Digital Twins in Application: A Survey by Sector and by Twin Type

1. Scope, reader, and how to read this survey

This survey is written for a researcher entering the digital twin field who needs a map of where digital twins are actually being built, what they are being built *of*, and how far the reported systems are from the claims made on their behalf. It is not an introduction to the concept and it does not teach the enabling technologies. It assumes the reader already knows roughly what a digital twin is meant to be and wants to know what the application literature supports.

The organising problem is that "digital twin applications" is two questions wearing one coat. The first is a question about **sector**: who is building these things — manufacturers, hospitals, city governments, wind farm operators? The second is a question about **type**: what class of thing is being twinned — a product, a process, a whole system, a human being? These two questions cut across each other. A pharmaceutical company and a bridge authority may both be building "process twins" while sharing almost no methods; a hospital and an aircraft operator may both be building "asset twins" of a single physical unit while facing utterly different validation regimes. Surveys that pick only one axis end up either with a list of industries that says nothing about engineering substance, or with a taxonomy that floats free of any actual deployment. This survey uses both axes deliberately, and adds a third — the **degree of integration** between physical and virtual — because that axis turns out to be the one that most often separates the claim from the delivery.

What is covered: manufacturing and the process industries; healthcare, medicine and the human body; the built environment, cities, civil infrastructure and mobility; energy, aerospace and maritime engineering; agriculture, food and environmental science; and the cross-cutting concerns of standards, platforms, security and validation that no sector escapes.

What is deliberately excluded, so that a reader can tell an omission from an oversight. First, this is a survey of *applications*, not of enabling technologies; sensing, communication protocols, model-order reduction and machine-learning methods appear only where an application's credibility turns on them. Second, it is bounded by a specific synced corpus, which has a pronounced centre of gravity in engineering, software engineering and cyber-physical systems. Section 10 reports which sectors that corpus genuinely does not cover, distinguishing those from sectors that are simply thin. Third, business-model, procurement and organisational-change questions are noted where sources raise them but are not surveyed in their own right. Fourth, and importantly, this survey does not attempt to adjudicate the definitional dispute described in Section 2; recording that dispute honestly is more useful to a new researcher than resolving it by fiat.

One framing observation before the map. The breadth of uptake is not in question. A *Nature Computational Science* focus issue records that "various areas of science — including, but not limited to, engineering, climate sciences, medicine, and social sciences — have realized the potential of digital twins," while noting in the same breath that "many challenges still need to be addressed before the research community can bring the promise of digital twins to fruition" [1]. The issue's editorial lists the hoped-for gains as better precision medicine, sharper weather and climate prediction and better-informed urban planning, and adds at

once that none of them comes without many unsolved problems [2]. That gap between realised potential and delivered promise is the recurring subject of the sections that follow.

2. Three axes for reading the application literature

2.1 The definitional problem is not academic

Any survey of digital twin applications inherits a problem from its sources: the sources do not agree on what they are describing. This matters practically, because a reader comparing a "digital twin of a jet engine" with a "digital twin of a cancer patient" needs to know whether the same noun is doing the same work.

The concept traces to Michael Grieves. VanDerHorn and Mahadevan locate its origin in "Michael Grieves' 2003 presentation on product life-cycle management based on his work with John Vickers," and record Grieves' original three components: a physical product in real space, a virtual representation of it in virtual space, and the connections of data and information tying the two together. Their assessment of what happened next is blunt: the subsequent proliferation of definitions "has diluted Grieves' original description," and their paper explicitly seeks to return to it, generalising it to three components — "(1) A physical reality, (2) a virtual representation, and (3) interconnections that exchange information between the physical reality and virtual representation" [3]. Grieves and Vickers themselves frame the purpose in terms that a purely representational reading misses: the digital twin exists to surface "unpredictable, undesirable emergent behavior in complex systems" — behaviours that only appear when components are aggregated into systems and that historically were discovered only when a physical object was already in service [4].

The most consequential recent definition comes from the US National Academies, whose consensus study defines a digital twin 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," adding the criterion that has become the field's most-cited discriminator: "The bidirectional interaction between the virtual and the physical is central to the digital twin" [5]. Two features of this definition deserve emphasis for a survey organised by sector. It explicitly admits *social* systems alongside natural and engineered ones, licensing the urban and societal applications of Section 5. And it makes bidirectionality constitutive rather than aspirational, which — as the following sections repeatedly show — disqualifies a large fraction of what the literature calls a digital twin.

Against this, the field's own reviewers concede that consensus has not arrived. Sabri notes flatly that "there is a lack of consensus on the definition, principles, and standards of digital twins in the industry" and that "the definition of digital twins is ever-changing" [6]. Jiang and colleagues identify the mechanism: the concept "has been successfully implemented for various purposes in many application spheres, such as aerospace engineering, robotics, smart manufacturing, renewable energy and process industry," and "the by-product of such widespread usage is that the concept of the DT has been interpreted in many different ways" [7]. Broad systematic reviews accordingly organise themselves around the definitional question rather than assuming it settled, asking "What is a Digital Twin?", "Where is appropriate to use a Digital Twin?" and "When has a Digital Twin to be developed?" as first-class research questions [8]. Further conceptual treatments and reviews of the evolution from concept to technology reach similar conclusions about the term's elasticity [9]–[11].

The practical advice for a reader is therefore: when a paper claims a digital twin, check which of the three axes below it actually satisfies.

2.2 Axis one: degree of integration

The single most useful discriminator in the literature is Kritzinger and colleagues' classification by level of data integration, which is worth stating precisely because it is so often cited loosely. A **Digital Model** does "not use any form of automatic data integration"; digital data from an existing physical system may inform the model, "but all data exchange is done in a manual way," and "a change in state of the physical object has no direct effect on the digital object and vice versa." A **Digital Shadow** adds automation in one direction: "if there exists an automated one-way data flow between the state of an existing physical object and a digital object, one might refer to such a combination as Digital Shadow. A change in state of the physical object leads to a change of state in the digital object, but not vice versa." A **Digital Twin** requires closure: "if further, the data flows between an existing physical object and a digital object are fully integrated in both directions, one might refer to it as Digital Twin" [12].

Their empirical finding is the reason this axis matters more than any other. Classifying the manufacturing literature against these levels, they report that "literature concerning the highest development stage, the DT, is scarce, whilst there is more literature about DM and DS" [12]. In other words, the field's own foundational classification paper found that most published digital twins were not digital twins.

Two later studies quantify that finding rather than repeating it. A cross-domain mapping study of 356 publications reports that only 21.63% "explicitly connect the Digital Twin with their real-world counterpart," and that among those making the interaction explicit, fewer than half support interaction in both directions [13]. An audit of 29 published manufacturing architectures finds control and actuation implemented in 17%, and states the conclusion without hedging: "current DT architectures tend to focus on functional aspects and describe digital shadows rather than twins" [14]. The same study's practitioner survey found 67% agreement with the proposition, one respondent remarking that "without a bidirectional interaction there is no digital twin," and another reporting an independent replication of Kritzinger's result: "I also investigated a large set of publications on DT. Most of them propose digital shadows [...] digital models that receive some measured data, but that do not give any feedback to the physical counterpart" [14]. The dissenting third disagreed definitionally rather than empirically, on the grounds that "there is no single definition of DT in the scientific community."

That finding has aged well, and it recurs domain by domain. In construction and the built environment, existing island solutions "only result in crafting digital models or shadows... which lack the ability to monitor or control (cf. digital shadow) the twinned system" [15]. A federated urban mobility architecture concedes that without stakeholder trust and adoption "the proposed system becomes merely a 'digital shadow' " [16]. Model-driven work on the construction sector identifies the same gap in building information modelling tooling, which "typically provide[s] limited support for bidirectional communication with physical systems" [17]. The distinction has been refined in subsequent conceptual work on digital shadows in manufacturing [18] and on asset-shell communication types, which differentiate model, shadow and twin precisely by their data flows [19]. Some authors now apply the label to their own systems: a friction-identification system for one axis of an industrial robot

is designed, named and reported as a purpose-driven digital shadow, because its job is to help engineers understand the machine rather than to command it [20].

Throughout this survey, where a reported system is a shadow rather than a twin on this criterion, that is stated.

2.3 Axis two: what is twinned — product, process, system

The second axis is the one the term itself obscures. Kritzinger and colleagues note the historical drift explicitly: "Digital Twin in its origin describes mirroring a product, while the state of the art allows processes (manufacturing, power generation etc.) to be as well subjects of virtual space reproduction ('twinning') in order to gain the very same benefits" [12]. The noun changed while the vocabulary did not.

The clearest statement of the resulting typology comes from the NIST manufacturing framework, which treats it not as a taxonomy of assets but as a scoping decision: "the viewpoint of the desired decision or control action determines whether a product, process, or system twin is needed" [21]. This is an important inversion. The type of twin is not a property of the physical thing; it is a consequence of the question being asked of it. The same framework couples this to a discipline that recurs throughout the credible literature — aim for "accuracy not complexity," because avoiding unnecessary components and functionality "saves time and effort and limits potential errors" [21].

