

COMMENTARY

On the scaling of digital twins by aggregation

Marcelo Pias¹ , Lawrence Bull^{2,3}, Daniel S. Brennan⁴, Mark Girolami^{2,3} and Jon Crowcroft² 

¹Federal University of Rio Grande (FURG), Rio Grande-RS, Brazil

²University of Cambridge, Cambridge, UK

³The Alan Turing Institute, London, UK

⁴The University of Sheffield, Sheffield, UK

Corresponding author: Jon Crowcroft; Email: jon.crowcroft@cl.cam.ac.uk

Received: 29 March 2024; **Revised:** 15 August 2024; **Accepted:** 06 October 2024

Keywords: aggregation; data-centric AI; decentralised; digital twins; federation; modular representation; peer-to-peer; systems thinking

Abstract

We discuss the emerging technology of digital twins (DTs) and the expected demands as they scale to represent increasingly complex, interconnected systems. Several examples are presented to illustrate core use cases, highlighting a progression to represent both natural and engineered systems. The forthcoming challenges are discussed around a *hierarchy of scales*, which recognises systems of increasing aggregation. Broad implications are discussed, encompassing sensing, modelling, and deployment, alongside ethical and privacy concerns. Importantly, we endorse a modular and peer-to-peer view for aggregate (interconnected) DTs. This mindset emphasises that DT complexity emerges from the *framework* of connections (Wagg et al. [2024, The philosophical foundations of digital twinning, *Preprint*]) as well as the (interpretable) units that constitute the whole.

Policy Significance Statement

Digital twins (DTs) are an emerging data-centric technology that creates virtual replicas of real-world objects, systems, and individuals. DTs are usually a hybrid of physics-based, process-based, and machine learning models validated by data. As the use of DTs becomes more widespread, it is crucial to establish a new data governance framework to ensure that ethical design issues are addressed. This article advocates for a balanced approach through a mindset of aggregation to scale DTs.

1. A panoply of digital twins

Digital twins (DTs) and their terminology, tools, and techniques have been gathering pace. We collect some recognised definitions to comment on how the concept has evolved through time. Our goal is to emphasise, rather than redefine, emerging themes and summarise them into common traits ([Section 2](#)) throughout a hierarchy ([Section 3](#)) to discuss a panoply of core technical ideas.

(2002) *DTs Introduced*: Michael Grieves proposed the DT conceptually, under a different name: a *Conceptual Ideal for Product Life Cycle Management* (Grieves and Vickers, 2017). While general, the concept refers to the three elements that underpin most DT definitions: (i) the real space, (ii) the virtual space, and (iii) the links for information flow between the two. Importantly, the virtual representation is

not static; it evolves throughout the system's life cycle. Considering assets or products, *life* refers to production (design and manufacture), operation (sustain/support), and disposal. Many Industry 4.0 applications fall under this category of DTs (Kappel et al., 2022).

(2012) *Another Core Definition*: NASA's definition (Glaessgen and Stargel, 2012) is more specific: a DT is an integrated multi-physics, multi-scale, probabilistic simulation of an engineered system—to mirror the life of its physical twin. Critically, the terminology considers interdependent sub-systems, including propulsion/energy storage, avionics, life support, vehicle structure, and thermal management (Glaessgen and Stargel, 2012; Negri et al., 2017). While still focussing on manufactured components, there is consideration of interactions between aggregated systems and virtual sub-spaces.

(2024) *Virtual constructs (NASEM, 2024)*: NASEM's recent definition of a DT focuses on a “set of virtual information constructs that mimic the structure, context, and behaviour of a natural, engineered, or social system (or system-of-systems). It is dynamically updated with data from its physical twin, has predictive capability, and informs decisions that realize value. The bidirectional interaction between the virtual and the physical is central to the digital twin.”

(2021) *Royal Society (Royal Society, 2020)*: A DT is a *virtual representation of a physical asset* that can be used to understand, predict, and optimise performance. Simulations can be run before manufacture, or during use, with the possibility of (near) real-time data feedback. Feedback enables a *control loop*, with the possibility to adjust the real-world set-up based on insights from *simulation*. In turn, *simulation* capabilities become a high priority, throughout the system's life cycle. DT solutions are proposed for industries such as agriculture, highlighting a growing interest in DTs that involve both natural and engineered systems. It is worth noting that DT applications in agriculture remain relatively unexplored (Jans-Singh et al., 2020; Pylaniadis et al., 2021).

(2021) *Towards Planetary Scales (Nativi et al., 2021)*: Emerging from the fields of Meteorology and Geoscience (Bauer et al., 2021a,b) the term *Earth DTs* (or *DTs of the Earth*) has been proposed within space agencies, climate research, oceanography, and others (Ye et al., 2023). Generally speaking, these are large representations of an Earth-scale process or phenomenon. Such large-scale DTs are obtained by merging continually monitored data from a variety of sources—for example, remote sensing, in-situ measurements, and synthetic data streams. Researchers tend to focus on data assimilation and ensemble forecasting techniques, which are prevalent in the associated fields. Continual updates in Earth DTs should help to reflect potentially dramatic changes associated with the highly nonlinear systems, which approximate natural phenomena. At the planetary scale, many representations are dominated by (rather than augmented with) naturally occurring systems.

Our main take is on the increase in scale of DT representations. This progression is also recognised in the Alan Turing Institute's white paper (Bennett et al., 2023), comparing a range of relatively simple representations (a single structure or asset) to complex interacting systems (transportation networks). Another fundamental concept is *simulations* via the virtual representation, allowing us to probe the design space (risk-free) for improved decision-making and insights: the ultimate purpose of a DT.

