Goodchild *et al. Urban Informatics* (2024) 3:16
https://doi.org/10.1007/s44212-024-00048-6

Urban Informatics
The International Society for Urban Informatics

ORIGINAL ARTICLE | Open Access

Digital twins in urban informatics

Michael F. Goodchild[1*], Dylan Connor[2], A. Stewart Fotheringham[2], Amy Frazier[3], Peter Kedron[3], Wenwen Li[2] and Daoqin Tong[2]

Abstract
Urban informatics appears to be a suitable area for the application of digital twins. Definitions of the term share some characteristics, but these definitions do not agree on what exactly constitutes a digital twin. The term has the potential to be misleading unless adequate attention is paid to the inherent uncertainty in any replica of a real system. The question of uncertainty is addressed, together with some of the issues that make its quantification problematic. Digital twins for urban informatics pose questions of purpose, governance, and ethics. In the final section the paper suggests some research issues that will need to be addressed if digital twins are to be successful.

Keywords Digital twins, Uncertainty, Governance, Data integration, User interface

1 Introduction

The term "digital twin" is now firmly established within the field of urban informatics: papers have appeared (Batty, 2018; Cureton & Dunn, 2021; Dawkins et al., 2018; Fotheringham, 2023; Wildfire, 2018), commercial companies use the term in advertising their products (e.g., Esri: https://www.esri.com/en-us/digital-twin/overview), and numerous workshops have been held (e.g., by ASU's SPARC: https://sgsup.asu.edu/node/9358). Yet much uncertainty remains in the domain of urban informatics, about the definition of the term, the purposes that a digital twin is intended to satisfy, the ownership and governance of digital twins, and the trust that users feel able to place in their models and predictions. This paper is intended to clarify many of these issues, to identify current research needs and unanswered questions, to encourage open discussion, to increase awareness of the ethical issues raised by digital twins, and to identify ways in which the next generation can be better prepared for the development and successful use of digital twins in urban informatics.

In February 2023 the Spatial Analysis Research Center (SPARC) of Arizona State University's School of Geographical Sciences and Urban Planning organized a workshop on digital twins that included eight keynote speakers from around the world, lightning talks, and extended discussions. The workshop website (https://sgsup.asu.edu/node/9358) includes the presentation slides and recordings of the sessions and the final report of the workshop. This paper synthesizes many of the issues raised at the workshop while focusing primarily on urban informatics.

The next section of the paper discusses the definition issue, and the relationship between the earliest use of the term in manufacturing and its applications in earth science and urban informatics. This is followed by an extensive discussion of accuracy in digital representations and the uncertainties present in all sources of data and in the models used for prediction (recalling George Box's famous dictum "All models are wrong but some are useful", Box, 1976).

The fourth section of the paper addresses the broader contextual issues of governance, ownership, ethics, and the factors underlying trust in digital twin predictions. The use of digital twins—and the use of the term itself—raises numerous issues of an ethical nature, and while

*Correspondence:
Michael F. Goodchild
good@geog.ucsb.edu
[1] Department of Geography, University of California, Santa Barbara, USA
[2] School of Geographical Sciences and Urban Planning, Arizona State University, Tempe, USA
[3] Department of Geography, University of California, Santa Barbara, USA

some of those issues might be addressed through technical solutions, more broadly it is the designers, developers, and users of digital twins who must take responsibility for ensuring that practices are ethical.

The final section of the paper examines the research questions that will need to be resolved if the promise of digital twins is to be realized. Comments are also made about educating the next generation of digital twin developers and users.

Throughout the paper it is assumed that the digital twins under discussion are geospatial: that is, they differentiate locations by the data used to characterize them, by the models that are applied, and by the predictions that are made. Thus the digital twins in this paper are intimately linked to the technology known as GIS (geographic information system; see for example Longley et al., 2015). This is consistent with the assertion by Esri that "A digital twin is a virtual representation of the real world, including physical objects, processes, relationships, and behaviors. GIS creates digital twins of the natural and built environments and uniquely integrates many types of digital models" (https://www.esri.com/en-us/digital-twin/overview).

2 Definitions

The idea of a replica of some physical system has a long history, and people were building analog models to scale long before the advent of digital technology. Engineers have built scale models of dams, scale models of cities have been constructed, and the building and operating of model railways have been popular hobbies for over a century. However the development and growth of digital technology over the past half century has provided an entirely new level of interest in modeling that has been fuelled in part by a simultaneous growth in the availability of fine-resolution data and computing power.

