our private and social lives.

## 5. The Abundance of Objects in the City

In modern cities there are many sources of data, from cartographic maps to interactive panoramas (e.g. Google Street View), from distributed sensors measuring flows of air, water and traffic to participatory GIS. We have learned earlier that all these data are collected by detecting devices which are built with some vision about how the world works and what we want to know about it. These visions do not need to be mutually coherent due to the variability of purposes, and are likely not compatible: some data are collected in discrete time (some at low frequency, others at high frequency), others in continuous time; some data are about a service as a structured object (e.g., the company's commitment to provide fresh water), others about the same service as a processual entity (activities like pumping and pressure control performed to deliver fresh water). Furthermore, the material level is only one side of the coin; ignoring the cognitive and the social levels would jeopardise our capacity to interact and coordinate with our world.

Consider, for example, an area for transportation services, a building present in any city and that we generically call a station. It is certainly a material entity composed of walls, openings, empty spaces and other elements aimed to make waiting comfortable and to share information about how the services work. Being a functional entity with a specific purpose, it is based on a design which was realized purposefully, that is, it is an artefact. Such artefact is justified by what is planned to happen there, which reminds us that it is also a location, a place to go, possibly stay for some time and to leave. This means that it is a kind of container as well as an orientation point. All this makes it a social entity which must be recognisable as such, and an institution that requires personnel, roles, duties and regulations.

Clearly a station is not a special case, our analysis applies to practically everything in our complex cities. One starts from the obvious and apparently inconsequential recognition of something in the city, and quickly ends up in a large network of interconnected things and concepts. Some of these are basic, like an amount of matter or an object. Others are complex to define, e.g., design, location, layout, structure, container, artefact, landmark, institution, organisation, and so on. Fortunately, we can turn to applied ontology to put order into this messy list of heterogeneous, yet interconnected, things: design is a concept, location is a portion of space, layout is a distribution of matter, structure is a relation across components, container is a functional object, artefact is an intentional object, landmark is a role, institution is a social entity and so is organisation.

Given that our cities are so rich in terms of what is there, we need to ask ourselves whether we take into account what the data we collect really talk about: Are we making clear which entities the data inform us about? Do we disentangle the data we collect in relation to the networks of entities that populate our world? With some embarrassment we must recognise that the standard developer of AI tools does not even think about the complexity of our world. And the standard user of AI tools follows suit. This should not come as a surprise since these people have important technical skills but have not been trained to analyse the subtleties of our world: they have learned to build decision-making systems, or to apply them. The most we can hope today is that in deploying their tools they take an engineering attitude: once a solution is suggested by whatever AI system they apply, that solution is suitably tested before being implemented in systems as complex as our cities.

The lack of attention to the relationships among our world, the data we collect and the powerful AI technology we have access to, leaves us vulnerable to misunderstanding the data and to misusing AI technology, and with

---

<sup>8</sup> The original sentence dates 15 August 1998 and states: "(T)hings, while being what they seem to us, never cease to be what they are..." (Saramago, 2018)

that come upsetting disappointments, unexpected failures and even dangerous scenarios. It is a call to rethink what data really tell us and how to use them.

## 6. Toward Ontology-Aware AI for Urban Environments and Planning

In light of what we have discussed so far, it becomes apparent that the abundance of data in urban environments, while seemingly a resource, can become a liability if not adequately framed within coherent conceptual structures. As has been highlighted by recent interdisciplinary research, AI systems designed for complex settings such as urban planning require not only access to vast datasets but also the ability to discern and formalise the relational structures between physical, functional, institutional, and perceptual aspects of the city (Stufano Melone et al., 2024). Urban planning is inherently a cognitive negotiation among multiple agents, where each stakeholder contributes with different values, constraints, and expectations. Traditional decision-making frameworks often rely on creative synthesis to resolve competing demands. However, when these human-centred processes meet data-driven AI systems, there is a risk of reducing the richness of urban knowledge to flattened abstractions (Zins, 2007; Romele, 2023), a concern repeatedly echoed in participatory urban design discourse.

In recent years, scholars have begun advocating for ontology-based AI systems that can support interdisciplinary planning tasks (Borgo et al., 2021). These systems act as mediators between raw data, human interpretation, and contextual knowledge. The aim is to go beyond simple numerical optimisation or predictive modelling and instead support the interpretive and dialogic nature of urban planning, especially in its “last mile”: the stage where divergent citizen needs must be synthesised into actionable interventions.

For instance, participatory scenarios, such as determining where to plant urban greenery, involve not just ecological or infrastructural parameters, but also aesthetic, functional, and emotional constraints. Some residents may prefer trees near their homes, while others may value open parking, and still others may desire central green plazas. Here, traditional data-driven AI falls short. Only by embedding ontological reasoning – that is, representing and reasoning over what a “tree”, “parking”, or “plaza” means in a social, spatial, and temporal context – can AI tools genuinely assist in this complexity (Mattern, 2021).