Both *process* and *data-centric* procedures must be utilised to enhance interpretation (Chantry et al., 2021). Since the DT themselves *reflect* systems and processes, we can only progress with deeper process understanding and the associated interactions. This is core to science in general. Intuitively, process models will be central when scaling DTs, as they inform the *framework* that connects DT components, which must also evolve alongside models and data. Following Blair (2021), we consider process modelling and data modelling as two points on a spectrum. An underlying consideration of DT design is to identify a sweet spot, whereby process and data understanding work together.

2. Three pillars

We present an alternative view of DT concepts based on three keystone features: Emulation, Evolution, and Aggregation (see Figure 1). Whilst seemingly abstract, these pillars aim to provide a general methodology to underpin the notion of a hierarchical aggregation of knowledge, an artefact emerging throughout the literature.

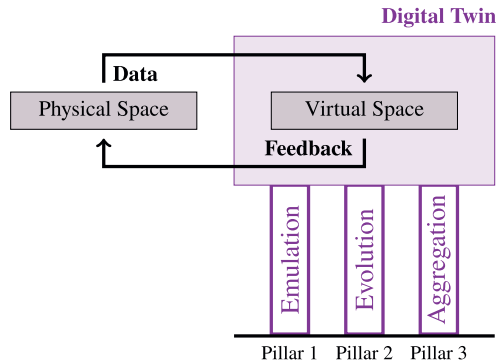


Figure 1. An alternative view of digital twin concepts based on three pillars: allowing extensions through aggregation.

2.1. [Pillar 1] Emulation: the mirror

DTs are an adaptive, trace-driven emulation: evolving in parallel with the real world and driven by sensory input. A DT should mirror the response of an associated real-world object, for a range of given inputs. The emulators may branch to trial *what-if* scenarios for hypothetical inputs. For example, rather than opening or closing the physical sluice gates of a dam (and observing reservoir and river water levels), we first experiment in the virtual. The results are utilised to inform decision analysis and avoid various flood risks.

Simulations of hypothetical scenarios are *what-if* mirrors. These are viewed as distorted reflections, used to explore potential future outcomes. Furthermore, a mirror's reflection is never perfect, representative of how DTs approximate reality.

2.2. [Pillar 2] Evolution: data pathways

A barrier in real-world applications is not always a lack of data, but a lack of data management and *interoperability*. Data are often divided across different mediums, formats, and data silos. For DTs to function at scale, these management obstacles must be overcome to facilitate readily available and exchangeable data. Alongside measurements, we consider model parameters, simulation results, asset specifications, and so forth, as DT data. That said, measured (sensor) data are always essential to validate DT representations through time. The “real-time” requirement is always application-specific. In fact, requirements will be some form of quasi-real-time: for example, industrial applications might consider shifts or production dates; on the other hand, oceanographers consider real time as days or even fortnights (Venkatesan et al., 2017).

2.3. [Pillar 3] Aggregation: modular knowledge

Deviating from the mindsets presented in *Nature* (Bauer et al., 2021a,b), DTs are an aggregate, probing a design space that scales from systems to ecosystems. The constituent parts can be a mix of *bottom-up* and *top-down* models. For example, traffic models (bottom-up) and rainfall models (top-down) might combine to form a wider smart city DT. Even at small scales, a single-asset DT represents sensing, feedback, control, and decision-making systems: such a representation is an *aggregate* of knowledge and models from the start. Two-way communication pathways between constituent parts are critical, recently posed as a knowledge-graph problem (Wagg et al., 2024).

We highlight the requirement to include natural as well as engineered systems, whereby larger groups of DTs or sub-systems interact. Subspace interactions will leverage methods from the aforementioned fields, as well as meteorology, behavioural science, and sociology. In all cases, the core idea is a computational system that tracks the real world as closely as possible, driven by data, which can (to some extent) be considered a continuously updating model. Aggregation accommodates a complete

representation of the *life cycle* within a changing environment, capturing the interactions between virtual components, decision analyses, and sensing systems. Critically, a DT's level of complexity should be designed around its application.

3. A hierarchy of scales

DTs started with modest ambition—in effect, a *product* of engineering asset management. The vision has evolved, and as we expand from smart homes to smart cities, the knowledge encapsulated within DTs increases. We now focus on the sensing, monitoring, and control of entire engineered ecosystems. We explore the consequences of this progression in the design space, extending the ideas of hierarchy, aggregation, and federation from Section 4. First, we summarise increased expectations as a hierarchy of scales (1–4):

1. **System-specific:** Typical procedures and tools for the monitoring and simulation of a single system or asset, relating to conventional asset management, prognostics, and control (Bruynseels et al., 2018; Barricelli et al., 2019; Furini et al., 2022). The *boundaries* of the virtual representation (or the limits of the DT) are intuitive, with clear approximations.
2. **Local collection:** Aligned with the ideas from the Internet of Things, including connected, distributed representations on a smaller scale. Examples include industrial applications (Kappel et al., 2022), urban farming in the United Kingdom and the United States (Jans-Singh et al., 2020), and land-based aquaculture (Lima et al., 2022). The boundary of the representation becomes less clear—for example, in agriculture, the soil depth and surrounding environment become increasingly complex.
3. **Ecosystems:** Typical applications include urban analytics, smart infrastructure (Oughton and Russell, 2020), population-based structural health monitoring (Worden et al., 2020), and Industry 5.0 applications (Braque et al., 2021). More recently, large-scale applications in health, medicine, and education are emerging (Björnsson et al., 2020; Furini et al., 2022). While these aggregates are typically constituted of engineered systems, with defined boundaries, the limit on the extent of the DT representation is harder to define.
4. **Planetary:** Distributed and parallel DTs combined to accomplish scaled-up tasks of complex interconnected systems (sensing for weather, pollution, large-scale energy modelling, etc.). Highly complex boundaries and interactions are associated with the aggregate.