The term "digital twin" first appeared around 2000 in connection with industrial design and manufacturing. The definition given by Grieves (2016) dates from that period and captures the meaning commonly assigned to the term in industrial design:

> *"A Digital Twin is a set of virtual information constructs that fully describes a potential or actual physical manufactured product from the micro atomic level to the macro geometrical level. At its optimum, any information that could be obtained from inspecting a physical manufactured product can be obtained from its Digital Twin."*

Several key points arise from this definition. The second sentence is in effect a restatement of the Turing Test devised by Alan Turing in 1950 and originally termed the Imitation Game (Turing, 1950). The definition implies that the information obtained from the digital twin would be identical to equivalent information that might be obtained from the physical product. To achieve such perfection it would be necessary for the digital twin to be a perfect replica of the physical system, down to the "micro atomic level." Such perfection is clearly impossible even with the greatest conceivable precision of manufacturing, and is far more problematic when the concept of a digital twin is implemented in fields such as urban informatics that must address both built forms and the environmental and social systems that modify them. Instead it would be necessary to invoke some concept of accuracy, and to identify the level of accuracy required for satisfactory performance in some defined use case—the concept often termed "fitness for use." This issue is explored in detail in the third section of the paper.

The Grieves definition links the concept of a digital twin to a range of scales: "from the micro atomic level to the macro geometrical level." Scale is also an important concept in urban informatics, where processes range from the individual scale of observable human behavior to the emergent properties that characterize entire cities and societies. Thus a digital twin for urban informatics will be faced with the challenge of reconciling these very different processes, and may perhaps be a platform for the discovery of new emergent properties. However the constant presence of uncertainty will make it impossible to determine if these new properties are actually present in reality, or merely the artifacts of various forms of uncertainty.

The term "replica" occurs frequently in definitions of digital twins, referring both to the form of a system and to the processes that control and modify it. Form is captured in a digital representation, and processes are represented in the software that operate on the data. Both would have to be perfect representations if the predictions of the digital twin are to hold into the future. Moreover, many definitions also emphasize the two-way relationships between the digital twin and the physical system that it is intended to replicate: the role of the digital twin in predicting future states of the physical system, and the role of the physical system in constantly updating and correcting the digital twin with new and revised information.

While most definitions focus on the role of digital twins as replicas, few address the many purposes for which digital twins might be built, from basic mirroring to monitoring, modeling, predicting, and even autonomously learning or forecasting behaviors. The ability to experiment on a virtual system rather than a physical one has clear advantages when failure in the physical system carries severe risks, as in the design of aircraft or of crewed space missions. Virtual cadavers have clear advantages over real ones in training surgeons or in courses in human physiology, and the potential ability to predict the effects of human activity on Earth systems, such as the burning of fossil fuels, is driving interest in digital twins in the environmental and conservation sciences. Digital twins may be cheaper, safer, and faster than physical systems or scaled analog models, but their benefits must be assessed in the context of achievable accuracy and the cost of building the necessary digital representations.

A recent consensus study by the National Academies of Sciences, Engineering, and Medicine sees digital twins as potentially transformative in many scientific disciplines, and urges supporting actions by several US federal agencies (NASEM, 2023). It defines digital twins as:

> *"A set of virtual information constructs that mimics the structure, context, and behavior of a natural, engineered, or social system (or system-of-systems), is dynamically updated with data from its physical twin, has a predictive capability, and informs decisions that realize value. The bidirectional interaction between the virtual and the physical is central to the digital twin." (p.18)*

In contrast to the earlier definition from Grieves, this definition places emphasis on the two-way interactions between the twin and its physical counterpart, and includes a reference to its value in making decisions. These themes occur throughout the NASEM report. While the definition is mute on the question of accuracy, the report includes extensive discussion of what it terms "VVUQ," for verification, validation, and uncertainty quantification; these terms are explored in detail in the third section of this paper.

While the term "digital twin" dates from around 2000, there are many references a decade earlier to the concept of a digital system as a replica or mirror world. David Gelernter's, 1993 book *Mirror Worlds* (Gelernter, 1993) describes a virtual replica that would "put the universe in a shoebox;" Neal Stephenson's *Snow Crash* (Stephenson, 1992) features a virtual world termed the "metaverse;" and Al Gore's *Earth in the Balance* (Gore, 1992) describes a future of environmental science that features a massive and accessible source of all available environmental data, to be termed "Digital Earth." Gore's concept was elaborated in a 1998 speech which was written when Gore was Vice-President (http://www.zhanpingliu.org/research/terrainvis/digitalearth.pdf), and which became the conceptual foundation of a biennial International Symposium on Digital Earth and the *International Journal of Digital Earth*.