The three types are best understood by what decision each serves.

A **product twin** (equivalently an asset twin) mirrors an individual manufactured or engineered artefact across its life, and serves decisions about *that instance*: its remaining life, its configuration, its maintenance. Tao and colleagues' framing of digital-twin-driven product design, manufacturing and service is the canonical statement of the product twin spanning the full lifecycle [22], and the product-lifecycle-management lineage is traced in detail by state-of-the-art surveys of the field [23].

A **process twin** mirrors an activity rather than an object — a production schedule, a chemical process, a clinical pathway, a construction sequence — and serves decisions about *how the activity should run next*. Digital-twin-based dynamic scheduling of cyber-physical production systems is the clearest manufacturing instance [24], and reconfiguration management is its close relative [25], [26].

A **system twin** mirrors an assembly of interacting elements and serves decisions that no component twin can answer alone. Jiang and colleagues give an unusually concrete illustration of the types working together: in a turbine prognostics case, "the product twin and the system twin provided two solutions to the operator for the optimization of the maintenance cost," learning from fifteen years of historical and fleet data, while at a lower level the same architecture instantiates "valve twins, the pump twins, the sensor twins and the tank twins" [7].

To these three, the application literature has effectively added a fourth that does not fit the original scheme: the **human twin**, treated in Section 4.5, where the twinned entity is a patient, a worker or a clinician. It is the type that most strains the definitions in Section 2.1, and the one where the ethical stakes are highest.

2.4 Axis three: scale

Cutting across type is scale. Tao and colleagues propose a three-level ladder and — usefully — tie each level to the class of decision it can support. Digital twins "fall into three types, including the unit-level DT, system-level DT, and SoS-level DT." A unit-level twin is "a minimum but independent individual, which cannot be further divided." For "the monitoring, fault prediction, and maintenance of a single piece of equipment, the unit-level DT is enough"; applications that "need to deal with different units, such as scheduling, progress control, and product quality control" require the system level; and coordination of an entire shop floor requires the system-of-systems level [27]. The same authors' five-dimension model — extending the classic physical-virtual-connection triad with services and twin data — is the most widely adopted structural model in the manufacturing literature [27], with later work extending the dimensional count further for complex industrial systems [28] and addressing multi-scale composition explicitly [29].

Scale is not merely descriptive; it predicts where the engineering difficulty sits. Böttjer and colleagues' systematic review deliberately restricts itself to the unit level precisely because that is "for which real-time control is most applicable," reviewing 96 papers on practical unit-level applications and focusing on "DTs of single production units such as traditional machine tools, additive manufacturing machines and advanced robotic applications" [30]. Their observation that manufacturing twins have been implemented "at different hierarchical levels, ranging from system of systems to unit level," and that no prior review had isolated the level where closed-loop control is actually feasible, is a useful corrective to surveys that treat scale as a detail.

2.5 A fourth reading: capability level

Integration, type and scale describe what a twin *is*. A fourth vocabulary, most developed in the energy sector, describes what it can *do*, and it resolves a good deal of apparent disagreement between papers. Stadtmann and colleagues adopt a six-rung ladder: standalone, descriptive, diagnostic, predictive, prescriptive and autonomous, the last defined simply as "closing the loop" [31]. A system that detects an anomaly is diagnostic; one that recommends an action is prescriptive; only the top rung acts. Aerospace uses a compatible four-term version — descriptive, diagnostic, predictive and prescriptive analytics [32].

The capability ladder is more useful than it first appears, because it separates two things that "digital twin" conflates. A twin can be highly integrated but low capability, streaming bidirectionally while only visualising; or high capability but poorly integrated, producing excellent prescriptions from data loaded by hand. Where the two axes are read together, the survey's central pattern emerges: most reported systems sit at diagnostic capability and shadow integration, while the claims made for them describe the autonomous, fully integrated corner that almost nobody occupies.

An independent formulation reaches the same verdict from the modelling side. Measuring reported systems against a five-dimensional definition, Thelen and colleagues find that "many of the digital twin models reported in the literature are not truly a digital twin model, as they only operate in three or four of the five total dimensions" [33]. That three separate criteria — Kritzinger's data flows, the aerospace requirement of a physical asset [32], and Thelen's five dimensions — independently reach the conclusion that most published digital twins are not digital twins is the single most robust finding in this survey.

2.6 A note on fidelity

One cross-cutting assumption deserves flagging before the sector sections, because it distorts reading of all of them: that a better twin is a more detailed twin. This is genuinely contested in the corpus, not merely under-examined.

The fit-for-purpose position is the majority one and has the strongest institutional backing. The National Academies conclude that "a digital twin should be defined at a level of fidelity and resolution that makes it fit for purpose," which "may lead to the digital twin including high-fidelity, simplified, or surrogate models, as well as a mixture thereof," and note the hard constraint that "in many cases, trusted high-fidelity models will not meet the computational requirements to support digital twin decision-making" [34]. The manufacturing standard adopts the same phrase, defining a twin as "a fit for purpose digital representation" of a manufacturing element [14]. Ali and colleagues state it flattest: "practical use does not always require high fidelity, as it can be costlier... The appropriate fidelity level is not necessarily the highest level of model fidelity feasible and is dependent on the use case" [11]. The NIST framing already pushes the same way with "accuracy not complexity" [21], and one survey names maximalism as an error: "a common misconception about the DT is that the virtual twin should reflect the physical twin in its entirety... these feats are not currently feasible, and certainly not always necessary" [35].

The maximalist position is nonetheless held explicitly. Rasheed and colleagues score twinning on a scale where "a score of 10 will correspond to a situation when a digital twin becomes indistinguishable from its physical counterpart to an observer," and frame the goal as "physical realism without any compromises" [36]. This is not a straw man; it is the intuition behind a great deal of visualisation-led and high-detail modelling work, and it is the assumption most readers arrive with.

Ali and colleagues offer the distinction that dissolves much of the disagreement: "DT fidelity pertains specifically to the suitable precision of the model's representation, [whereas] DT validity encompasses the broader notion of whether the DT serves as a fitting and efficient tool for its intended application" [11]. On that split a twin can be high-fidelity and invalid — precise about the wrong things.

Empirically, the fit-for-purpose camp has the better of it. Bettencourt notes that detailed epidemic simulations during H1N1 and COVID-19 "cast some doubt on the advantages of very detailed simulations, at least when compared to aggregate population models and real-time data assimilation methods," because "large agent-based models are very brittle" [37]. Reviews of civil infrastructure twins agree from the engineering side: "as the number of components increases, IDTs provide broader services; however, they become more complex and less convenient" [38]. And assessment methods are themselves uneven, being "mathematically mature for some classes of models" while "less mature for assessing the fidelity of other classes of models (particularly empirical models) and coupled multiphysics, multi-component systems" [34]. Section 8.5 returns to how fidelity might be measured at runtime rather than asserted at design time [39].

3. Manufacturing and the process industries

3.1 Why manufacturing dominates, and what that distorts

Manufacturing is the field's home sector and remains its centre of gravity by a wide margin. The cross-domain mapping study puts numbers on it: of 356 classified publications, 252 — 70.79% — are manufacturing, with information and communication a distant second at 13.2%, and energy, construction, mining, human health, transport, education and agriculture together making up most of the remainder. For twelve further economic sectors, "we did not find research on Digital Twins" at all [13]. An independent two-part review reports the same skew from a different sample, with "107 out of 230 papers" coming from manufacturing [40]. A comparative study of commercial platforms confirms it on the adoption side, ranking manufacturing "the most prominent, followed by Automotive, Aerospace, and Marine, with Logistics and Agriculture being the least common," and noting that the dominant platforms — Azure Digital Twins, GE Predix, Siemens MindSphere, PTC ThingWorx, 3DEXPERIENCE, Ansys Twin Builder — were largely built for industrial customers [41]. The consequence for a reader is that generic digital twin reference architectures tend to encode manufacturing assumptions: a well-instrumented asset, a cooperative owner, a decision cycle measured in seconds to hours, and a validation regime inherited from engineering rather than from medicine or public policy. Broad reviews of concepts, technologies and industrial applications [42], of smart manufacturing system design [43], and of enabling technologies and tools [44] all inherit this framing, as do reference-model treatments of digital-twin-based manufacturing systems [45].

3.2 Product twins across the lifecycle

The product twin is manufacturing's founding application and remains its most mature. Tao and colleagues' account of digital-twin-driven product design, manufacturing and service with big data set the template for treating design, production and in-service phases as a single information continuum rather than three disconnected model-building exercises [22]. State-of-the-art surveys situate this squarely within the product lifecycle management tradition from which Grieves' original formulation emerged [23], and more recent work on engineering design examines how digital twins are developed and leveraged as design instruments in their own right [46]. A systematic review that sorts 188 manufacturing twin applications by lifecycle phase tempers the picture of maturity: most turned out to be prototypes giving real-time observability of the physical system for optimisation or predictive maintenance, and the review judges deployment in manufacturing to be at an early stage despite how widely the term is used [47].