Unsurprisingly, it becomes difficult to define large-scale DTs in plain practical terms. The concepts of aggregation (Grieves and Vickers, 2017) are steadily increased from the combination of virtual subsystems (for a single artefact) to many nested groups of interacting models. Various issues are related to combining and coupling DTs: emergent properties, impedance matching, and uncertainty proportions. The required systems thinking will have both cultural and theoretical issues lying in wait, regarding modelling and data management. All of the above highlight a (pre-established) necessity to formalise message passing and knowledge transfer between DT components. In other words, a framework for the combination of models, data, and outputs from many related DTs.

3.1. Knowledge transfer

An active research considers sharing information between DTs via *knowledge transfer*—a natural *modelling* pathway to address the concerns of aggregation (Bennett et al., 2023; Bull et al., 2023). Here, we use *knowledge transfer* as an umbrella term to refer to tools that aid in sharing information (in some sense) between DT representations: for example, multitask learning, transfer learning, multi-view learning (Hernandez-Leal et al., 2019), or *federated learning* (Zhang et al., 2021). Broadly speaking, these procedures can support inferences from multiple related datasets and, in turn, offer some mechanism to combine the virtual sub-components of DT representation. Project Odysseus was one example, presenting part of an urban DT; it federates many (data-driven) models of different components to

monitor activity during the lockdowns of COVID-19 in London (<https://www.turing.ac.uk/research/research-projects/project-odysseus-understanding-london-busyness-and-exiting-lockdown>).

Federation is often considered a silver bullet to the interoperability of knowledge domains (via models + data) while prioritising privacy protection. However, it is only suitable provided that connections between constituent parts are adequately defined and implemented. Practically, DT aggregates will suffer from scalability issues if a robust, extensible, and flexible scheme is not in place. Ad hoc peer-to-peer interactions between DTs could present an alternative decentralised system, which should naturally integrate with a modular and systems-thinking philosophy (e.g., consider distributed ledgers and their impact on alternative currencies). Furthermore, decentralised storage of *digital assets* would attract less opposition from asset owners, when it comes to knowledge transfer. Our view is that knowledge and insights extracted from dissimilar models could be implemented as a dual-role system, allowing each agent to produce and consume knowledge: the exchange of models, data, and contextual information within a wider network. These concepts are further discussed in [Section 5](#), with regard to privacy.

We introduce a hypothetical example to support the hierarchy of scales and DTs by aggregation.

3.2. An example of increasing scale: wind energy

[Figure 2](#) presents one natural example of DT aggregation, progressing through the hierarchies of scale with lighter shades of purple.

1. **System-specific:** A single turbine—even at this scale, the DT is an aggregate of data + models, for example, corresponding to different components and their interactions (the nacelle, blades, and foundations).
2. **Local collection:** The wind farm of constituent turbines. The boundary becomes more complex: where is the limit of representation and how do we capture the interactions between other turbines, the grid, and the environment (on or off-shore)?
3. **Ecosystem:** Alternative power sources are now included, alongside human agents. Interoperability of DT information becomes key, especially between organisations. The boundary extends to both natural (human as an agent) and engineered system behaviour.
4. **Planetary:** Repeated extensions (into the background). The DT combines several *whole energy systems*; this could be international. Note that the elements of the aggregate can be shared between parts (e.g., the weather models).

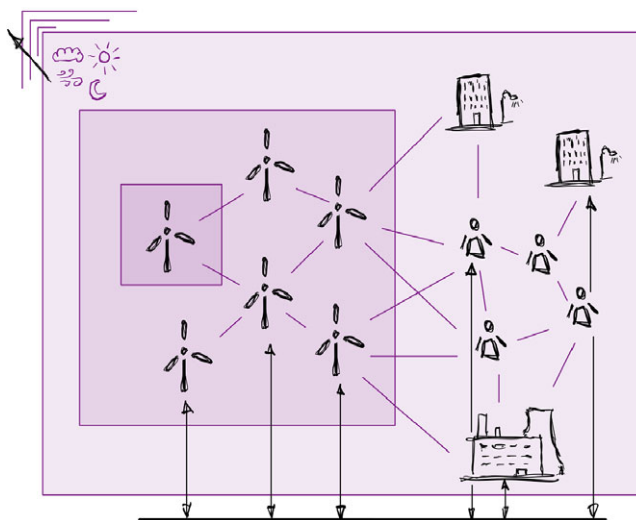


Figure 2. Scaling digital twins by aggregation: wind energy systems.

4. The design space: models + data

Traditionally, machine learning researchers have adopted a model-centric mindset; their focus is on model architectures, tuning hyperparameters, or tailored loss functions to improve *prediction*. An alternative mindset is emerging, which places the data themselves at the centre of the philosophy—*Data-centric ML* (or data-centric AI; Strickland, 2022). We see this approach as harmonious with conventional statistical, scientific, and engineering processes. Given our process understanding, encoded in causal or scientific models (domain expertise), *how do we collect data* to test our hypotheses? Experimental design and data curation become a key focus at least as important as model design. We have already discussed how *process understanding* is critical for DTs at scale, and a data-centric mindset conveniently integrates with this philosophy to curate datasets that further our understanding of the properties and behaviours of the systems of interest.

A simple consequence of this view is that we consider the design space as an extension of the model space, as shown in Figure 3. We encompass data and deployment, as well as models/algorithms, to constitute the evolving virtual representation. Process understanding must inform the design of each aspect.

To demonstrate, we discuss these extensions around a few important dimensions.