3 Accuracy

The term "twin" shares etymological roots with "two", but has specialized meaning as one of a pair of human or animal offspring that are delivered at the same birth. "Identical twins" occur when the same fertilized egg splits, and thus share the same DNA. However different environmental factors, both prenatal and postnatal, ensure that so-called identical twins are never exactly identical. The digital twin and the real system, by contrast, do not share a literal genetic connection; the former is a representation or imitation of the latter, based on computational rather than physical principles. So while the term "digital twin" certainly implies some level of relationship with the real system it attempts to replicate, it cannot imply the kind of perfection that is required by the Turing test or the Grieves definition quoted earlier. Alternative terms have been suggested that might avoid any misunderstanding resulting from the use of "twin", including "avatar" in the sense of an embodiment or representation of an idea, and "doppelganger" as a person's double.

If perfection in digital twins is impossible, then some means of determining the degree of imperfection of a particular digital twin would clearly be useful. Moreover a digital twin might pass the Turing test for some kinds of information but not others, so it would be helpful to know which kinds. From an application perspective, it would be useful to know the degree of uncertainty or error that is present in any predictions made from the digital twin. Thus many authors have called for greater attention to uncertainty. The NASEM report cited earlier (NASEM, 2023) argues that this "is an area of particular need that necessitates collaborative and interdisciplinary investment to advance the responsible development, implementation, monitoring, and sustainability of digital twins" (p.22). The report identifies this area as "VVUQ", for verification, validation, and uncertainty quantification. But although the terms verification and validation are often encountered in disciplines that make use of measured observations, they both imply a question with only two possible answers: is the analysis valid or verified, or not? Any measured observations inherit the uncertainties in the measuring instrument, and all models include some degree of inaccuracy, so the terms "valid" and "verified" clearly cannot imply perfection, but must instead depend on some specific context: the uncertainties or errors must be sufficiently small not to disrupt the objectives of a particular use. This argument thus leads to the doctrine of "fitness for use"; the digital twin is valid or verified against a specific use case. That implies a well-defined purpose for the digital twin; as noted earlier, purpose is not addressed in most of the commonly used definitions of digital twins. The ethical implications of this issue are examined further in the fourth major section of this paper.

The earliest geographic information technologies were focused on creating digital representations of the contents of maps, and doing so as accurately as possible (Goodchild, 2018). Yet as Korszybski pointed out many years ago (Korzybski, 1933), "the map is not the territory", in other words a map is no more than a representation of reality. It has long been accepted that all maps are generalizations, approximations, and abstractions of reality, and a map at a given cartographic scale will suppress all detail that is too fine to show at that scale. It was not until the 1980s that researchers began to investigate the accuracy of geographic information technologies by directly comparing the results of representation, analysis, and modeling with corresponding values drawn not from maps but from the physical world (Burrough, 1986; Goodchild & Gopal, 1989). Initially the terms "accuracy" and "error" were used. But while many types of geographic data originate as measurements and inherit measurement errors, other types such as maps of soils or land cover are not replicable, in the sense that two experts would not produce identical maps of the same area. Moreover many concepts that are commonly used in the geographic sciences are not rigorously defined. Thus "accuracy" and "error" have generally been replaced by "uncertainty" in the literature (Zhang & Goodchild, 2002).

Several fundamental problems arise when attempting to quantify uncertainty in geographic information technologies. First, how should one deal with the uncertainties that result because many geographic concepts are not rigorously defined or readily measurable? This is both an ontological and operational issue for digital twin construction. For example, what is a neighborhood in an urban digital twin? We commonly represent neighborhoods using polygons that follow administrative boundaries. Are their boundaries so clean? Is this even a useful representation for the predictive or explanatory purpose of the given twin? The uncertainty we acquire as we measure these effects is always rooted to how these units and relationships are defined.