Service-oriented formulations extend the product twin into additive manufacturing, where the process is difficult to describe with purely physics-based models and the twin's role is partly to reduce the cost of part characterisation and inspection [48]. Digital product passports represent an emerging regulatory-facing variant, implemented via the Asset Administration Shell to carry lifecycle and end-of-life information between supply chain partners [49].

3.3 Process twins: scheduling, reconfiguration, control

The process twin is where manufacturing's closed-loop ambitions are most explicit and most often realised. Villalonga and colleagues present a decision-making framework for dynamic

scheduling of cyber-physical production systems in which the twin performs "automated decision-making… while the production is running," motivated by the observation that the condition of manufacturing equipment "may in fact lead to schedule unfeasibility or inefficiency, thus requiring responsiveness to preserve productivity" [24]. This is a genuine bidirectional application on Kritzinger's criterion: the twin's output changes what the physical system does next.

Reconfiguration is the adjacent problem. A systematic review of reconfiguration management in manufacturing separates allocation, sequencing and parameterisation of the production process, distinguishing scheduling as "a time-wise decision" [25], and framework-level work addresses reconfiguration management concretely [26] alongside digital-twin-enabled reconfigurable modelling for smart manufacturing systems [50]. Visualised architectures for flexible manufacturing systems address the interface between shop-floor protocol layers and twin applications [51], and digital twin services toward smart manufacturing formalise the service layer through which such decisions reach the floor [52]. Two further strands push toward the system scale of Section 2.4. One builds a hierarchy of machine twins that mirrors the structure of the whole plant, so that applications address the hierarchy instead of each device [53]. The other treats reconfiguration as the place where the loop is closed, with the twin supporting a reconfiguration decision that is taken either autonomously or by human stakeholders [54].

A methodological caution worth carrying into other sectors comes from work bridging discrete event simulation and digital twins, which finds that the requirement for continuous synchronisation "hinders the seamless integration" of established simulation practice into twin architectures [55]. Much of what is called a manufacturing digital twin is a simulation model with a data feed attached, and the difference is not cosmetic.

3.4 Predictive maintenance: the sector's most-populated application

If one application dominates the manufacturing twin literature it is predictive maintenance. Van Dinter and colleagues' systematic review of 42 primary studies is the anchor, and its findings map directly onto this survey's axes. By sector, manufacturing leads with 25 studies — subdivided into computer-controlled machine tools, industrial robots and semiconductor fabrication — followed by energy with 10 and aerospace with 9. The aerospace figure comes with a useful methodological caveat, being "related to the fact that NASA provides several datasets in this field" [56]. Sectoral distribution in this literature partly measures data availability rather than industrial activity, which is worth holding in mind wherever this survey reports counts by sector.

By scale, the same review finds component-level twins in 19 studies and system-level in 17, but system-of-systems in exactly one, which "discuss[es] using a system-of-systems Digital Twin for predictive maintenance of a whole shop floor" [56]. Tao's three-level ladder from Section 2.4 is, empirically, a two-level ladder with an aspiration on top.

The review's ranked challenges are computational burden, data variety, and the complexity of assets, data and models, followed by data scarcity and quality, with the authors summarising that "the key challenges are a computational burden, the complexity of data, models, and assets, and the lack of reference architectures and standards," alongside "a low adaptation level of industrial Digital Twin platforms" [56]. One caution for readers: the same paper's conclusion asserts that predictive maintenance with twins "dramatically reduces the number

of maintenance activities and machines' downtime while increasing machine lifetime" without attaching a figure, case or citation. It is an expectation, not a result, and should not be cited as evidence of benefit.

The same group derived a reference architecture for digital-twin-based predictive maintenance [57] and instantiated it for cable joint degradation [58]. Complementary overviews cover the power and process industries [59], and foundation-model approaches to prognostics and health management represent the current methodological frontier [60].

Van Dinter and colleagues also make a typological observation supporting Section 2.3: digital twins "have many variants, such as Twins that simulate components, assets, systems, and processes," and the abstraction level chosen "aims to serve various employees in the business" [56]. Type follows decision, and decision follows role.

3.5 How often does the loop actually close?

Manufacturing is the sector where closed-loop control is most plausible, most claimed, and best measured — which makes it the right place to ask what fraction of reported twins actually command their physical counterpart. Six independent studies answer. They agree on the direction, and only one of them gives a hard count.

Böttjer and colleagues, reviewing 96 unit-level applications where real-time control is most applicable, find that "only a minority of publications established virtual-to-physical feedback and were able to command the physical entity by the virtual entity," that "the majority of publications focused on the physical-to-virtual connection, and few considered the virtual-to-physical connection," and that "entering the control loop and achieving virtual-to-physical feedback was rarely achieved in the reviewed literature" [30]. They also project this forward — such feedback "will continue to be rarely achieved for the foreseeable future" — and identify the causes: closed machine controllers without interfaces, inadequate sensing, the cost of physically adjusting equipment, and distrust of model accuracy. Their claim is qualitative; they publish no percentage.

Van Dinter and colleagues supply the cleanest number. Classifying their 42 predictive-maintenance studies against architectural patterns, they find 30 implement a monitoring pattern, 9 a shadow pattern, and just **two** a control pattern — and note that one of those two never built the decision support it described [56]. Their monitoring pattern is explicitly open-loop: such twins "provide predictive information of the physical object while leaving the decision-making to the machine operators."

A review of additive manufacturing twins appears to point the other way, classing most process-control systems in that field as autonomous, bidirectional twins. It then qualifies the classification in a way that removes the discrepancy: "most current process control only involves basic control signals, such as pause, resume, and alerts. Real-time dynamic control still requires further research" [48]. Counting an alert as closed-loop control gives a very different picture from counting parameter adjustment, and the gap between the two readings is itself a finding about how loosely "closed-loop" is used.

Liu and colleagues reach the same conclusion at larger scale and add a useful device, separating what each paper *claims*, what it *describes*, and what it *actually implements*: "over half of literatures described and studied digital model or digital shadow, although the authors all claimed a digital twin in their papers," and on the strictest reading "most implementations

of the 'claimed digital twin' in literatures are actually digital model or digital shadow" [42]. Since the "over half" applies to the described twin, the actually-implemented fraction is lower still. A lifecycle-oriented review of 188 manufacturing applications reaches the same verdict from a different sample, finding that most were prototypes providing real-time observability rather than control [47].

Finally, practice confirms the literature. A 2025 study at a working automotive engine plant assesses its own result against a four-level capability ladder and reaches only the bottom two, noting that "advancing to higher levels… requires significant technical expertise and innovation." Even the physical-to-virtual direction is manual there: the manufacturer "has opted for manual data retrieval from the PLCs and cloud service," with the authors listing "human in the loop" and "delayed data availability" among the drawbacks [55]. Reconfiguration research reports a similar picture against its own requirements, with only 5% of examined publications scoring in the top quartile and interoperability satisfied by 10% [25].

The convergent answer, then, is that in the sector best placed to close the loop, a small minority of reported twins do so. The only hard count is two control-pattern studies in forty-two, roughly one in twenty [56], and the sources that report larger bidirectional shares do so by counting pause, resume and alert signals as control [48]. The stricter the definition, the smaller the number. This is the quantitative backbone of Section 2.2's argument, and it comes from manufacturing's own reviewers.

3.6 The process industries: a different set of obstacles

The continuous process industries — oil and gas, chemicals, pulp and paper — are usually folded into "manufacturing" and should not be. Perno, Hvam and Haug's review of 79 papers identifies obstacles that discrete manufacturing largely does not face, including system integration problems concerning "the transition from old and outdated legacy equipment and systems," and data quality issues that are "particularly challenging in the process industry" given the nature of continuous production [61]. Their ranked barriers put security and privacy first, followed by lack of system integration, difficulty achieving real-time communication performance, and interoperability; their ranked enablers are led by industrial Internet of Things instrumentation, cloud services, sensing and actuation, and machine learning [61]. Their most telling finding, though, is about the literature itself: of the 79 papers, "only 10… either focused on the process industry or included it within the scope of the study." The authors are candid that this is their principal limitation, since "it is possible that certain enablers and barriers have yet to be identified." Their conclusion that digital twins have "reached maturity" sits oddly against their own composition — 34 of 79 papers are conceptual — and against every other review cited in this section.

