Article:
Birkin, M., Ballantyne, P., Bullock, S. et al. (2025) Digital twins and AI for healthy and sustainable cities. Computers, Environment and Urban Systems, 120. 102305. ISSN: 0198-9715

https://doi.org/10.1016/j.compenvurbsys.2025.102305

Title page

Digital twins and AI for healthy and sustainable cities

Mark Birkin[1,2], Patrick Ballantyne[3], Seth Bullock[4], Alison Heppenstall[2,5], Heeseo Kwon[3], Nick Malleson[1,2], Jing Yao[5], Anna Zanchetta[2]

(1) University of Leeds; (2) Alan Turing Institute; (3) University of Liverpool; (4) University of Bristol; (5) University of Glasgow

Manuscript

This paper offers recent thinking from the UK national programme in urban analytics (UA), hosted by the Alan Turing Institute (the Turing) as the national institute for data science and artificial intelligence (AI). A key objective of the programme, in line with the broader ambitions of the Turing, is to promote cutting edge research (drawing on the latest innovations in AI and machine learning (ML)), which generates impact through direct applications to policy and society. It is hard to overstate the contemporary importance of AI, given the breadth of social, economic and environmental challenges or 'wicked problems' we continue to try and resolve. For instance, the past 5 years have seen a global pandemic, a growing climate emergency and wars in Europe and beyond – with economic impacts such as increasing prices and a loss of energy and food resilience within the UK and much of Western Europe, and wider humanitarian repercussions across the globe. These are just a few of the 'wicked problems' that demand more robust and sophisticated evidence and modelling which can generate better insights into the design of effective mitigation strategies. Digital twins[1] and AI[2] can help meet this demand by offering decision-makers the possibility to interrogate a range of potential futures that can be expressed in social, financial or physical outcomes (e.g. poverty, healthy life expectancy, living standards, net zero; UDG11).

At a recent symposium drawing together members of the Turing's Urban Analytics Network, we assembled a panel session that included voices drawn from different disciplines and career stages. We wanted to leverage their combined expertise to evaluate the current and future potential of both AI and digital twin technologies to address wicked challenges, framing the conversation around both the 'what' (how do we understand these methods) and the 'so what' (why do we need them?) questions. These questions have a topical importance as we sit on the cusp of an AI revolution with more data than ever, and driven forward by a 'current craze for Digital Twins'[i] .

A persistent criticism of methodological innovation is that it is no more than 'old wine in new bottles'. In relation to AI, this could manifest itself in the view that AI is just the latest box of tricks which facilitates some degree of prediction of a complex system without delivering an attendant degree of understanding. Against this we argue that the achievements of AI over a long period are demonstrable and profound – for example, 50 years ago the perceived wisdom was that no computer programme could play chess to the standard of a human master, or replicate the conversational interaction between human companions. Now such things are established and widely exploited. Whilst large language models (LLMs) may be something less than a panacea, they are rapidly moving towards a demonstrable utility to support a huge range of tasks, from text summarisation and evaluation to data analysis and programming. As they become more refined, LLMs and related tools may become the de facto interface to data, offering the ability to gain insight from data that has previously only been possible using advanced skills and technologies.

One thing which has until recently been lacking here, perhaps, is a recognition of the challenges in extending such technologies beyond text and into truly spatial datasets (e.g. area-based counts and profiles, networks and flows, positional data, satellite images, etc.). Previous developments in GeoAI/spatial machine learning can here be grouped into two categories: applications of ML to

---

[1] A digital twin is a simulation of a real-world complex system, built for the purpose of predicting its behaviour in a range of counterfactual scenarios. Digital twins aspire to achieve a degree of empirical and ecological validity sufficient to make their predictions valuable to real-world system stakeholders.

[2] Artificial Intelligence refers to a wide array of technologies ranging from machine learning analytics, tools for automated reasoning and inference, chatbots and artificial agents, to so far unrealised artificial systems with human or super-human levels of general intelligence. In this paper we use "AI" in its broadest sense, encompassing potentially any of these technologies.

spatial data (most common in image recognition) and new ML approaches developed for spatial data (e.g., accounting for spatial heterogeneity or spatial autocorrelation). Recent reviews can be found in[ii,iii,iv]. The challenges associated with these approaches include how to include neighbourhood relationships in AI/ML models, scalability; transparency and explainability (this is important if DTs and AI as tools are to be used by policymakers), algorithm/data bias, how to link different urban subsystems (transportation, housing, energy, etc.) in DTs, and how to incorporate qualitative data like social values.