This calls for a fundamental paradigm shift: moving away from viewing AI merely as a solution engine and instead treating it as an epistemic collaborator. Such a collaborator must be capable of distinguishing between binding constraints—such as legal zoning rules—and negotiable preferences, like the color of urban furniture. Hybrid frameworks, which integrate symbolic reasoning with data-driven inference, are increasingly being explored in decision-support applications (Calegari et al., 2020). These approaches offer a promising path toward developing explainable and socially grounded urban intelligence.

Yet, these efforts remain in their infancy, and the full potential of ontology-based modeling is still not well understood. To appreciate how ontological modeling enhances AI capabilities, consider the new types of questions these systems enable us to address. Because ontological models explicitly represent the entities, their relationships, and their behaviors within a given scenario, they allow for much richer analyses than traditional approaches. For example, when examining a city’s traffic system, ontology-based models move beyond simple forecasting and balancing of input/output parameters, such as vehicle flows, parking occupancy, and pollution levels at key locations. Instead, they enable us to ask and investigate more nuanced questions and interconnected issues:

- How might traffic patterns shift in response to ongoing or scheduled cultural and social activities? (This could involve cross-referencing traffic data with information from cultural mailing lists, library departments, and local associations.)
- Are the health and security services in a given area sufficient, given the change in traffic, number of people gathering and the nature of activities in those locations? (This requires matching facility capacities with the expected behaviors of different population groups, like teenagers, families, sports fans, etc.)
- How might changes in weather forecasts alter public behavior, potentially overloading or even overwhelming urban services in the area?
- What are the possible actions that can be taken to reduce the above issues and which stakeholders should coordinate to minimise the problems?

By enabling such inquiries, ontology-based models open up new possibilities for understanding and managing the complex, interconnected dynamics of urban life.

Ultimately, what emerges is the need for ontological models and interdisciplinary workflows, designed not merely to process data but to reflect the conceptual richness of the environments we inhabit and shape. In this sense,

applied ontology is not a luxury—it is the responsible foundation on which to build trustworthy and meaningful AI systems for our cities.

## 7. Complexity, Reflexivity, and the Epistemic Turn in AI Urban Planning

Reframing AI's role in the city requires engaging with a broader epistemological shift, one that acknowledges the complexity and indeterminacy of urban systems. Drawing on the complexity theory of cities, authors such as Batty (2007; 2013), Portugali (2021), and Bettencourt (2014) have demonstrated that cities are not static, linear constructs, but rather adaptive, emergent systems governed by feedback loops, nonlinearity, and path dependency. In such contexts, the governance of urban processes cannot be reduced to control mechanisms or closed models; it must instead embrace uncertainty and multiplicity.

In this light, urban AI cannot be a decision-making mechanism or an actuator of automated decisions as we see it today. Instead, it should develop to become a co-participant in a distributed and evolving epistemic ecosystem. This aligns with the conceptual innovations introduced by Federico Cugurullo, who identifies the emergence of AI Urbanism as a radical departure from conventional smart city paradigms (Cugurullo et al., 2024). While the smart city focuses on automating optimisation through data collection and automation, AI urbanism highlights the multiple effects of emerging AI systems with increasing autonomy, capable of influencing governance, infrastructures, and urban imaginaries. Rather than replacing human decision-making, these systems should reconfigure the conditions under which decisions are made, mediating between data, knowledge, and collective negotiation.

This shift has profound ontological and political implications. In his critique of Frankenstein Urbanism, Cugurullo (2021) warns against constructing cities from disjointed technological modules without a coherent ontological framework. The outcome is a fragmented and dysfunctional urban environment, technically efficient but socially disjointed. More recently, in his analysis of AI ideology, Cugurullo (2025) demonstrates how AI systems embody implicit ideological logics, converting political and ethical issues into algorithmic decisions. Within this context, AI becomes not just a tool but an ideological actor—an entity that determines what is visible, what can be governed, and ultimately, what is *real*.

Such critiques align with rising concerns about algorithmic bias, opacity, and accountability (O'Neil, 2016; Safransky, 2020). Predictive policing, smart zoning, and resource allocation algorithms often embed and replicate structural inequalities, resulting in urban futures that are less diverse and more deterministic. Without reflexive design and transparent ontology management, AI can unintentionally diminish the epistemic diversity essential for fair and inclusive urban planning.