Wanasinghe and colleagues' overview of 82 oil and gas articles is the most complete sectoral treatment. They rank applications by prevalence, with asset integrity management, project planning and lifecycle management, and drilling at the top, followed by offshore platform monitoring, pipeline design, and virtual training. They describe twin-based operation as six iterative steps — "create, communicate, aggregate, analyze, insight and act" — while noting that the final step "may be subjected to human intervention (review)," which places most deployments short of autonomy [62]. Their eleven named challenges lead with scope and focus, lack of standardisation, cyber security, and data ownership and sharing. Two structural observations transfer beyond oil and gas: "the majority of the publications

presented theoretical concepts rather than the industrial implementations," and the literature is written overwhelmingly by suppliers rather than operators, with roughly 95% of industrial contributions coming from the supply chain and 5% from operating companies. Adoption itself is "limited to isolated and selective applications instead of industry-wide implementation" [62]. When a technology's literature is produced mainly by those selling it, the enthusiasm asymmetry noted in Section 10.4 has a supply-side explanation as well as a publishing one. System-of-systems framings for digital asset lifecycle management address the specific difficulty that an oil and gas "asset" may comprise several fields and their infrastructure, requiring uncertainty to propagate across constituent systems [63].

Sector-specific instances elsewhere in the process industries include roll-to-roll label printing, where the twin is framed as a dataspace and business-ecosystem problem as much as an engineering one [64], and beer fermentation, where a twin monitors and controls a sampling process through repeated cycles [65].

3.7 Standards: ISO 23247 and its acknowledged hole

Manufacturing is the only sector with a dedicated digital twin framework standard. ISO 23247 has four parts: "(1) overview and general principles, (2) reference architecture, (3) digital representation, and (4) information exchange" [66]. Part 1 supplies the fit-for-purpose definition quoted in Section 2.6. Part 2 gives the reference architecture, organised into observable manufacturing, data collection and device control, core, and user domains, with a cross-system entity providing data translation, assurance and security across them. Part 3 defines the information attributes of manufacturing elements — personnel, equipment, material, process, facility, environment, product and supporting document — of which only the identifier is mandatory. Part 4 covers information exchange across user, service, access and proximity networks [66]. A useful distinction the standard's analysts draw is between a *driven* twin, updated by sensor data from the element, and a *driving* twin, which "describes a plan or an operation to produce a product" [67] — roughly, the two directions of Kritzinger's data flow given separate names.

The NIST analysis is candid about the gap that matters most here: "the current ISO 23247 series provides a framework and guideline for implementing digital twins in manufacturing; however, it does not cover the aspect" of credibility assessment [67]. It names further omissions that recur in every sector of this survey — no digital-thread or lifecycle guidance, since twins "were often developed in silos for one particular functional area"; no guidance on integrating multiple twins; and no component reuse, so that "almost all digital twins are implemented from scratch, which makes implementations time-consuming and costly" [67]. Model-based engineering work adds that the architecture is pitched too abstractly to build from, offering "concepts for components" while "concrete implementation is still a task for a DT developer" [68]. A framework standard exists; a standard for deciding whether a given twin can be trusted does not. Section 8.1 returns to how poorly the standard is actually followed [14].

3.8 Quantified benefit, and its scarcity

Manufacturing produces more quantified benefit claims than any other sector, and reading them carefully is instructive about the genre.

First-hand figures do exist. A shop-floor deployment in chemical fibre textiles reports that a

twin "reduced equipment fault handling time by about 30% and manual operation costs by about 20%" [69]. A flexible manufacturing system reports "saving the site programming and debugging time of PLC by 30%" through virtual commissioning [51]. From the standards side, a named industrial user "used [ISO 23247] to develop a cutting tool life optimization digital twin that helped increase tool life for milling operations by 15%," though this is reported second-hand from pilot projects rather than measured by the authors [67].

The most detailed source is a European programme reporting four small-manufacturer pilots, including a paint and mortar producer whose twin of raw-material and production-order flow replaced a paper process, and a components manufacturer twinning a prosthesis adapter across design, analysis, additive manufacture and test [70]. Reported gains include a roughly 20–30% reduction in the time taken to communicate production order status between departments and a 50% reduction in factory-to-laboratory communication.

Two cautions attach to that source, and they generalise. Its headline claim that a partner "reduced major error by 90%" appears in the executive summary as an achieved result but in the pilot body as a forecast — errors "should be identified as soon as possible, which should reduce major error by 90%." And the summary and the underlying results table transpose which figure applies to cost and which to process iterations. Neither is dishonest; both are the ordinary drift of a number as it moves from a table to a headline, and both are invisible to anyone citing the summary. The report is also unusually honest about the limits of its own measurement, noting that "not all listed KPIs will become fully representative while addressing the digital twin only in the pilot project" and that it "would take at least two production cycles to completely understand the impact" — and, for one pilot, that "there is no direct financial benefit compared to conventional approaches with regard to reduction of effort" [70]. It also records a cost the benefit literature almost never does: a per-unit consumable cost for the tagging that made the twin possible.

Widely circulated sector-level figures are weaker still, being vendor and analyst projections passed between papers rather than measurements — a 10% effectiveness increase attributed to an industry analyst [61], and a claim that companies forgo "35% to 65% of possible value of digital twin investments" attributed to a consultancy [67]. Where this survey reproduces a number, it reproduces the source's framing with it; readers compiling a business case should trace each figure to the pilot that produced it, and check whether it was measured or expected. Comparative analysis of platforms for net-zero manufacturing situates such gains against the high upfront capital costs of deploying digital technology [71]. The general pattern — that benefits are reported per-pilot, rarely with a counterfactual, and rarely audited afterwards — is one of the field's clearest evidential weaknesses and recurs in every sector below.

4. Healthcare, medicine, and the human body

4.1 The sector's structural position

Healthcare is a fast-growing application literature and the one where the gap between publication volume and clinical deployment is widest. It is not, on the cross-domain evidence, the second-largest: the mapping study of Section 3.1 places human health behind manufacturing, information and communication, energy and construction [13]. Katsoulakis and colleagues' scoping review searched PubMed for digital twins in health from 2016 to 2023, retrieving 220 papers of which "85 representative references… form the basis" of their

analysis, and tracked publication growth showing healthcare digital twin output rising from zero in 2016 to 81 in 2023 against 321 for digital twins overall. That is roughly a quarter of the digital twin papers PubMed indexes, a database that over-represents health by design, beginning about three years later [72]. Their verdict on maturity is unambiguous: "although various DT initiatives have been underway in the industry, government, and military, DT4H is still in its early stages" [72].

Five independent sources, the scoping review and the four below, converge on the same lag, which is worth stating plainly because it is one of the few claims in this survey with near-consensus support. Laubenbacher and colleagues write that "today, nearly every industry that deals with complex technologies use sophisticated 'digital twin' computer simulations... The medical analog is only in its infancy" [73]. Viceconti and colleagues note that although "the concept emerged in healthcare more or less concurrently with other industrial sectors... its clinical adoption has proceeded at a much slower pace" [74]. Emmert-Streib and colleagues observe that "the idea of a digital twin originated in manufacturing, and most applications to date are found in that field and in engineering domains" [75]. Niederer and colleagues characterise the whole field as "still in the custom production phase" [76].

The reason for the lag is not funding or enthusiasm. Emmert-Streib and colleagues identify the epistemic asymmetry directly: "digital twin simulation models for manufacturing and engineering are based on descriptions of the mechanical problem and are often related to physical theory. In contrast, our understanding of problems in biology, medicine, and related fields is limited compared with physical theory" [77]. And there is a data asymmetry on top of the theory asymmetry: "unlike engineered systems, which are designed to be measured and controlled, the human body is opaque, fragile, and ethically constrained. One cannot embed sensors into every organ or sample tissues continuously," so "the ideal of real-time, high-resolution data that is central to engineering digital twins is largely unattainable in clinical settings" [75].

4.2 Patient and organ twins

The dominant clinical domains are cardiovascular, oncological, pulmonary and neurological. Katsoulakis and colleagues catalogue cardiac twins including the Living Heart project and vendor cardiac platforms, lung applications including ventilator allocation, spine, Alzheimer's disease work using synthetic control arms, breast imaging, oropharyngeal cancer, and a type 2 diabetes randomised trial [72]. Niederer and colleagues take cardiology as their exemplar, noting that "physics-based computational models of the heart have moved from being a research technique to a clinical tool and are now used in prospective clinical studies" [76]. Outside the organ-model tradition, an early cloud framework for elderly care treats the patient twin as a monitoring service, fed by wearable electrocardiogram, blood-pressure, temperature and pulse sensors and aimed at continuous health management rather than at simulating treatment [78].

Two exemplars illustrate the range of what "patient twin" currently means, and how far apart the two ends are.

Chaudhuri and colleagues present a predictive digital twin for high-grade glioma radiotherapy that is mechanistic and decision-theoretic: a tumour growth model with a linear-quadratic radiation response, calibrated by sequential Bayesian data assimilation from MRI

at days 0, 20 and 27, embedded in a probabilistic graphical model, with treatment regimens selected by risk-aware multi-objective optimisation. The reported results are a median increase in time to progression of around six days at equal total dose, or at equal tumour control "a median reduction in radiation dose by 16.7% (10 Gy) compared to SOC total dose of 60 Gy," rising to a 77.8% median dose reduction for late progressors. The authors state the decisive limitation themselves: the cohort is *in silico*, with growth and response parameters "sampled from literature values and governed by our model," and "a clinical trial would be necessary to establish the benefit of digital twin adapted RT regimens" [79].