Second, positive spatial dependence is almost always present in geographic information, a property expressed by Tobler (1970) as "nearby things are more similar than distant things", and a property that distinguishes the field of geostatistics or regionalized variables (Goovaerts, 1997) from classical statistics. The errors present at some location $\mathbf{x}$ will likely be strongly and positively correlated with the errors present a short distance away at $\mathbf{x}+\boldsymbol{\delta}$. This will greatly reduce the effective number of degrees of freedom in any data set, and likely violate the key independence assumption of inferential analysis.

Third, the quantification of uncertainty requires a comparison between the results of a query or analysis of a digital representation, and equivalent results from the physical world. Those equivalent results might come from a source of higher accuracy or lower uncertainty, or perhaps a source of finer spatial resolution, if certain assumptions can be made. But for many types of geographic information there can be no possibility of an equivalent result that is absolutely true. Moreover for a dynamic digital twin making predictions of the future, any effective quantification of uncertainty is clearly highly problematic, if not impossible.

Third, quantification of uncertainty in a complex model requires appropriate methods for the propagation of uncertainty, including that present in the input data and that inherent in the model. Some methods of propagation are available (Heuvelink, 1998), by making what are often draconian assumptions about the relationships between input data sets. For example, two input data sets may have been derived from the same source, and may therefore have inherited some of the same errors, while other data sets may have come from statistically independent sources. The effects of uncertainty are clearly much greater in the second case.

Fourth, visualization might be an attractive option for communicating uncertainty to users. However research has shown several problems with this approach. The presence of strong positive spatial dependence is difficult to visualize, though experiments with animation show promise; a waving blanket provides a useful metaphor because it similarly preserves strong positive spatial dependence. Efforts have also been made to incorporate uncertainty into choropleth maps (Mu & Tong, 2022). Unfortunately users of maps do not expect uncertainty, believing instead that the Earth's surface has been accurately mapped, at least since the end of the nineteenth century and with the growth of Earth-observing satellites. Thus any display of uncertainty, such as the use of blurring or greying, must be accompanied by explicit notes to the user to avoid misunderstanding.

In short, while it is important to recognize the existence of uncertainty in digital twins, and important that the user understands its significance and associated dangers in the form of misinterpretation and unjustified trust, the practical quantification of uncertainty is an extremely challenging problem. There is a substantial body of literature on the topic, but much of it is highly mathematical and thus likely inaccessible to many users.

4 Contextual issues

4.1 Purpose

The discussion to this point has implied the possibility of several distinct purposes for digital twins. First, a digital twin might be built for the purposes of prediction, by allowing its users to anticipate the effects of planning decisions. The temporal scale of such effects might be minute to minute, as when a digital twin is used for predicting traffic congestion in response to emergencies or major events, or it might be year to year, as when a digital twin is used to predict the long-term impacts of new housing or highway developments or the burning of fossil fuels. In both cases, and more importantly in the first, it will be necessary for the computation required by the digital twin to occur more rapidly than the predicted changes. If continuous input from the physical world is planned, then it will similarly be necessary for the digital twin to ingest and process that input much more rapidly than the rate of acquisition of new data. Moreover, since more accurate predictions from the digital twin will likely take more computation, it will be important to find an effective compromise between accuracy and computation time. Such considerations make digital twins expensive to construct and operate. Thus it is important that the benefits of the improved decision making supported by the digital twin exceed the costs to build and maintain it. Digital twins must yield an adequate return on investment.

A digital twin may also find value as a repository. As an illustrative example, Wright et al. (2021) argue that large conservation preserves commonly attract hundreds of researchers and their projects during field seasons, but have no means of preserving the data, models, and results of these projects in a unified, accessible repository. Thus valuable digital assets return with their researchers to home institutions, making it difficult for future projects to leverage them for further and possibly more integrated research, or for land managers to use them for monitoring and intervention. When any technical and ethical issues associated with these assets are resolved, a digital twin can serve to aggregate and share such data. While the preserve discussed by Wright et al. is focused on ecological and anthropological conservation, the same concept could be applied to cities, where a digital twin might combine the products of research with data acquired by the city and its contractors. Thus the ideas advanced by Wright et al. represent a distinct purpose for digital twins that has no obvious implications for prediction, but could have great value as a vision for digital twins in urban informatics. The technical issues raised by this vision are discussed in the final section.