Modular data collection: Modular representations will combine many partitioned systems—their data are likely to be siloed and/or incomplete; however, the whole is something *besides* its parts (*Metaphysics*, Aristotle, 1908, translated by W.D. Ross). Here, a data-centric mindset has great potential, where data slicing includes cleaning, wrangling, and other data transformations in view of the collected representation rather than individuals (Data-centric AI: <https://spectrum.ieee.org/andrew-ng-data-centric-ai>). Aside from the push for *big-data* and universal models, we believe that most successful DTs will focus on curated, high-quality data and a combination of relatively specialised models to provide interpretable representations.

Encoding prior knowledge: The formulation of models must be carefully considered. This allows us to encode our prior intuition within DT representations, given knowledge of the underlying processes. Typically, this knowledge is encoded in two ways: (i) through interpretable model structures and (ii) our prior belief of the underlying distribution of unknown parameters. The prior representation is critical since it exists *before any data are observed*: before the model has learnt from data. With reference to Pillar 1, the prior ensures the quality of the emulation, constraining and informing the mirror, so that predictions *align* with reality. In a survival model example, predicting a negative age for life expectancy is clearly unreasonable: the DT representation must not permit such a prediction. The emergent areas of physics-informed machine learning (Mao et al., 2020) will prove critical here to adjust/constrain the model space to reflect reality. Importantly, the prior is essential to inform experimental design and data collection (implicitly or directly) as it underpins a data-centric mindset and informs our curated data. Caution is required, however. Like data, if our prior knowledge is incorrect, this can be detrimental to the DT representation.

Cyber-physical components: Of course, cloud computing (private or public) presents a natural starting point for DT models. This follows the scaling of data computation/storage, alongside various new




Figure 3. *Data-centric machine learning: folding models, data, and process understanding into the design space.*

software for parallel, distributed, and federated learning. Not only does a cloud-based residency for DT models facilitate the distribution of knowledge outside traditional data silos, it supports the necessitated scale required to achieve aggregated, planetary scale DTs. The alignment of distributed sensors to cloud-aware technologies presents an exciting opportunity. Sensors and actuators have been integrated within engineering solutions for hundreds of years, with the theory of control feedback dating back to James Clerk Maxwell (1868). As such, a distributed approach (Fuller et al., 2020) to data transfer, storage, and processing should naturally integrate with DTs at scale. With on-device data analytics receiving notable attention (Bagchi et al., 2019) and discussed in Section 5, the opportunity to tailor data collection at the edge, around the requirements of the DT, is presented.

Another interesting trend follows smaller components and the capability to run on scavenged power. These systems do not require connection to the power grid, enabling sensing at increasingly remote locations. Additionally, the drive to make bio-degradable (or at least neutral) devices for environmental sensing is emerging (Zhu et al., 2023).

Systems heterogeneity: Two concepts are important—heterogeneity of the physical twin, alongside that of the virtual. In the virtual sub-spaces, structured and parameterised models are more easily categorised into groups that inform knowledge transfer and the potential for data sharing, as similar categorisation would be useful for physical systems (Gosliga et al., 2021). An intuitive example returns to the wind farm, with turbines of the same make and model. These systems will have interpopulation variances—local environment, manufacturing differences, and location within the farm. Another wind farm might include alternative turbine models (e.g., analogous to a population of similar species). Despite many differences between the machine groups, these two populations should still share characteristics (or parameters)—for example, the parametrisation of their power curves (Bull et al., 2023). The design of both data and models informs the suitability and manner of DT aggregation.

Openness of the system at the edges: We can attempt to measure all aspects of a railway system, but the entire ecosystem is far too large and complex. We must draw boundaries to constrain *both* measurement and modelling, rather than modelling alone. For example, if a DT estimates the response of a bridge from train loading, how do we model the varying mass of commuter trains? Do we average, or model passenger dynamics at rush hour? We see that boundary definitions become increasingly difficult as DTs include naturally occurring systems, that is, human behaviour or weather/ocean models (UN DITTO programme on DTs of the Ocean <https://ditto-oceandecade.org/>). Another consideration of boundaries is purely virtual, concerning legacy software; for example, this could be resolved with wrappers for archived code in evolutionary architectures.

Levels of autonomy: One aim of DTs is to reduce the reliance (and overspending) on human interactions. Progress will require investigations of control and decision-making informed by data (Langtry et al., 2023), likely towards more AI-based control, such as reinforcement learning (Cronrath et al., 2019). By guiding interactions and data collection with the virtual representation, we are once again (quite naturally) within a data-centric ML mindset: whereby models, actions, and data collection evolve alongside their physical counterparts. Autonomy implies decreasing levels of human input; however, the rate at which human interactions are replaced will depend on the type of action. For example, the formulation of domain expertise will be very difficult to automate (and illogical) in most settings.

An extensive system design challenge exists for a scalable DT architecture. From a purely technical standpoint, we propose that systems researchers should consider a clean-slate approach to scale distributed processing. That said, considerable investment in the current systems could make any new architecture financially prohibitive. Discussions would be similar to those surrounding new Internet architectures (Feamster et al., 2014).

5. Beyond data privacy

Privacy must go beyond the data themselves to encompass the evolving aspect of aggregate models. While the data typically come from critical infrastructure or commercial services, they should not be the only concern regarding privacy. As DTs develop, we require a new type of governance to ensure secure




Figure 4. Privacy concerns when scaling digital twins—smart energy meters.

knowledge exchange from sensing to deployment. Most obviously, this includes both *data* and *modelling* components of continually evolving representations. During DT design, it will be important to incorporate privacy-enhancing technologies, both hardware and software, to protect the models and data within the core system. The approach should minimize the possession of personal data (Bruynseels et al., 2018) at any of the four stages of the hierarchy of scales. Whilst the sophistication of knowledge increases at each scale—through the aggregation of previous levels of knowledge and data—so does the repercussion of any privacy implications due to residual personal data. We discuss privacy considerations while progressing through the hierarchy, illustrated in Figure 4.