Osipov and colleagues' Molecular Twin is a different animal entirely. It integrates ten data modalities — clinical pathology, DNA variants, copy number, RNA expression and fusions, tissue and plasma proteomics, plasma lipidomics and computational pathology — across a resected pancreatic ductal adenocarcinoma cohort, reporting accuracy of 0.85 (95% CI 0.73–0.96) and positive predictive value of 0.87 against 0.59 accuracy for the standard-of-care CA 19-9 biomarker alone, with a parsimonious model retaining performance on 589 features. The authors frame cheap modalities such as computational pathology and plasma proteomics as a route to the "democratization of precision oncology" with "clinical implications in resource-poor geographies," while conceding it is a proof of principle whose external validation degrades substantially [80].

The Molecular Twin is a supervised classifier called a twin. It has no dynamics and simulates no intervention. Whether it belongs in this survey at all is precisely the dispute recorded in Section 4.6.

4.3 Infrastructure rather than models: the Virtual Human Twin

The most architecturally interesting proposal in the healthcare literature is a refusal to build a big twin. Viceconti and colleagues are explicit that "the VHT is not a gigantic digital twin: the VHT is an infrastructure that makes it easier to develop and validate digital twins" — "a distributed and collaborative infrastructure, a collection of technologies and resources (data, models)… and a collection of Standard Operating Procedures that regulate its use" [74]. The motivation is a diagnosis of why the earlier Virtual Physiological Human initiative fell short: despite its systemic ambition, "the initiative produced mainly what is called Digital Twins in Healthcare today, focusing on a single organ system" [74].

The paper's treatment of credibility is the most developed in the corpus and generalises beyond medicine. "The most challenging part for in silico methodologies is the credibility assessment, which aims to quantify the accuracy and reliability of the prediction," with existing standards suggesting a one-off assessment is possible for knowledge-driven models while "for data-driven models, the topic is still being debated." The proposed response extends credibility from models to data, annotating every data object with a credibility level, with new objects entering "at the lowest level of credibility (non-qualified data)" [74]. The authors list seven barriers: lack of advanced models, lack of representative validation data, unclear regulatory pathways, poorly informed stakeholders, poor scalability, workforce gaps, and immature business models — and acknowledge that "the vision in this position paper is still incomplete and, to some extent, immature" [74].

Laubenbacher and colleagues' immune-system roadmap is the comparable effort for a system that resists organ-level decomposition. The difficulty is stated squarely: the challenges stem "from the inherent complexity of the immune system and the difficulty of measuring many

aspects of a patient's immune state in vivo," so "the choice of application is crucial for success." Their proposal is a four-stage, eleven-step workflow explicitly mirroring industrial twin development, including uncertainty quantification as a named step, supported by a federated consortium and reusable model components, at a projected cost and complexity "comparable to the Cancer Moonshot Program" over seven years [73]. Their own caution is that "a medical digital twin will likely never be a finished product."

4.4 Hospitals as process twins

A distinct and under-recognised strand twins clinical *processes* rather than patients, and it maps cleanly onto Section 2.3's process twin. Burattini and colleagues describe an ecosystem of twins for operating room management, structured hierarchically from hospital to operating suite to room to device, with a Surgery twin instantiated at appointment time and linked to a Patient twin through a ten-step perioperative lifecycle, built over an HL7 FHIR server with a rule layer raising inconsistency warnings for missing timestamps and double-booked rooms. Their motivation is that "OR allocation is nowadays scheduled mostly based on the experience of the surgical team," and they note that "to the best of our knowledge, the literature does not propose an application of DTs in the context of ORM." The system is an early prototype without quantitative evaluation, and — importantly for the integration axis — "the current version of the system still requires a human input to label each step with the current time of execution, [so] the real-time acquisition is not guaranteed" [81].

Sieve and colleagues' BedreFlyt twins a hospital ward's bed-bay allocation, combining an ontology-backed knowledge base, actor-based simulation of patient arrivals, and an SMT solver for the allocation itself, under constraints of monitoring requirements, room sharing, gender and contagion isolation. Using anonymised historical data from Oslo University Hospital, an allocation for 100 patients arriving over 30 days "was computed by the DT in approximately 30 seconds," scaling to larger scenarios in under twenty minutes in most configurations. The authors are clear that the ward model "is fixed, i.e. it is not evolving over time, and models only patient trajectories within one ward," and flag a deployment barrier that generalises across healthcare: "there may be practical and legal barriers to feeding the twin with patient data from the ward's data system" [82].

A third variant twins software rather than either patients or processes. EvoCLINICAL is a cyber-cyber digital twin of the Cancer Registry of Norway's automated validation system, using a neural surrogate so that experiments can be run without perturbing the production registry — "querying the target GURI with a large number of unlabeled messages should be avoided since it can interfere with the normal operation" — achieving F1 scores above 0.91 across three system evolution processes [83]. A fourth twins the devices around a care service so that its software can be tested at scale: model-based twins of the medicine dispensers Oslo City issues to residents stand in for the physical units when testing a healthcare IoT application, and behave more than 92% like them on the authors' fidelity measure [84]. Both are reminders that the twinned entity need not be a patient, and need not be physical at all.

4.5 Human twins of workers

The human twin appears in two quite different guises, and the literature does not always distinguish them. In one, the human is genuinely modelled. Villani and colleagues build an Operator Digital Twin whose state class is biometric — electrocardiogram, electrodermal

activity, electroencephalogram and electromyogram signals from wearables — inferring fatigue, discomfort and posture states to modulate collaborative robot behaviour [85]. Azevedo and colleagues twin an expert clinician rather than a patient, modelling a nurse-practitioner preceptor's expertise from think-aloud protocols, eye tracking, facial expression and wrist physiology, motivated by practitioner burnout and attrition [86]. Karnoub and colleagues propose a "digital triplet" in which the human entity is captured through a control glove for prosthetics applications [87], and human-centric AI architectures for Industry 5.0 describe twins "mirroring humans (e.g. to monitor fatigue or emotional status)" [88].

In the other guise, the human is not modelled at all — a *system* is twinned and humans are placed in the loop as roles. Shangguan and colleagues' triple architecture comprises physical objects, digital twins and humans mapped to expected, interpreted and physical worlds [89]; Leirmo and colleagues twin the factory and locate human-centricity in interface and role design, including personalised ergonomic instruction [90]; Toth and colleagues represent the human as an agent class in a collaboration ontology [91]. All seven papers in this subsection use human-centric digital twin vocabulary. Karnoub and colleagues attack the distinction directly, arguing that in a triple human-digital twin architecture "the relationship between humans and DTs is still not close enough, as the DTs only take commands from humans and execute them without replicating human behavior in DT systems" [87]. Wider treatments of digital twins in Industry 5.0 [85], [92] inherit this ambiguity.

The ethical treatment across this cluster is uneven in a way a survey should record. The most developed discussion invokes GDPR, the Data Governance Act and the EU AI Act's risk tiers, alongside secure processing environments, differential privacy and synthetic data — while conceding that "we provide no systemic solution from an architectural point of view" [88]. A dedicated privacy subsection elsewhere acknowledges that "the direct management of raw biometric signals and personal information by the ODT potentially opens to privacy and security critical issues and threats," then defers the problem as out of scope [85]. Other papers in the cluster treat autonomy, privacy and accountability in a sentence of future work [89], [90] or not at all. Continuous biometric monitoring of workers is being architected faster than it is being governed.

4.6 Disagreements the healthcare literature has not resolved

Four disputes are live, and a reader should not be led to believe otherwise.

What counts as a digital twin in medicine. Emmert-Streib and colleagues require mechanism: digital twins "are built on mechanistic models that explicitly represent underlying biological and physiological processes," and the decisive consequence is that "black-box prediction models, due to their abstract and opaque nature, do not support meaningful structural modifications. As a result, the ability to test virtual interventions is a distinctive feature of mechanistic models, which is not present in black-box models" [75]. Viceconti and colleagues agree for the Virtual Human Twin specifically, requiring models that "express an expected causal relation between quantities," which "excludes, for example, all data-driven models" [74]. Against both, a pure multi-omic classifier is named a twin [80] under a definition broad enough to admit it [72]. Niederer and colleagues occupy the middle, allowing statistical, data-driven or mechanistic models "or a combination of these" while insisting that "big data alone are not enough; digital twins must bring together big data, big theory and big computational models" [76]. There is no agreed boundary.

Whether real-time coupling is required. Emmert-Streib and colleagues devote a section to the proposition that digital twins require real-time data [75]. Viceconti and colleagues explicitly reject it as a criterion, observing that "only a minority of digital twins in healthcare described in the literature have a real-time component" and assuming that twins "can or cannot be driven in real-time by sensors, depending on the use case" [74]. This is a direct contradiction, and Section 5.3 shows the same fault line running through civil infrastructure.