Value for a digital twin may also lie in its potential role in data sharing, especially in urban informatics. A digital twin can be an effective way of realizing objectives of openness and transparency in cities, by making information readily accessible to city managers, citizens, community groups, and NGOs. In this context information should be readily understood, through a user interface that includes access to definitions of terms, visualizations and infographics, the provenance of data, the inner workings of models, and details of uncertainties. Developed in this way, digital twins may be a platform for the type of open and global urban science called for by Acuto et al. (2018). Given their capacity as platforms that can advance the truth claims of potentially competing interest groups, digital twins are also inevitably susceptible to politicization and discord.

4.2 Governance

The concept of a digital twin implies that all data about a given domain, whether it be a farm, preserve, city, nation, or the entire globe, be accessible through one or more portals. For a given city, however, it is to be expected that data resources are held by numerous owners, who may or may not be amenable to sharing through a portal. The city will possess some types of data, especially those related to city management and taxation, but other types will be in proprietary hands. Utility companies, for example, tend to hold data closely, and to be reluctant to share data even in emergencies. While strenuous efforts have been made to ensure that data can be shared during emergencies (see, for example, NRC, 2007), and many cities have organized elaborate portals to improve access to data, the results fall far short of the universal accessibility demanded by digital twins. Projects such as Open Street Map (https://www.openstreetmap.org/) have provided an effective response to this situation through the use of volunteers, for certain types of data that can be obtained from individual use of GPS tracking and fine-resolution satellite imagery, and cellphone data have opened access to vast resources of information about individual spatial behavior in some circumstances. But the sharing of data through digital twins remains a largely unsolved and intractable problem at this time.

Similar concerns exist about access to the models needed to enable the replication of processes and simulations in digital twins. Many models are proprietary or subject to intellectual property restrictions. Moreover many available models may address the same process, perhaps using different data sources and reflecting differing theories about the operation of human and physical processes. Notably, the primary fields of research that are focused on human systems are already divided on the fundamental assumptions necessary for modeling human interactions and social processes. The knowledge on which it might be possible to choose between alternative models—predictive accuracy, computing requirements, supporting software, peer review—are often not available or incomplete. Moreover, at this time there are few potential mechanisms for searching for models, and few catalogs except in narrowly defined areas (Hill et al., 2001).

In short, the apparently straightforward objective of a digital twin—to mimic "the structure, context, and behavior of a natural, engineered, or social system" (NASEM, 2023)—will be beset by numerous practical realities, many of which may appear intractable. Moreover for any given natural, engineered, or social system there will almost always be more than one possible digital twin, offering different levels of predictive accuracy, spatial resolution, cost, and computing requirements. The potential for multiple and overlapping digital twins also has a geographic dimension, since a digital twin for a city may find itself in competition with digital twins for specific neighborhoods, or for a containing county or state. This reinforces the possibility of digital twin predictions as an *ensemble*, a concept that is now widely known in the modeling of global climate change, where as many as 50 different models may offer varying predictions of the same climate future. An ensemble approach to prediction will still face considerable challenges in the prediction of uncommon and surprising events, or what are referred to in the literature as black swans. Furthermore, it also raises the somewhat worrying possibility of a search across overlapping digital twins for the one that best supports a hidden purpose. Cherry picking results from complex digital twins would represent a form of selective reporting or selective inference, a practice that is already well documented in existing literature. In contrast, developing reliable ensemble models requires the principled weighting of alternative models and the tracking and communicating of uncertainty across the set. How easily practices from climate modelling would translate to urban ensembles is an open question.

4.3 Ethics

It is of course desirable that the development and use of digital twins be ethical. This section opens the door to this topic by discussing some of the more immediately relevant issues, but no claim is made that the set of issues discussed here is complete. Suggestions are also made for technical developments that might reduce the likelihood of unethical behavior. Since many of these issues are also found with many other forms of geospatial technology, recent efforts by the American Association of Geographers (https://www.aag.org/program/locational-information-and-the-public-interest/) and the American Geographical Society (www.ethicalgeo.org) might be helpful.

The topic of uncertainty was discussed at some length in Section 3, and the point was made that any predictions from a digital twin must necessarily be subject to some level of uncertainty. To ignore uncertainty, and to present a prediction as perfect, is therefore unethical; yet the very use of the term "twin" may oversell the technology and lead to unreasonable expectations. More problematic, perhaps, is the difficulty of actually measuring uncertainty with any degree of confidence, given the many contributions to uncertainty from the data and simulation models, and the positive spatial dependence that will be present. The doctrine of fitness for use is helpful in this context, but it assumes an ability to determine both the uncertainty present in predictions, and the level of uncertainty that would be acceptable to a given decision. It also implies that repurposing of a digital twin, without an associated determination of fitness for the new use, is similarly unethical.