We are particularly drawing on concerns about how well LLMs and other AI technologies can become spatially literate, and the impacts of integrating such technologies into workflows which are dependent on a well-developed awareness of geographical space and theory. However, this may change as foundation models (a model that is trained on large data sets and can be applied to multiple case-studies), that have shown remarkable success in natural language processing and computer vision, continue to be developed for spatial contexts. Spatial foundation models that offer the ability to seamlessly train on diverse datasets (satellite images, spatially referenced text reports, sensor data, and so on; see[v] for an example) could potentially offer a more holistic understanding of spatial phenomena and the capability to detect more complex patterns than existing methods are able to. They also offer opportunities for transfer learning. Here huge models trained on vast databases using specialist hardware can be made available to others for application without the need for retraining. They can also be trained in data-rich areas and applied to those that are data poor. That said, foundation models are temporally static and are entirely data-driven; there are currently no mechanisms for including theories of human behaviour, decision-making or other system features that will be essential components for a robust urban analytics programme. Hence there may be a growing role for methods that can include 'soft' social features, such as travel demand or housing market aspirations that might be captured more effectively by urban digital twins.

The idea of a 'digital twin' (DT) has captured the imagination of policy-makers and funders, particularly within the UK, in recent years. At a simple level, the digital twin can be understood as an attempt to replicate a real-world system in silico, with the ambition that the dynamics or cross-sectional impact of real-world change is then subject to cheap, powerful and flexible simulation within the computational laboratory. One of the key advantages of DTs over conventional models or representations of (urban) systems is that they are highly dynamic[vi], where a change in the physical or 'real' system is accurately reflected in real-time by the DT, and vice versa. This is something that Kitchin and Dawkins (p6) described: "they become dyadically intertwined, with a change in one direction directly affecting the other"[vii]. Furthermore, there are other key distinguishing features of DTs, including their ability to be multidimensional or cross-sectoral, and their ability to operate in a fully autonomous fashion[viii,ix]. These ideas are considered as having the potential to be highly transformative[x], and has already attracted considerable interest in the modelling of complex physical entities, notably in the fields of engineering and medicine, with applications to human organs, bridges and aeroplane engines[xi] .

A similar level of appeal for urban systems is easy to understand, for example as a means to forecast and compare the impact of alternative policy interventions ranging from the global (e.g. levying a sugar tax) to the local (e.g. building a new road junction). In some cases, the parallel to engineered systems is a strong one. For something like a network of buses in a city which follows well-defined patterns, we argue that the feasibility of a DT is relatively obvious. Through the utilisation of large volumes of historical and real-time network data, we can simulate a number of 'what if' scenarios in an effort to make substantial network improvements, evaluating these against known mobility requirements of the population[xii]. Operational examples can be found in cities around the world including Singapore[xiii], whereby DTs are used to make key predictions, and inform better urban planning in these locations. Whilst the practical obstacles to successful deployment of such urban DTs remain considerable (e.g. in relation to data sharing, transparency), perhaps the greatest challenge relates to the digital representation of complex human behaviours.

Human behaviour is 'messy'[xiv] and does not follow a distinct series of rules that can be easily modelled and digitally represented. Nevertheless we do have at our disposal established approaches to the representation of human behaviour in urban DTs, such as spatial agent-based modelling[xv] and the use of synthetic populations[xvi]. Analytical approaches to the COVID-19 pandemic is a compelling instance in which DTs of a policy environment – specifically through nonpharmaceutical interventions (i.e. lockdowns) – can be seen as both necessary and achievable in advancing the art of spatial modelling[xvii]. Elsewhere, the concept of the 'social DT' has been advanced as a means for the synthesis of population or household data into urban DTs[xviii]. Such efforts are in line with Goodchild et al.'s emphasis on the importance of scale in urban DTs, "where processes range from the individual scale of observable human behaviour to the emergent properties that characterize entire cities and societies"[xix].