5.1. (Scale 1) System-specific privacy

Smart energy meter programs, implemented in various countries, aid households in managing efficient energy usage and reducing costs and CO₂ emissions. These programs utilise household-specific smart meters, collecting high-resolution data and providing valuable insights for decision-making. They also present the opportunity for simple energy-use DTs (at the household level). Computation at the edge, such as generating a basic forecast given historical data patterns, eliminates the need to transmit data or models. When transmission occurs, techniques like differential privacy, data obfuscation, and homomorphic encryption (Yin et al., 2022) facilitate secure data sharing while protecting user privacy.

5.2. (Scale 2) Local collection privacy

A more sophisticated DT might attempt to represent constituent appliances within a user's home. Non-intrusive load monitoring (NILM) algorithms are already capable of accurately disaggregating total energy consumption, by deducing the use of appliances from changes in the voltage and current, monitored by smart meters. In fact, these methods can predict appliance use with up to 95% accuracy (McKenna et al., 2012; Teng et al., 2022). In certain cases, such data might expose consumers' daily habits, posing a privacy risk. Importantly, while the underlying data types remain the same between each scale, the knowledge embedded in the DT and privacy concerns have increased, as further insights are enabled.

5.3. (Scale 3) Ecosystems privacy

The transition to ecosystems leads to complex boundaries and thus complex privacy concerns, with the potential for multiple institutions or organizations. Non-disclosure aspects of trading and intellectual property become pertinent. Notably, in the context of data transformation pipelines, it is essential to maintain transparency and comply with provenance guidelines. This includes documenting the transformations from primary data to derived representations and models (Barricelli et al., 2019). Edge computing holds high potential at this scale, as these schemes can deploy ML on low-powered devices for privacy preservation (Warden et al., 2023). Such edge devices offer the potential to scale computation through federated models rather than data. However, this typically involves complex model aggregation servers, especially when dealing with non-IID data and potential adversaries. Returning to smart meters,

recent advances in the aggregation of city and region-wide metering pose a model-agnostic solution to federation. This framework enables collaborative training of NILM models while addressing privacy concerns at the ecosystems level (Shi et al., 2023). Similarly, DARE UK is a nationwide initiative that exploits the federation of trusted research environments (TREs) to manage sensitive data for research purposes securely—also relating to TRUST principles discussed in Section 6. Exploratory case studies are underway to examine federated networks of TREs, where data remain localised and the analysis is shared.

5.4. (Scale 4) Planetary privacy

Any feasible, resilient approach when scaling DTs internationally will require increasingly decentralised computation rather than centralised aggregation servers. One approach, such as the Internet's root servers, decentralises the knowledge of top-level Domain Name System (DNS) records. Peer-to-peer learning should be combined with the aforementioned privacy-preserving mechanisms, for example, obfuscation (Teng et al., 2022) to reduce the information revealed during model updates. Shi et al. (2023) suggest a learning framework based on peer-to-peer communication between participant nodes. The decentralised paradigm offers one promising path to address privacy issues for sensor measurements, processed data, and simulated data; however, it is yet to be validated at larger scales. In the context of power systems, another example might consider submarine cables, which would prove essential in the realisation of interconnected power system DTs. Since these cables often span multiple jurisdictions, adherence to various national and international laws presents a number of challenges to ensure privacy compliance (Helbing and Argota Sanchez-Vaquerizo, 2022) (e.g., General Data Protection Regulation (GDPR)). Privacy protection that traverses country-specific jurisdictions has been proposed in the context of distributed ledger systems (Roberts et al., 2023).

6. Ethical considerations

The proposed three pillars of DT representation rely on data pathways to update the mirror (emulator) evolving through time. We can then virtually trial *what-if* scenarios (Schluse et al., 2018), investigating how alternative interactions might influence real-world processes and avoiding costly *in-situ* experiments. While the experiments are virtual, the results inform *real-world* decisions. Helbing and Argota Sanchez-Vaquerizo (2022) have recently discussed the ethical implications of viewing DTs in this way as *duplicates* (or imitations) of reality that interfere with our individual thoughts, decisions, and behaviour. Consider a refocus towards humans, rather than automation, in Industry 5.0 (Commission et al., 2021). As we bring humans *back-in-the-loop* critical concerns relating to human rights and dignity are raised. That is, we expect to interact within increasingly virtual spaces (Commission et al., 2021). Importantly, since the outputs of any virtual representation can have a very real effect on human experience, they must be subject to ethical scrutiny. As we aggregate DTs, the concerns of misuse progress from individual (or corporate) to societal. To aid discussion, we adopt the considerations of Helbing and Argota Sanchez-Vaquerizo (2022) and Caldarelli et al. (2023) to reflect on some ethical design issues (EDIs) when scaling DTs.

6.1. EDI 1 (complexity and uncertainty)

As we increase complexity, it is vital that we represent the cumulative uncertainty. Relatively speaking, smaller scales might present manageable variations over time, though this will propagate for highly dynamic and interconnected systems. For robust decision analysis, all *real-world* actions must rest upon appropriate quantification of uncertainty in the representation.

6.2. EDI 2 (optimisation in twinning)

Intuitively, the cost function of an optimisation process will become more intricate at scale. The objectives become complex, regarding the boundary and environment variables—including human behaviour,

natural resources, and interactions. These optimisers must also consider legal, corporate, and cultural constraints—limitations that are not immediately quantifiable (or recognised). Helbing and Argota Sanchez-Vaquerizo (2022) suggest that coordination approaches are appropriate over more conventional control methods. When embedded in the DTs, optimisers should rarely attempt to map real-world complexity to one dimension. Such a compression will rarely be sensible and highlight why humans are required in the loop.