Relatedly, identification of fitness for use necessarily raises the question of who determines the purpose, use, and fitness of a digital twin. Resolving this issue is at best difficult as many groups likely have a claim to influence these decisions if digital twins of urban systems will affect their lives. The analogous issue of algorithmic governance and emerging frameworks on digital privacy and location tracking may be a useful guide for this issue. However, the diversity of people living in cities and their varied power and capacity to engage with agenda setting will likely remain a barrier to ethical inclusion.

Applications of digital twins in urban informatics may also raise ethical issues when they deal with people at the individual level, as for example in the simulation of individual spatial behavior. Data about individuals will raise questions of individual privacy and concerns for surveillance. Aggregation is a traditional solution to the privacy problem, but will result in the loss of spatial and temporal detail. Alternatively various approaches have been suggested for masking spatial and temporal detail, replacing accurate locations and times with randomly generated ones.

When a digital twin is constructed by a highly qualified team, it seems inevitable that the result will lack full transparency. Many of the data sets may lack adequate information on provenance, and may have been pre-processed to achieve compatibility—for example by resampling—without full documentation. The models used as digital replicas of processes may not be fully described and subject to peer review. And if neural nets have been used in the modeling, as is increasingly common given the popularity of AI, then it is likely that many aspects of its solutions have not been exposed. This raises concerns about replicability, specifically whether the results generated by one digital twin can be replicated with another (Goodchild & Li, 2021). Furthermore, cutting-edge AI foundation models, such as OpenAI's Sora, have demonstrated strong capabilities in creating virtual realities based on text descriptions. Applying similar techniques in a digital twin context could lead to the generation of a faked replica of physical reality. Without proper provenance of data and models, the trustworthiness of a digital twin and its predictions could be negatively affected (Anselin et al., 2014). In summary, transparency is a strong dimension of trust, and essential if the digital twin is to be subjected to intensive peer review.

Some of these ethical issues might be approached with technical solutions that operate in the user interface. For example, a user might be prevented from querying a digital twin containing individual data, but allowed to see aggregated maps or statistical summaries. The use of maps to visualize uncertain data or predictions might be avoided, because of the danger that maps will be misinterpreted, and instead results might be made available in the form of statistical data with associated uncertainty warnings.

A final high-level challenge in the development of digital twins will be the resolution of contradictions between different forms of normative ethics when these technologies are applied to urban systems. As in all areas of application, it is often unclear how an individual should act when his or her choices create outcomes that would be valued differently under different ethical systems. For example, a digital twin may well face instances of the Trolley Problem (Foot, 1978), which has different resolutions when taking consequentialist and certain duty-based ethical approaches.

5 Concluding points and ways forward

As noted at the beginning of Section 2, researchers have long been attracted by the building of replicas as a means of prediction and as a way of increasing our understanding of natural and human-influenced processes. Initially such models were scaled physical analogs, but the development of digital technologies has given modeling much greater momentum, through the availability of massive amounts of fine-resolution data and of software that is capable of emulating processes in the physical and social worlds. But if the replication of real processes has such a long history, what might explain the recent adoption of the term and the recent enthusiasm around the concept? Has some kind of threshold been passed, in the availability of data, models, and high-performance computing?

In urban informatics it is now common to capture fine-resolution representations of cities, often using airborne and ground-based LiDAR. Such techniques clearly have value in planning, since they can be edited and modified at will in order to create simulated images of physical appearance. A fine-resolution image might well be said to pass a limited form of the Turing test, since it would be next to impossible to determine if the image was generated from a digital database, or captured in a photograph. Perhaps this explains the sudden interest in digital twins in urban informatics, since it comes close to Grieves' definition in the narrow sense of visual inspection: "any information that could be obtained from inspecting (the real world) can be obtained from its Digital Twin" (Grieves, 2016). But it is essentially static, technically straightforward, and falls far short of including and modeling the behavior of the city's occupants (a far more complex and difficult problem; Fotheringham, 2023), or its vehicles, infrastructure, and long-term development.

Moreover some of the more recent compute-intensive developments in planning—photorealistic visualizations of species-specific trees, vehicles, humans, and even graffiti—have often been attacked as overkill and no more helpful than the traditional "artist's impression." Why, one might ask, should a simple binary decision about a plan require terabytes of data and gigaflops of computation in order to reach the conclusion that some like the plan and some don't?