Whether the field is mature. Bibliometric maturity and deployment maturity are being conflated. One review infers maturity from a steady publication rate [10], while every source reporting on deployment describes early-stage work [72]–[74], [76].

Where the equity risk lies. Cheap modalities are argued to *reduce* disparity by making precision oncology accessible in resource-poor settings [80]. The scoping review warns the technology may *widen* it: while digital twins can "function as a social equalizer, they can also be a driver for inequality as patterns identified across populations may lead to segmentation such that the DT technology may not be accessible to everyone," and "building human DTs by using biased datasets would exacerbate the existing bias" [72]. Both positions are defensible; the tension is unresolved.

4.7 Why validation is the binding constraint

Niederer and colleagues supply the mechanism behind healthcare's lag rather than merely asserting it, and their diagnosis is the sharpest general statement on digital twin validation anywhere in this corpus: "there are no widely accepted standards for validating a digital twin, nor is there a process that can be applied to determining whether a twin is credible. There are no widely agreed upon tests that should be performed to provide estimates of numerical accuracy and stability, nor uncertainty quantification and propagation" [76]. They add the practical constraints — clinical studies "contain ten or fewer patients, often requiring computing and analysis facilities unavailable in a hospital," against a clinic in which "a new patient may be seen every 10–15 minutes" — and note that open-source platforms of the kind that industrialised finite element analysis and machine learning "are not widely available for digital-twin applications" [76]. Their governing analogy is that the finite element method took roughly seventy years to move from expert practice to commodity tool, via mathematical foundations, error analysis, accessible computing and standard software. Digital twins have had none of that time.

5. Society, cities, and the built environment

5.1 The conceptual objection

Urban digital twins are where this survey's evidence is most polarised, and the polarity is instructive.

Bettencourt provides the strongest reasoned objection, and it is conceptual rather than technical. His verdict is that "urban digital twins remain, at this point, shallow explanations of urban processes despite their great level of representational detail and growing data assimilation power" [37]. The core asymmetry is captured in a single line: "while digital twins disaggregate, cities aggregate." Cities' "most important properties are emergent, arising from many interactions over time, which are extremely difficult to compute from

the bottom up" — so the very move that makes an engineering twin more faithful makes an urban twin less explanatory.

He formalises this as a "planner's problem," explicitly echoing an older argument in economics: the computational complexity of urban environments makes "optimal decisions and planning provably impossible," because the strategy space "would grow extremely fast, typically faster than exponentially," so "any actual set of simulations becomes a vanishing, unrepresentative fraction of all possibilities as time goes by" [37]. From this follows a normative warning that no engineering framing anticipates: cities are "predicated on diversity and variation," which is "functional — a necessary source of innovation, creativity, and resilience," so using urban twins "for homogenizing urban environments in terms of human behavior or physical spaces… would, in effect, kill a city."

He also states the scope condition under which twins *do* work, and it is worth quoting because it explains this survey's whole sectoral pattern: successful applications "have in common a relatively short time horizon, the absence of significant behavioral change, and relatively well posed objectives. Such conditions are not typical of cities" [37]. That sentence explains why manufacturing succeeded first and why medicine and urban policy have not.

Two further points. Verification and validation are "challenging in cities because most computational models are designed for (some sort of) average prediction," which "blurs the line between anomaly detection (and elimination) and functional surprise." And the observational apparatus itself "raises concerns about surveillance and privacy." Notably, Bettencourt does not make the reflexivity argument that citizens change behaviour in response to the model's own predictions; that argument does not appear in this corpus and should not be attributed to him.

5.2 Governance, participation and the digital divide

The counterweight to the conceptual critique is practitioner evidence about institutions. Alexandridis and LaFontaine, writing from inside Orange County Public Works, describe the Orange County Digital Twin: a reality-capture workflow, BIM models "linked and integrated into real-time relational asset management workflows through object identification and standardization," a 3D visualisation engine, and an institutional anchor in the county's seven-year Capital Improvement Plan, whose public site gives citizens geospatial and cost data years ahead of project start [93].

Their governance findings are the most useful in the corpus because they are neither boosterish nor abstract. Governance must be designed in, not retrofitted: "governing the digital twin must be considered in the design phase of its development or important components such as inclusiveness, equality, empowerment, and citizen participation could be excluded." And participation alone is insufficient: "often, citizen participation and inclusion may not simply be enough to initiate and promote realized public benefits," citing "the decoupling of participation from democratic decision making for critical climate-related infrastructure" [93]. Equity is analysed with census-tract data rather than asserted, and the consequence is framed institutionally — sharp demographic and economic divisions bear on "the democratic and the representative legitimacy of institutional actions." The digital divide cuts both ways: visualisation tools "can potentially create a digital divide" across generations, yet immersive interfaces may "reduce disparities in the ways that users and communities perceive digital data." Their stated limitation is that "a systematic or standardized framework

approach that includes widely accepted principles, computer and engineering standards and protocols is yet to emerge."

Broader governance framings identify data privacy and security, ethics and bias, interoperability and standardisation, intellectual property, and accountability as the five dimensions requiring institutional treatment, and describe national-scale efforts including a national digital twin roadmap and proposed coordinating authority [94]. Work on smart-city twin development identifies fragmentation as the core structural problem — such twins "continue to be fragmented in terms of scale, scope and location and deployed on a case-by-case basis with little interaction between the different systems in use" — and itemises risks including privacy, cybersecurity, cost, skills, data quality, the practical difficulty of scaling, and over-reliance, which "could also render city operations highly vulnerable to system failures" [95].

5.3 Civil infrastructure: maturity measured honestly

Naderi and Shojaei's review of civil infrastructure digital twins covers 85 studies from 2012 to 2022 and is unusual in offering a calibrated maturity verdict rather than an impression. Their finding is that infrastructure twins "are in their early stages and far from complete adoption in the industry," with current work concentrated on frameworks for a particular service, mostly in the operations and maintenance phase of bridges and roads [38]. Technology frequencies are informative: building information modelling (BIM) appears in 53% of studies, geographic information systems 19%, Internet of Things instrumentation 42%, LiDAR and point clouds 23%.

Two of their findings bear directly on this survey's axes. On the BIM confusion: "many researchers refer to DTs interchangeably with BIM. Others consider DT as a simple 3D representation of their assets," and the discriminating criterion they offer is Kritzinger's in different words — "effective data communication between twinning technologies can distinguish a real IDT model from monolithic 3D models" [38]. On standards: "most efforts for creating digital twin standards are proprietary, and a truly open standard is yet to be developed," with hopes pinned on forthcoming schema versions while most current implementations extend silo-based predecessors.

Torzoni and colleagues provide the most methodologically complete civil twin: an asset-twin system encoded as a probabilistic graphical model, with "the time-repeating observations-to-decisions flow… modeled using a dynamic Bayesian network," diagnostics from deep learning on raw vibration data, digital state updated by sequential Bayesian inference, and a reduced-order model used offline to generate training data and precompute a health-dependent control policy. The validation case is a real named structure, an integral concrete portal-frame railway bridge on a Swedish line, under realistic train loading. The framework "promptly suggest[s] the appropriate control input, within at most two time steps of when the (unknown) ground truth structural health demands it." The authors state the limitation that matters: both cases are synthetic — "in the absence of experimental data, the tests have been carried out considering high-fidelity simulation data, corrupted with an additive Gaussian noise," and the reduced-order model exists precisely "to overcome the lack of experimental data for civil applications" [96]. Their earlier multi-fidelity surrogate work for Bayesian model updating is the methodological precursor and is equally explicit that experimental validation remained future work [97].

Chacón and colleagues supply the complementary case: real measurements without a closed loop. Their work covers load testing across nineteen viaducts spanning roughly eight kilometres on a Spanish high-speed railway branch, with accelerometers, displacement transducers and strain gauges, modal identification, and measured-to-predicted frequency ratios reported per mode. Their conceptual contribution is a direct challenge to real-time orthodoxy: new measurements integrated into a bridge twin "are not needed on a 'real-time' basis. Rather, this information is needed on a 'right-time' basis, which may take days, weeks or even years. As a result, the challenge for DT in bridges would rather be its open integration for various stakeholders rather than the development of high-fidelity, low computational power, and real-time numerical advanced models" [98]. They also make an economical argument worth noting: load tests on new bridges "present a good opportunity to establish primeval yet useful digital twins of these type of structures with insignificant additional operative costs."

The same group's requirements analysis, grounded in ten named European demonstration sites spanning bridges, buildings, an airport runway, a stadium roof and a quay wall, states the sector's position with unusual bluntness: "technology is ready for isolated measurements but not for integrated systems." They identify the absence of a single pipeline from structural health monitoring into twins, procurement as an under-recognised barrier — "interoperable systems need to be properly tendered during procurement" — and the sheer coverage problem that "the vast majority of assets do not have their own virtual twin" [99].