6.3. EDI 3 (humans and things)

As humans are folded back into the loop of digitisation, engineers must consider the responsibilities of representing societies and individuals. Best practices have long been central to the fields of sociology, anthropology, and epidemiology; however, DTs present some novel concerns. Bruynseels et al. (2018) argue that the mere existence of one's DT might discriminate against real people, of whom the twins are a representation. Would a DT ever respond in a manner which would negatively impact a human population or human life? Though in practice, rather than replacing humans, we see the potential to utilise *machines* (virtual or not) cooperating with humans, as *cobots*, or *collaborative robots* (Kumar et al., 2021). In agreement with Ginni Rometty, we view the tools and insights of AI as an opportunity to augment our intelligence, rather than replace it.

6.4. EDI 4 (data)

Increased data sharing is essential; however, this must happen *responsibly*. While FAIR principles remain at the top of the agenda, more scrutiny, regulation, and preventative measures are required, especially when sharing data between organisations at the ecosystem level. More industry-focussed and open guidelines are emerging, for example, the TRUST principles (Lin et al., 2020), which consider the interests of stakeholders and intellectual property. These ideas are critical to realise the vision of a platform whereby data and models are no longer trapped in *silos* and utilised across the hierarchy. Some propose a *logic data highway* that connects regional and global DTs (Royal Society, 2020); however, such convergences must respect the principles of FAIR and TRUST. We call for interfaces with open standards and protocols to facilitate the operation of DTs between countries. The use case should be paramount when sharing data. For example, vehicle driving data might be used to inform driving practice or autonomous vehicles to work towards safer cities. On the other hand, insurance companies may try to access these data to determine insurance premiums. Balancing these interests clearly requires responsible and secure data sharing.

6.5. EDI 5 (data-centric AI)

Transparency and interpretability of DTs should be paramount, and while the model-centric (more black-box) mindset should be utilised (Brown et al., 2020), this should only happen when necessary—Occam's razor prevails. Aside from the fact that training procedures can serve a significant cost to the environment (Strubell et al., 2019), model interoperability remains problematic. These costs and lack of interpretability are counterproductive when simpler models provide the same insight. Wherever possible, alternative representations based on small, curated data should be utilised. A data-centric mindset prioritises the requirements of the job at hand and the design space of the DT. Transparency ensures *interoperability of DTs*, which will be essential for robust knowledge exchange between constituents.

7. Concluding remarks

Rather than redefine, we have gathered emerging concepts in the DT literature—highlighting requirements and areas that should prove critical when moving forward. In doing so, we have identified several important themes: model/data aggregation, design space boundaries, privacy, and ethical concerns. While most researchers recognise these elements as crucial, they are usually discussed separately, across

different communities, and rarely considered holistically. The backbone of our discussion follows a *hierarchy of scales* that correspond to increasingly aggregated units. Following this modular approach, the focus when collecting DTs is on the framework of connections, rather than the models themselves. As such, local expert knowledge (and ownership) of the real-world system can stay connected to local DT units, while the aggregate is naturally decentralised, with shared control. Despite a modular structure, the whole can be iteratively refined through the pillars of emulation and evolution, as discussed. This mindset of DT *representation* naturally addresses concerns throughout many themes from practical implementation and sensing to ethics and privacy. By embracing these ideas, the vision of an interconnected DT ecosystem becomes more attainable.

Data availability statement. Data availability is not applicable to this article as no new data were created on in this study.

Author contribution. Conceptualisation: M.P., J.C., L.B., M.G., D.S.B.; Writing—original draft: M.P., J.C., L.B., M.G., D.S.B. All authors approved the submitted draft.

Acknowledgements. The authors want to thank the anonymous reviewers for their contributions to improving this article.

Funding statement. L.B., D.S.B., and M.G. acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC) through the ROSEHIPS project (Grant No. EP/W005816/1). L.B. was also supported by Wave 1 of the UKRI Strategic Priorities Fund under the EPSRC Grant EP/W006022/1, particularly the *Ecosystems of Digital Twins* theme within that grant and the Alan Turing Institute. M.P. acknowledges the support of the European Commission through the H2020 ASTRAL project (Grant No. 863034).

Competing interest. The authors declare no competing interests exist.