Of course the parallels between a scientific and a social system in this debate are far from trivial. The power of contemporary computational hardware is sufficient to make the faithful and complete representation of a physical entity (such as an aircraft wing or a combustion chamber) a realistic ambition. However, the richness and complexity of human behaviour would suggest to most that a similar level of fidelity in social systems models will remain a subject for science fiction[xx] rather than political fact into the long-term future. This is far from a novel observation, with echoes back to the early days of the quantitative revolution in geography and social science[xxi] . This raises questions about the language and construction of a digital twin which are far from semantic. Whilst the dictionary definition of a twin as 'one of two persons or things closely related to or resembling one another' (Webster) admits a degree of variation between template and image, the underlying intentionality is clear. The language of 'twin' may also mislead the ultimate users of the technology – policymakers in particular – into perceiving what is ultimately a probabilistic and uncertain simulation as a 'promise' regarding the true state of the world or its future state. The notion of 'digital cousins for robust policy' (Dai et al, 2024, https://arxiv.org/abs/2410.07408) goes some way towards addressing this deficiency, although perhaps not far enough to satisfy the hardliners for whom the digital twin is no more than the latest costume in which to dress established underpinning technologies (or just 'models' in the accepted language of social simulation)[xxii].

For all of this, there are also compelling arguments for sticking with the digital twin framing as a vehicle for policy innovation: first, that the DT has a proven capability to capture the imagination of government and policy-makers themselves, perhaps even the general public; second, that as with AI the rate of transformation in the associated methods – and particularly data – is clearly rapid, substantial and ongoing; third, to the extent that the DT metaphor is driving advances in physical modelling of fully replicable systems, then continuing to co-develop alongside such advances may be rational in its own right.

A further question which arises here is the legitimacy of the digital twin vision in the first place. At one extreme we might ask whether the optimisation of policy insight through an all-seeing robotic overseer serves the interest of a free society, while if the answer to this is in the negative then what is our motivation in further development of such technology? The answer to this conundrum is surely that while policy will always be seen as a necessary interplay between the achievable and the desirable in a well-governed society, the means by which agreed outcomes can be achieved (robustly, efficiently and equitably) is still uncertain and challenging, for all of the reasons cited above in relation to systemic complexity. Policymakers need support in this process, so that the rationale is simply one of enabling better decisions through scientific augmentation[xxiii]. Whilst AI and DT approaches are also potentially exposed to criticism for their 'black box' approach there is also a counter-argument that the power of technology may also be used to aid the interpretability of decisions and the underlying evidence. Might it be possible, for example, as in certain recent experiments to leverage LLM-style functionality onto vast repositories of social simulations to determine and narrate the outcome of alternative policy choices which fuels a constructive debate regarding those choices and their consequences?[xxiv] Could the right kind of AI enable a productive policy-relevant interplay between detail-rich "twins" with their realism and predictive validity (but opaque impenetrable complexity and fragility) and their semi-automatically generated unsophisticated toy-model "cousins" with their robustness and explanatory transparency but lack of realism[xxv]. Regardless of their specific practicability, such possibilities give a sense of the potential benefits from these research avenues which is to some extent independent of immediate short-term policy concerns.

In conclusion, we urge early adopters to keep pioneering DTs and related AI technologies to foster significant leaps in the understanding and management of wicked problems in cities, advocate for continued critical examination of these methodologies from technical, practical and political perspectives, and encourage the undecided to commit more fully to engaging with the full spectrum of these efforts. The politics of social decision-making is complex, contested and multifaceted, but obscuring the relation between cause and effect or choice and outcome does nothing to improve the state of our world. As social scientists this is surely our ultimate motivation, and one which can only benefit from greater critical engagement with the new and powerful emerging technologies.

---

[i] Fotheringham, A.S., 2023. Digital twins: The current "Krays" of urban analytics?. Environment and Planning B: Urban Analytics and City Science, 50(4), pp.1020-1022.

[ii] Georganos, S., Grippa, T., Niang Gadiaga, A., Linard, C., Lennert, M., Vanhuysse, S., ... & Kalogirou, S. (2021). Geographical random forests: a spatial extension of the random forest algorithm to address spatial heterogeneity in remote sensing and population modelling. *Geocarto International*, *36*(2), 121-136.

[iii] Sun, K., Zhou, R. Z., Kim, J., & Hu, Y. (2024). PyGRF: An improved Python Geographical Random Forest model and case studies in public health and natural disasters. *Transactions in GIS*, *28*(7), 2476-2491.

[iv] Nikparvar, B., & Thill, J. C. (2021). Machine learning of spatial data. *ISPRS International Journal of Geo-Information*, *10*(9), 600.