5.4 Construction: what separates BIM from a twin

The clearest statement of the BIM-versus-twin distinction has three parts: "first, BIM models are predominantly static representations focused on design and construction phases, lacking mechanisms for real-time data integration… Second, BIM tools typically provide limited support for bidirectional communication with physical systems… Third… it lacks native capabilities for modeling runtime aspects of Building Control Systems," creating "a fundamental gap between BIM's design-time focus and the operational requirements of DTs" [17]. The same group's service platform generates twins automatically from industry-standard building models via semantic transformation into asset administration shells, and criticises general-purpose commercial modelling languages on the grounds that they yield "information loss as these languages are not tailored to the specific needs and requirements" of the architecture, engineering, construction and operations (AECO) sector [15].

Two observations from practice sharpen the sectoral picture. First, an emerging role: analysis of twinning an office building identifies four team leaders including "the emerging position of the DT Manager," noting that "no prior literature has been identified regarding the DTM position within the AECO industry," and that current BIM practice "do[es] not yet include an active connection between the physical asset and virtual models" — the twin being "the living version of a BIM model" [100]. Computational design tooling to lower the barrier to producing twin-ready models addresses the same skills constraint [101]. Cognitive digital twin models for building lifecycle management represent the knowledge-graph-centric end of the same effort [102]. Software-engineering practice is arriving in the sector as well: a DevOps framework for built-asset twins, which calls the domain's twins "relatively nascent", sets out a microservice architecture and continuous-delivery pipeline and demonstrates them on a single climatic chamber rather than on a building [103]. At campus scale, a preliminary twin of 175 university buildings brings live, historical and forecast energy and

air-quality data into one platform, mainly for building energy use [104].

Second, an honest failure report. A BIM-and-twin monitoring system for a solar-powered noise barrier tunnel in South Korea — motivated by a fatal tunnel fire — reports that one air quality sensor lost 30.65% of records, a second "did not transmit any data," and a water detection sensor "did not transmit at all," with intermittent power problems causing data loss, concluding that "the current system is deficient in the realm of real-time recovery and buffering mechanisms" [105]. Such reporting is rare and disproportionately valuable: the field's reliability claims are almost never tested against sensor attrition of this magnitude.

The most striking deflationary datum in the built-environment literature comes from a practitioner speaking against commercial interest — the head of a firm selling a twin platform, who estimates that "there are fewer than 50 true building twins in the world," adding that "we are at the beginning of the DT journey in the built environment" and "we're very much in the early adopter/innovator stage, so it's a leap of faith" [106].

5.5 Mobility

Mobility twins sit at the opposite end of the latency spectrum from bridges. Wang and colleagues define a Mobility Digital Twin as "an AI-based data-driven cloud–edge–device framework for mobility services," twinning three physical elements — human, vehicle and traffic — across four implementation layers. Their case study, personalised adaptive cruise control on an instrumented production vehicle over a real highway trip, reports takeover events occurring "86.4% less frequently" with the learned controller across three drivers, with median end-to-end latency of 80 ms for light microservices but a cloud compute path requiring a sixteen-second warm-up. Their stated limitations include that "there is no universal definition of this technology, let alone existing standardization" [107]. Industry agrees: a questionnaire of practitioners in the German automotive industry starts from the observation that automotive has no common understanding of digital twins and no established techniques for modelling them [108].

At the urban scale, a federation architecture for low-car transformations proposes four pillars and six functional modules across macroscopic, mesoscopic and microscopic model layers for an Amsterdam district, developed through five stakeholder co-design workshops — but reports no performance results and describes itself as "a high-level vision" resting on the assumption that federated modules can be integrated effectively [16]. Work on autonomous vehicles and built-environment change delivers the sector's most quotable maturity verdict alongside three specific gaps: no reviewed study used real-time data, none used machine learning methods, and many "did not explicitly validate their results," so that "from a real-time modeling and simulation perspective, the current state-of-the-art remained conceptual" [109]. Two papers address the infrastructure layer beneath such applications: an edge–cloud smart-city architecture that places latency-sensitive mobility and road-safety services at the edge [110], and fluid twins that migrate computation across the edge–cloud continuum, prototyped and evaluated in the Modena Automotive Smart Area testbed [111].

5.6 Utilities, supply chains and the wider economy

Two applications extend the societal sector beyond cities proper.

Water distribution provides one of the few genuinely operational municipal twins: a digital twin of the water distribution network of a Spanish city and its metropolitan area serving 1.6 million inhabitants, described as "one of the first DTs built of a water utility, being currently in operation," built on an established hydraulic simulator and applied to leak localisation, energy efficiency, water quality and maintenance planning [112]. It is a system twin in Section 2.3's terms, and it is in production.

Supply chains represent a different use of the concept, closer to scenario analysis than to mirroring. Digital twins here "allow managers to create alternate realities, and explore vulnerabilities that exist in their networks," described as "digital enablement of traditional 'wargaming' " against disruption drivers including climate-related shipping constraints and supply chain intrusion [113]. Buildings-as-batteries work applies twins to distributed energy participation across a multi-city network, while conceding that "there is still insufficient data to capture the success or struggles" of the approach and that many such projects "are still being applied in small scale, conceptual, or test environments" [114].

6. Energy, aerospace and maritime engineering

6.1 Wind energy: the corpus's strongest validation

Energy is the second-most-populated domain in predictive maintenance work [56], and offshore wind supplies this survey's single best-validated twin.

Branlard and colleagues implement, verify and validate a physics-based digital twin for a floating offshore wind turbine "using measurement data from the full-scale TetraSpar prototype," estimating aerodynamic loads, wind speed and tower section loads "with the aim of estimating the fatigue lifetime of the tower" from commodity sensors alone — power, pitch, rotor speed, tower acceleration, inclinometers and GPS [115]. Full-scale validation against a real deployed structure is rare enough across every sector in this survey that it is worth flagging as an exemplar of the standard.

What makes it exemplary, though, is less the validation than the honesty of the error reporting, which is worth setting out because it calibrates every other claim in this survey. In numerical experiments the twin recovers damage equivalent loads for tower fore-aft bending "with an accuracy of approximately 5 % to 10 %"; against the physical prototype "the errors increased to 10 %–15 % on average" [115]. Underneath those averages sits a wide distribution: across 576 ten-minute samples the best case reaches 0.4% error, yet "the tower damage equivalent loads are, on average, within ±10 % of the values obtained from the measurements, but some cases show errors ranging between ±50 %." The authors name what fails and why — "our current method fails at capturing the high-frequency content of the signals, which can have a significant impact on the accuracy of the damage equivalent loads," the twin "lacks sufficient information to fully capture the tower-top loads," and the quasi-steady aerodynamic estimator "cannot fully capture the dynamic aerodynamic state of the rotor in floating conditions." They also concede that the numerical error levels are "a best-case scenario because we assumed that no noise or biases were present in the measurements," and that structural health diagnostics and maintenance decisions "are postponed to future work" [115]. It is, precisely, a load-estimation twin rather than a decision twin — a distinction most papers blur.

Stadtmann and colleagues supply the field's clearest capability vocabulary alongside a

genuine operational finding. Their six-rung ladder — standalone, descriptive, diagnostic, predictive, prescriptive and autonomous, the last defined as "closing the loop" — places their own system at rung two: "a digital twin equipped with such capabilities is called a diagnostic digital twin," able to "assess asset conditions and detect anomalies" [31]. Applied to a year of one-minute operational data from a full-scale floating turbine, normal-operation neural models flagged a generator temperature anomaly after which, "7.8 hours after the detected anomaly, the turbine stopped for 4.2 days" — a coincidence they put at under 0.2% probability by chance, and, critically, "a temperature anomaly that had not been identified in the commercial system of the asset." Their methodological point is the one this survey keeps returning to: "in contrast to most existing research on floating offshore anomaly detection, which relies on turbine simulations, this work uses data from an operational floating wind turbine" [31]. They are equally clear about the limits — validation "has been limited to one year of historical data from a single turbine," and an alternative recurrent model "was not able to clearly identify the anomaly."

Two adjacent contributions should be read for what they are rather than what a keyword search suggests. Population-level wind farm monitoring argues that "turbine groups with sparse data (automatically) borrow statistical strength from those that are data-rich," a genuinely useful idea for fleets — but the work frames itself in terms of digital *representations* rather than twins, and verifies its central concept "on a simulated blade case study" [116]. Blade research and demonstration platforms are not twins at all but the physical reference artefacts a twin would need, motivated by the observation that "the existing reference wind turbine concepts are virtual wind turbines without measurement data to compare with" [117].