References

- Bagchi S, Siddiqui M-B, Wood P and Zhang H (2019) Dependability in edge computing. *Communications of the ACM* 63(1), 58–66. <https://doi.org/10.1145/3362068>.
- Barricelli BR, Casiraghi E and Fogli D (2019) A survey on digital twin: Definitions, characteristics, applications, and design implications. *IEEE Access* 7, 167653–167671. <https://doi.org/10.1109/ACCESS.2019.2953499>.
- Bauer P, Dueben P, Hoeffler T, Quintino T, Schulthess T and Wedi N (2021a) The digital revolution of Earth-system science. *Nature Computational Science* 1, 104–113. <https://doi.org/10.1038/s43588-021-00023-0>.
- Bauer P, Stevens B and Hazeleger W (2021b) A digital twin of earth for the green transition. *Nature Climate Change* 11, 80–83. <https://doi.org/10.1038/s41558-021-00986-y>.
- Bennett H, Birkin M, Ding J, Duncan A and Engin Z (2023) Towards ecosystems of connected digital twins to address global 31 challenges. <https://doi.org/10.5281/zenodo.7840266>. Published May 18, 2023 | Version v1.
- Björnsson B, Borrebaeck C and Elander N (2020) Digital twins to personalize medicine. *Genome Medicine* 12, 4. <https://doi.org/10.1186/s13073-019-0701-3>.
- Blair GS (2021) Digital twins of the natural environment. *Patterns* 2(10), 100359. <https://doi.org/10.1016/j.patter.2021.100359>.
- Braque M, De Nul L and Petridis A (2021) *Industry 5.0: Towards a sustainable, human-centric and resilient European industry*. European Commission, Directorate-General for Research and Innovation. <https://data.europa.eu/doi/10.2777/308407>.
- Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler D, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I and Amodei D (2020) Language models are few-shot learners. In Larochelle H, Ranzato M, Hadsell R, Balcan M and Lin H (eds.), *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., pp. 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf.
- Bruynseels K, Santoni F and van den Hoven J (2018) Digital twins in health care: Ethical implications of an emerging engineering paradigm. *Frontiers in Genetics* 9, 31. <https://doi.org/10.3389/fgene.2018.00031>.
- Bull L, Di Francesco D, Dhada M, Steinert O, Lindgren T, Parlikad AK, Duncan A and Girolami M (2023) Hierarchical Bayesian modelling for knowledge transfer across engineering fleets via multitask learning. *Computer-Aided Civil and Infrastructure Engineering* 38(7), 821–848.
- Caldarelli G, Arcaute E, Barthelemy M and Helbing D (2023) The role of complexity for digital twins of cities. *Nature Computational Science* 3, 374–381. <https://doi.org/10.1038/s43588-023-00431-4>.
- Chantry M, Christensen H, Dueben P and Palmer T (2021) Opportunities and challenges for machine learning in weather and climate modelling: Hard, medium and soft AI. *Philosophical Transactions of the Royal Society A* 379, 20200083. <http://doi.org/10.1098/rsta.2020.0083>.
- Clerk Maxwell J (1868) On governors. *Proceedings of the Royal Society of London* 16, 270–283. <https://doi.org/10.1098/rsp1.1867.0055>.

- Cronrath C, Aderiani AR and Lennartson B** (2019) Enhancing digital twins through reinforcement learning. In *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, pp. 293–298.
- European Commission for Research and Innovation, Breque M, De Nul L and Petridis A** (2021) *Industry 5.0: Towards a Sustainable, Human-Centric and Resilient European Industry*. Publications Office of the European Union. <https://doi.org/10.2777/308407>.
- Feamster N, Rexford J and Zegura E** (2014) The road to SDN: An intellectual history of programmable networks. *SIGCOMM Computer Communication Review* 44(2), 87–98. <https://doi.org/10.1145/2602204.2602219>.
- Fuller A, Fan Z, Day C and Barlow C** (2020) Digital twin: Enabling technologies, challenges and open research. *IEEE Access* 8, 108952–108971. <https://doi.org/10.1109/ACCESS.2020.2998358>.
- Furini M, Gaggi O, Mirri S, Montangero M, Pelle E, Poggi F and Prandi C** (2022) Digital twins and artificial intelligence: As pillars of personalized learning models. *Communications of the ACM* 65(4), 98–104. <https://doi.org/10.1145/3478281>.
- Glaessgen E and Stargel D** (2012) The digital twin paradigm for future NASA and US air force vehicles. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA*, p. 1818.
- Gosliga J, Gardner P, Bull L, Dervilis N and Worden K** (2021) Foundations of population-based SHM, part II: Heterogeneous populations—graphs, networks, and communities. *Mechanical Systems and Signal Processing* 148, 107144.
- Grievens M and Vickers J** (2017) Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In: Kahlen J, Flumerfelt S and Alves A, (eds), *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches*. pp. 85–113. Springer: Cham. https://doi.org/10.1007/978-3-319-38756-7_4.
- Helbing D and Argota Sanchez-Vaquerizo J** (2022) Digital twins: Potentials, ethical issues, and limitations. In *Elsevier Handbook on the Politics and Governance of Big Data and Artificial Intelligence*. <https://doi.org/10.2139/ssrn.4167963>.
- Hernandez-Leal P, Kartal B and Taylor ME** (2019) A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 33(6), 750–797. <https://doi.org/10.1007/s10458-019-09421-1>.
- Jans-Singh M, Leeming K, Choudhary R and Girolami M** (2020) Digital twin of an urban-integrated hydroponic farm. *Data-Centric Engineering* 1, e20. <https://doi.org/10.1017/dce.2020.21>.
- Kappel G, Brecher C, Brockmann M and Koren I** (2022) Internet of production: Entering phase two of industry 4.0. *Communications of the ACM* 65(4), 50–51. <https://doi.org/10.1145/3514093>.
- Kumar S, Savur C and Sahin F** (2021) Survey of human–robot collaboration in industrial settings: Awareness, intelligence, and compliance. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 51(1), 280–297. <https://doi.org/10.1109/TSMC.2020.3041231>.
- Langtry M, Zhang C, Ward R, Makasis N, Kreitmair MJ, Conti ZX, Di Francesco D and Choudhary R** (2023) *Value of information analysis for rationalising information gathering in building energy analysis*. Preprint, [arXiv:2305.16117](https://arxiv.org/abs/2305.16117).
- Lima A, Royer E, Bolzonella M and Pastres R** (2022) Digital twins for land-based aquaculture: A case study for rainbow trout (*Oncorhynchus mykiss*). *Open Research Europe* 2, 16. <https://doi.org/10.12688/openreseurope.14145.1>.
- Lin D, Crabtree J, Dillo I, Downs RR, Edmunds R, Giaretta D, De Giusti M, L'Hours H, Hugo W, Jenkyns R, Khodiyar V, Martone ME, Mokrane M, Navale V, Petters J, Sierman B, Sokolova DV, Stockhouse M and Westbrook J** Placeholder Text (2020) The trust principles for digital repositories. *Scientific Data* 7(1), 144.
- Mao Z, Jagtap AD and Karniadakis GE** (2020) Physics-informed neural networks for high-speed flows. *Computer Methods in Applied Mechanics and Engineering* 360, 112789.
- McKenna E, Richardson I and Thomson M** (2012) Smart meter data: Balancing consumer privacy concerns with legitimate applications [Modeling transport (energy) demand and policies]. *Energy Policy* 41, 807–814. <https://doi.org/10.1016/j.enpol.2011.11.049>.
- NASEM** (2024) *Foundational Research Gaps and Future Directions for Digital Twins*. The National Academies Press: Washington, DC. <https://doi.org/10.17226/26894>.
- Nativi S, Mazzetti P and Craglia M** (2021) Digital ecosystems for developing digital twins of the earth: The destination earth case. *Remote Sensing* 13(11), 2119. <https://doi.org/10.3390/rs13112119>.
- Negri E, Fumagalli L and Macchi M** (2017) A review of the roles of digital twin in cps-based production systems [27th international conference on flexible automation and intelligent manufacturing, FAIM2017, 27–30 June 2017, Modena, Italy]. *Procedia Manufacturing* 11, 939–948. <https://doi.org/10.1016/j.promfg.2017.07.198>.
- Oughton EJ and Russell T** (2020) cdcam: Cambridge digital communications assessment model. *Journal of Open Source Software* 5(45), 1911. <https://doi.org/10.21105/joss.01911>.
- Pyliandis C, Osinga S and Athanasiadis IN** (2021) Introducing digital twins to agriculture. *Computers and Electronics in Agriculture* 184, 105942.
- Roberts JD, Defranco JF and Kuhn DR** (2023) Data block matrix and hyperledger implementation: Extending distributed ledger technology for privacy requirements. *ACM Distributed Ledger Technologies*. 2(2), 1–11. <https://doi.org/10.1145/3585539>.
- Royal Society** (2020) *Digital Technology and the Planet: Harnessing Computing to Achieve Net Zero*. The Royal Society: London.
- Schluse M, Priggemeyer M, Atorf L and Rossmann J** (2018) Experimentable digital twins—Streamlining simulation-based systems engineering for industry 4.0. *IEEE Transactions on Industrial Informatics* 14(4), 1722–1731. <https://doi.org/10.1109/TII.2018.2804917>.
- Shi Y, Li W and Chang X** (2023) On enabling collaborative non-intrusive load monitoring for sustainable smart cities. *Scientific Reports* 13, 6569. <https://doi.org/10.1038/s41598-023-33131-0>.