In response, recent research promotes ontology-aware and co-constructed AI systems models that incorporate normative and symbolic knowledge and are developed through interdisciplinary and participatory approaches. These systems can support pluralistic reasoning, accommodate conflicting worldviews, and enhance explainability, particularly in value-laden decision-making contexts (Calegari et al., 2020). Rather than reducing complexity, these models aim to navigate it, allowing the city to remain a space of negotiation, emergence, and collective imagination.

## 8. Conclusions

This article has addressed the profound ontological and epistemological challenges that arise when artificial intelligence engages with the urban domain. By revisiting the distinction between symbolic and data-driven AI, we have highlighted how the meaning and reliability of data cannot be taken for granted but must be grounded in conceptual frameworks that reflect the complexity of cities as socio-technical systems. Applied ontology emerges as a powerful tool to bridge raw data and urban knowledge, transforming sensors and databases into genuine sources of information, and thus enabling AI to act as an epistemic collaborator rather than a mere computational engine.

In doing so, we have finally situated the debate within the broader perspective of complexity theory, where urban systems are viewed as adaptive and emergent, requiring tools capable of navigating uncertainty and multiplicity. Within this horizon, the emergence of urban AI, as theorised by scholars such as Batty and Cugurullo, signals not only a technical shift but also an ontological and political one, where AI acts as a co-producer of urban reality. Such a shift demands awareness of the risks associated with algorithmic reductionism and ideological embedding, but also opens up possibilities for more reflexive, participatory, and ontology-aware approaches to urban planning.

In this sense, the contribution offered here aligns with the aims of this special issue *PlanAir* (which stand for “Planning and Artificial Intelligence in Urban Research”) and the debate on urbanism and AI. We frame AI not only as a set of optimisation tools but as a conceptual and practical challenge for urban studies, planning, and design. By foregrounding ontologies as mediators between data, human interpretation, and collective negotiation, the article highlights how urban AI can be mobilised to support more transparent, plural, and democratic forms of urban governance. Far from considering AI as a dystopian threat or utopian solution, we argue for a pragmatic and critical perspective in which artificial intelligence becomes an integral part of the epistemic turn in urbanism, capable of enhancing our collective capacity to imagine and shape future cities.

## References

- Batty, M. (2007). Cities and complexity: understanding cities with cellular automata, agent-based models, and fractals. The MIT Press.
- Batty, M. (2013). The new science of cities. MIT Press.
- Bettencourt, L. M. (2014). The uses of big data in cities. *Big data*, 2(1), 12-22.
- Borgo, S., Borri, D., Camarda, D., & Stufano Melone, M. R. (2021). An ontological analysis of cities, smart cities and their components. In *Technology and the City: Towards a Philosophy of Urban Technologies* (pp. 365-387). Cham: Springer International Publishing.
- Borgo, S., Ferrario, R., Gangemi, A., Guarino, N., Masolo, C., Porello, D., Sanfilippo, E. M., & Vieu, L. (2022). DOLCE: A descriptive ontology for linguistic and cognitive engineering. *Applied Ontology*, 17(1):45–69.
- Calegari, R., Ciatto, G., & A. Omicini (2020). On the integration of symbolic and sub-symbolic techniques for xai: A survey. *Intelligenza Artificiale*, 14(1):7–32.
- Cugurullo, F. (2021). “Frankenstein Urbanism: Artificial intelligence and the city”, in Marvin, S., Luque-Ayala, A., and McFarlane, C. (eds), *Smart Urbanism: Utopian Vision or False Dawn?*, Routledge, London, pp. 200–215
- Cugurullo, F., Caprotti, F., Cook, M., Karvonen, A., M'Guirk, P., & Marvin, S. (2024). The rise of AI urbanism in post-smart cities: A critical commentary on urban artificial intelligence. *Urban studies*, 61(6), 1168-1182.
- Mattern, S. (2021). *A City is Not a Computer: Other Urban Intelligences*. Princeton, NJ: Princeton University Press.
- Portugali, J. (Ed.). (2021). *Handbook on cities and complexity*. Edward Elgar Publishing.
- Romele, A. The datafication of the worldview. *AI & SOCIETY*, 38:2197–2206, 2023.
- Saramago, J. (2018). *Último Caderno de Lanzarote. O diário do ano do Nobel*. Ed. Porto Editora, pp. 272. ISBN: 978-972-0-03128-0
- Stufano Melone, M. R., Borgo, S., & Camarda, D. (2024, July). City Interactions in Urban Planning: The Square Example from an Ontological Analysis Point of View. In *International Conference on Computational Science and Its Applications* (pp. 448-455). Cham: Springer Nature Switzerland.
- Romele, A. (2023). *Digital habitus: a critique of the imaginaries of artificial intelligence*. Routledge.
- Zins, C. (2007). Conceptual approaches for defining data, information, and knowledge. *Journal of the American society for information science and technology*, 58(4), 479-493.