It may be worth noting that conceptions of digital systems as replicas emphasize digital objects as reflections of the real world that are perhaps useful as mirrors for describing or predicting events. However, following Abrams's (1954) literary critique, it is perhaps more constructive to imagine digital twins as "lamps": purpose-built objects that can be used to direct human creativity and ingenuity to the explanation of and intervention in real systems. This shift in perspective is useful in a number of ways. First, it moves emphasis from representational accuracy as a goal to representational accuracy as a necessary means of measuring and controlling uncertainty, to allow for reliable action. Linking VVUQ measures to these goals can likewise act as a clarifying device that might indicate which types VVUQ are most important to pursue for a given twin. Second, understanding digital twins as lamps makes plain that their construction and use is inextricably tied to ethical questions related to specific actions.

If digital twins represent a desirable goal, then it is appropriate for the research community to ask: What do we not yet know how to do, and what is inhibiting progress toward that goal?

- *Building the database.* Simulation models have their own requirements for data structures, yet we have made little progress in developing methods of data integration and fusion. Data sets will likely need to be resampled, and perhaps downscaled, raising additional issues concerning uncertainty.
- *Search for process models.* We need appropriate methods for describing models of process (a metadata for models), for documenting their provenance, for finding and evaluating models, and for determin-

ing the inherent uncertainty associated with each model.

- *Interoperability of process models*. We need more research on how models can work together effectively, simulating different processes in a digital twin. We need to know more about the integration of different spatial and temporal resolutions in dynamic models.
- *Description of digital twins.* We need research on standards for the description of digital twins (e.g., geographic coverage, the uses for which a given digital twin is fit, approaches to the estimation of uncertainty, spatial and temporal resolution).
- *The user experience*. We need research and examples of the kinds of user interfaces that are suitable for given types of users, given purposes, and protection of confidential information. Is it possible for the user interface to anticipate and flag potentially unethical behavior on the part of the user?

Many questions remain about the use of the term digital twin in urban informatics. What gives someone the right to claim that a certain combination of data, models, and user interface constitutes a digital twin? What exactly is implied by the term, and is its use ethical? Hopefully this paper has shed some light on some of these questions. In essence, however, the use of the term is a matter of human behavior and the sociology of science: it can mislead, oversell, and create unreasonable expectations as a direct result of its transfer from a practical idea in manufacturing to an elusive goal in the social and environmental sciences.

Acknowledgements
Not applicable.

Authors' contributions
Prof. Goodchild drafted the paper. All listed authors were equally involved in the design, editing, and review of the paper.

Funding
Funding for the workshop that led to this paper was provided by Esri and by Arizona State University.

Availability of data and materials
No data were used in this research.

Declarations

Ethics approval and consent to participate
Not applicable.

Consent for publication
Not applicable.

Competing interests
Prof. Goodchild is an Editorial Board member of *Urban Informatics*. He was not involved in the peer review or handling of the manuscript. The authors have no other competing interests to disclose.

Received: 29 March 2024 Revised: 25 April 2024 Accepted: 25 April 2024
Published online: 16 May 2024