- Strickland E** (2022) Andrew ng, ai minimalist: The machine-learning pioneer says small is the new big. *IEEE Spectrum* 59(4), 22–50. <https://doi.org/10.1109/MSPEC.2022.9754503>.
- Strubell E, Ganesh A and McCallum A** (2019) Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3645–3650.
- Teng F, Chhachhi S, Ge P, Graham J and Gunduz D** (2022) *Balancing Privacy and Access to Smart Meter Data*. Imperial College London. www.imperial.ac.uk/energy-futures-lab/reports/briefing-papers/paper-9/.
- Venkatesan R, Tandon A, D’Asaro E and Atmanand MA** (2017) *Observing the Oceans in Real Time*. Cham: Springer. <https://doi.org/10.1007/978-3-319-66493-4>.
- Wagg D, Burr C, Shepherd J, Conti ZX, Enzer M and Niederer S** (2024) The philosophical foundations of digital twinning. Preprint. <https://engrxiv.org/preprint/view/3500>. The philosophical foundations of digital twinning. DOI: <https://doi.org/10.31224/3500>.
- Warden P, Stewart M, Plancher B, Katti S and Reddi VJ** (2023) Machine learning sensors: A design paradigm for the future of intelligent sensors. DOI10.1145/3586991. Published: November 1, 66(11), 25–28.
- Worden K, Bull LA, Gardner P, Gosliga J, Rogers TJ, Cross EJ, Papatheou E, Lin W and Dervilis N** (2020) A brief introduction to recent developments in population-based structural health monitoring. *Frontiers in Built Environment* 6, 146.
- Ye X, Du J, Han Y, Newman G, Retchless D, Zou L, Ham Y and Cai Z** (2023) Developing human-centered urban digital twins for community infrastructure resilience: A research agenda. *Journal of Planning Literature* 38(2), 187–199. <https://doi.org/10.1177/08854122221137861>.
- Yin X, Zhu Y and Hu J** (2022) A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. *ACM Computing Surveys* 54(6), 131. <https://doi.org/10.1145/3460427>.
- Zhang J, Guo S, Ma X, Wang H, Xu W and Wu F** (2021) Parameterized knowledge transfer for personalized federated learning. In Ranzato M, Beygelzimer A, Dauphin Y, Liang P and Vaughan JW (eds.), *Advances in Neural Information Processing Systems* Placeholder TextPlaceholder Text, Vol. 34. Red Hook, NY: Curran Associates, Inc., pp. 10092–10104. <https://proceedings.neurips.cc/paper/2021/file/5383c7318a3158b9bc261d0b6996f7c2-Paper.pdf>.
- Zhu J, Wen H, Zhang H, Huang P, Liu L and Hu H** (2023) Recent advances in biodegradable electronics—from fundament to the next-generation multi-functional, medical and environmental device. *Sustainable Materials and Technologies* 35, e00530. <https://doi.org/10.1016/j.susmat.2022.e00530>.