[v] Balsebre, P., Huang, W., Cong, G. and Li, Y. (2024). City Foundation Models for Learning General Purpose Representations from OpenStreetMap. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management* (CIKM '24), October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3627673.3679662

[vi] Wright, L. and Davidson, S., 2020. How to tell the difference between a model and a digital twin. *Advanced Modeling and Simulation in Engineering Sciences*, *7*, pp.1-13.

[vii] Kitchin, R. and Dawkins, O., 2025. Digital twins and deep maps. *Transactions of the Institute of British Geographers*, *50*(1), p.e12699.

[viii] Caprotti, F., Cugurullo, F., Cook, M., Karvonen, A., Marvin, S., McGuirk, P., & Valdez, A.- M. (2024). Why does urban Artificial Intelligence (AI) matter for urban studies? Developing research directions in urban AI research. Urban Geography, 45(5), 883–894.

[ix] Lartey, D., & Law, K. M. (2025). Artificial intelligence adoption in urban planning governance: A systematic review of advancements in decision-making, and policy making. *Landscape and Urban Planning*, *258*, 105337.

[x] Goodchild, M.F., Connor, D., Fotheringham, A.S., Frazier, A., Kedron, P., Li, W. and Tong, D., 2024. Digital twins in urban informatics. Urban Informatics, 3(1), pp.1-9.

[xi] Niederer, S.A., Sacks, M.S., Girolami, M. and Willcox, K., 2021. Scaling digital twins from the artisanal to the industrial. Nature Computational Science, 1(5), pp.313-320.

[xii] Son, H., Jang, J., Park, J., Balog, A., Ballantyne, P., Kwon, H.R., Singleton, A. and Hwang, J., 2025. Leveraging Advanced Technologies for (Smart) Transportation Planning: A Systematic Review. *Sustainability*, *17*(5), p.2245.

[xiii] https://geospatialworld.net/prime/case-study/national-mapping/virtual-singapore-building-a-3d-empowered-smart-nation/.

[xiv] Fotheringham, A.S., 2023. op cit.

[xv] Manson, S., An, L., Clarke, K.C., Heppenstall, A., Koch, J., Krzyzanowski, B., Morgan, F., O'Sullivan, D., Runck, B.C., Shook, E. and Tesfatsion, L., 2020. Methodological issues of spatial agent-based models. Journal of Artificial Societies and Social Simulation, 23(1).

[xvi] Wu, G., Heppenstall, A., Meier, P., Purshouse, R. and Lomax, N., 2022. A synthetic population dataset for estimating small area health and socio-economic outcomes in Great Britain. Scientific Data, 9(1), p.19.

[xvii] Spooner, F., Abrams, J.F., Morrissey, K., Shaddick, G., Batty, M., Milton, R., Dennett, A., Lomax, N., Malleson, N., Nelissen, N. and Coleman, A., 2021. A dynamic microsimulation model for epidemics. Social Science & Medicine, 291, p.114461.

[xviii] Yossef Ravid, B. and Aharon-Gutman, M., 2023. The social digital twin: the social turn in the field of smart cities. Environment and Planning B: Urban Analytics and City Science, 50(6), pp.1455-1470.

[xix] Goodchild et al, op cit:2

[xx] Eg Asimov, I., 1951. Foundation, Gnome Press, New York City.

[xxi] Harvey, D., 1973. Social justice and the City, Johns Hopkins University Press; Sayer, R.A., 1979. Understanding urban models versus understanding cities. Environment and Planning A, 11(8), pp.853-862.

[xxii] Fotheringham, A.S., 2023, op cit.

[xxiii] Castelvecchi, D., 2016. Can we open the black box of AI?. Nature News, 538(7623), 20.

[xxiv] Xu et al., 2024. Leveraging generative AI for urban digital twins: a scoping review on the autonomous generation of urban data, scenarios, designs, and 3D city models for smart city advancement, https://link.springer.com/article/10.1007/s44212-024-00060-w; Güzay, Özdemir and Kara (2023) A Generative AI-driven Application: Use of Large Language Models for Traffic Scenario Generation; https://ieeexplore.ieee.org/document/10415934/authors

[xxv] Bullock, 2014, Levins and The Lure of Artificial Worlds, The Monist, Volume 97, Issue 3, 1 July 2014, Pages 301–320, https://doi.org/10.5840/monist201497320.