References

Abrams, M. H. (1954). *The Mirror and the lamp: Romantic theory and the critical tradition*. New York: Oxford University Press.
Acuto, M., Parnell, S., & Seto, K. C. (2018). Building a global urban science. *Nature Sustainability, 1*, 2–4. https://doi.org/10.1038/s41893-017-0013-9
Anselin, L., Rey, S. J., & Li, W. (2014). Metadata and provenance for spatial analysis: the case of spatial weights. *International Journal of Geographical Information Science, 28*(11), 2261–2280.
Batty, M. (2018). Digital twins. *Environment and planning B: Urban analytics and City Science*, *45*(5). https://doi.org/10.1177/2399808318796416.
Box, G. E. P. (1976). Science and statistics. *Journal of the American Statistical Association, 71*(356), 791–799.
Burrough, P. A. (1986). *Principles of geographical information Systems for Land Resources Assessment*. Oxford: Clarendon Press.
Cureton, P., & Dunn, N. (2021). Digital twins of cities and evasive features. In A. Aurigi & N. Odendaal (Eds.), *Shaping smart for better cities* (pp. 276–282). Academic Press. https://doi.org/10.1016/B978-0-12-818636-7.00017-2
Dawkins O, Dennett A, Hudson-Smith A (2018) Living with a digital twin: Operational management and engagement using IoT and Mixed Realities at UCL's Here East Campus on the Queen Elizabeth Olympic Park. Available at: www.researchgate.net/publication/324702983_Living_with_a_Digital_Twin_Operational_management_and_engagement_using_IoT_and_Mixed_Realities_at_UCL%27s_Here_East_Campus_on_the_Queen_Elizabeth_Olympic_Park . (Accessed 6 December 2023)
Foot, P. (1978). *Virtues and vices and other essays in moral philosophy*. Oxford: Basil Blackwell.
Fotheringham, A. S. (2023). Digital twins: The current "Krays" of urban analytics? *Environment and Planning B: Urban Analytics and City Science, 50*(4), 1020–1022. https://doi.org/10.1177/23998083231169159
Gelernter, D. (1993). *Mirror worlds*. New York: Oxford University Press.
Goodchild, M. F. (2018). Reimagining the history of GIS. *Annals of GIS, 24*(1), 1–8. https://doi.org/10.1080/19475683.2018.1424737
Goodchild, M. F., & Gopal, S. (1989). *Accuracy of spatial databases*. Basingstoke UK: Taylor and Francis.
Goodchild, M. F., & Li, W. W. (2021). Replication across space and time must be weak in the social and environmental sciences. *Proceedings of the National Academy of Sciences, 118*(35), e2015759118. https://doi.org/10.1073/pnas.2015759118
Goovaerts, P. (1997). *Geostatistics for natural resources evaluation*. New York: Oxford University Press.
Gore, A. (1992). *Earth in the balance: Ecology and the human Spirit*. Boston: Houghton Mifflin.
Grieves, M. (2016). *Origins of the Digital Twin Concept*. https://doi.org/10.13140/RG.2.2.26367.61609 Available at www.researchgate.net/publication/307509727_Origins_of_the_Digital_Twin_Concept (Accessed 19 December 2023) .
Heuvelink, G. B. M. (1998). *Error propagation in environmental modelling with GIS*. Basingstoke, UK: Taylor and Francis.
Hill, L. L., Crosier, S. J., Smith, T. R., & Goodchild, M. F. (2001). A content standard for computational models. *D-Lib Magazine*, *7*(6).
Korzybski, A. (1933). *Science and sanity. An introduction to non-Aristotelian systems and general semantics* (pp. 747–761). The International Non-Aristotelian Library Pub. Co.
Longley, P. A., Goodchild, M. F., Maguire, D. J., & Rhind, D. W. (2015). *Geographic information science and systems* (4th ed.). Hoboken NJ: Wiley.
Mu, W., & Tong, D. (2022). A general model for creating robust choropleth maps. *Computers, Environment and Urban Systems, 96*, 101850. https://doi.org/10.1016/j.compenvurbsys.2022.101850
National Academies of Sciences, Engineering, and Medicine. (2023). *Foundational research gaps and future directions for digital twins*. Washington DC: National Academies Press.

National Research Council. (2007). *Successful response starts with a map: Improving geospatial support for disaster management*. Washington, DC: National Academies Press.
Stephenson, N. (1992). *Snow crash*. Bantam Books.
Tobler, W. R. (1970). A computer movie simulating urban growth in the Detroit region. *Economic Geography, 46*, 234–240. https://doi.org/10.2307/143141
Turing, A. (1950). Computing machinery and intelligence. *Mind, 59*(236), 433–460. https://doi.org/10.1093/mind/LIX.236.433
Wildfire, C. (2018). *How can we spearhead city-scale digital twins? Infrastructure Intelligence*. Available at: www.infrastructure-intelligence.com/article/may-2018/how-can-we-spearhead-city-scale-digital-twins (Accessed 6 December 2023) .
Wright, D. J., Easterday, K., Reynolds, M., Merrifield, M., Bell, M., & Goering, M. (2021). *Building a digital twin of the Dangermond preserve, the last coastal wilderness in Southern California*. New Orleans: AGU Fall Meeting. Available at: https://ui.adsabs.harvard.edu/abs/2021AGUFMSY14A.02W/abstract (Accessed 22 December 2023) .
Zhang, J. X., & Goodchild, M. F. (2002). *Uncertainty in geographical information*. New York: Taylor and Francis.

Publisher's Note