6.2 Aerospace: where the concept was formalised

Aerospace gave the digital twin its first institutional definition, and the AIAA and AIA joint position paper remains the reference statement: "a set of virtual information constructs that mimics the structure, context and behavior of an individual / unique physical asset, or a group of physical assets, is dynamically updated with data from its physical twin throughout its life cycle and informs decisions that realize value" [32]. Its most useful contribution to this survey is a boundary condition stated more sharply than anywhere else: "the essential elements of a Digital Twin are a virtual representation (model), a physical realization (asset), and a transfer of data / information (connected) between the two. Hence to have a Digital Twin requires a physical asset." On that criterion a large share of the papers surveyed here — architectural proposals, prospective case studies, simulation-only frameworks — are not digital twins at all. The paper's barriers are grouped as business and transactional, where "the value of a Digital Twin is still not clearly understood or articulated"; technical and analytical, including standards, since "the majority of these commercial sector implementations are proprietary," and "verification/validation/accreditation, certification and uncertainty quantification"; and cultural, including a workforce spanning disciplines "rarely taught within the same academic curriculum," and the regulatory obstacle that for certification by analysis "the current Regulatory requirements, policies and landscape do not allow for full realization of this benefit" [32].

Kapteyn, Pretorius and Willcox provide the field's most cited mathematical foundation for predictive twins. The formalism is precise: "the model presented here is a dynamic decision network: a dynamic Bayesian network with the addition of decision nodes," relating

physical state, digital state, control input, observational data, quantity of interest and reward. The asset is a fixed-wing unmanned aerial vehicle with purpose-built composite wings, calibrated in three stages — as-manufactured geometry, a static tip-load and displacement test updating the material stiffness parameter, and an initial-condition response test fitting point masses and damping to the first two bending modes [118]. What the twin decides is narrow and concrete: a binary in-flight manoeuvre choice between a 2g and a 3g load factor, informed by strain at 24 gauges near two defect regions, yielding a policy to "fly the more aggressive 3g maneuver until $z_1 \geq 60$, at which point it should fall back to the more conservative 2g maneuver."

The scoping is worth reading carefully, because it recurs across the corpus: calibration uses real experimental data collected in a laboratory with the wings detached, while the mission in which decisions are actually taken is simulated. This split — real calibration, simulated decisions — is easy to miss and common. The authors name the deeper obstacle themselves: "limitations of the proposed framework include the challenge of defining and parameterizing the models comprising the digital twin. A central aspect of this challenge is a need to quantify and manage model inadequacy" [118]. The group's earlier work on the same fixed-wing aircraft shows how the asset twin is meant to scale. A library of component-based reduced-order structural models is meant to be shared across many assets, and a classifier trained on sensor data picks which library models make up each aircraft's updated twin, so that the vehicle can replan its mission after damage or degradation [119].

Thelen and colleagues' two-part review is the most thorough treatment of modelling and twinning enabling technologies [40] and of uncertainty quantification and optimisation, anchored by a battery twin case study [33]. Part 1 also supplies this survey's hardest number on sectoral concentration: of 230 papers reviewed, "107 out of 230 papers come from the manufacturing domain" [40]. Part 2's battery twin determines "the optimal time to retire a Li-ion cell from its first life application in advance," using a particle filter over a capacity-fade model feeding a utility optimisation — though the authors note the twin "mostly underestimates the RUL of the cells," and that although the particle filter yields a probabilistic prediction, "this case study used a deterministic method to determine the optimal time to retire a cell" [33]. Related probabilistic machine learning work on battery health is a review rather than a twin, and warns that laboratory accuracy figures "are unlikely to generalize to real-world applications" [120]; a companion augmented model-based framework for remaining-useful-life prediction, validated across 237 real cells, does not use the digital twin label at all and is better read as prognostics [121].

Two claims from Part 2 deserve to be lifted out of the battery context, because they are the most useful general statements in this survey about what a digital twin is and how it can be checked. On the first: measuring reported systems against a five-dimensional definition, the authors conclude that "many of the digital twin models reported in the literature are not truly a digital twin model, as they only operate in three or four of the five total dimensions," and more pointedly that "many of the research papers we reviewed which claim to have created a digital twin have actually proposed a prognostic model for predicting the RUL of an engineered system. However, a prognostic model on its own is not fully a digital twin" [33]. This is Kritzinger's finding arriving independently, from a different discipline and a different criterion, four years later.

On the second — and this is the sharpest explanation anywhere in the corpus of *why* twin validation is structurally hard rather than merely neglected: "validation of a digital twin

is not the same as validation of a general purpose prediction model," because "the model updated at one time instant can only be validated with future data; thus validation is a lagging indicator of digital twin model quality" [33]. A twin that updates itself has no fixed artefact to certify. That is a different problem from the one existing verification and validation practice was built to solve, and it explains why every sector in this survey reports the same gap.

6.3 Maritime and offshore

Maritime applications combine the structural concerns of civil engineering with the operational concerns of manufacturing, and the sector supplies an unusually clean illustration of the integration axis. A shipboard crane twin on a research vessel is a co-simulation integrating a six-degree-of-freedom vessel model with crane multibody dynamics and a pendulum payload model, and its declared purpose stops short of control: it exists to "allow the operators to practice demanding operations in a risk-free immersive environment" and to "identify the optimal operating conditions, predict potential risks, and provide onboard support to the operator" [122]. The only quantified agreement reported is a roll comparison against the vessel's motion reference unit differing "around 0.2 degrees," from a single day's sea state.

The same vessel appears in a cross-domain case-study collection that states the integration level explicitly: "the PT/DT connection is currently running in an 'open-loop' manner, where no automatic actions have been implemented yet," making it, by the authors' own classification, a digital shadow [123]. That collection's broader methodological complaint is one this survey has felt throughout: "case studies in experience reports often do not fully report on all relevant details. This not only obscures the DT engineering process, but also does not allow the complexity of the DT-enabled system to be revealed."

Reviews of structural health monitoring for offshore and marine structures situate twins as a middle path rather than a replacement — "the digital twin approach stands in between these two techniques as a 'grey-box model' " — and give the economic motive for the whole cluster: "more than 50% of the installed offshore and marine structures in the Norwegian, UK and Gulf of Mexico shelves exceed their design life" [124]. Industrial practice from a major offshore consultancy builds the twin as a Bayesian-updated finite element model, arguing that once analytical and experimental modal parameters agree, "this is the crucial step in generating a True Digital Twin of the real structure" [125]. That claim is contested within this very corpus, with the monitoring review observing that the term "true" was used "because they considered the platform's original model as a digital twin" [124]. The industrial work is also candid that its own validation is unquantified: "the FEMU of the prediction models results in improvement of the prediction models, but the question remains: by how much?" [125]. A deployed twin across many real platforms, with no error metric — the mirror image of the academic pattern of precise error metrics on synthetic assets.

6.4 What the heavy-engineering sectors share

Three features distinguish this cluster from manufacturing proper and explain its methodological character. The assets are expensive, long-lived and individually instrumented, which makes per-asset product twins economically rational in a way they are not for mass-produced goods. Failure consequences are severe, which forces uncertainty quantification

from a nicety into a requirement. And the physics is well understood, which means mechanistic models genuinely carry predictive weight — the opposite of the situation in medicine described in Section 4.1. Where those three conditions hold, the digital twin concept works close to as advertised. Where any one fails, the sections above show what happens.

7. Agriculture, food, environment and science

7.1 Living systems as twinned assets

Agriculture occupies an interesting middle position: biological uncertainty comparable to medicine, but without medicine's ethical and regulatory constraints on instrumentation, which makes it a natural proving ground.

Livestock twins are the clearest instance. IUMENTA is "a generic framework for animal digital twins," which models **energy balance** — metabolic rate, heat production, oxygen consumption — through random-forest soft sensors fed by a wearable combining heat flux, skin temperature and accelerometry. Its three demonstrations all use real animal data: pigs in a respiration chamber held below, at and above thermoneutral temperature, with early-life measurements of an individual pig predicting its own later energy expenditure; salmon oxygen consumption inferred from accelerometry in swim tunnels; and shellfish respiration from shell opening and heart rate. The authors describe it as a proof of concept, note that "whereas our proof-of-concept ADTs allow automated inference, we have not yet handled the integration of larger groups of animals," and report no error metrics in text [126]. It is also not closed-loop: the path runs animal to sensor to model to researcher, with no actuation.

GreenhouseDT is the counter-example, and one of the very few genuinely closed-loop twins in this entire corpus. Running on commodity hardware with moisture, temperature and light sensors and **water pumps as actuators**, "a decision procedure is executed for each plant to determine whether, and for how long, the pump should be activated," with the simulation driver relaying decisions to the actuators [127]. Its contribution as an exemplar is to separate *behavioural* self-adaptation — tuning pump frequency and duration — from *architectural* self-adaptation, meaning reaction to changes in what the physical system structurally *is* as plants are added, moved or replaced. An ontology-backed asset model in a knowledge graph is checked against running program state by "defect queries" that detect drift and trigger repair. A model-based control extension exports a physical model as a functional mock-up unit and detects model drift automatically: "if the difference is too big (i.e., the model drifted too far), the model is automatically reset with the last, correct sensor value." The honest limitation is that the asset model does not update itself — "we let a human operator — generally an engineer but, for GreenhouseDT, a gardener — perform these update operations" [127]. Companion work on declarative lifecycle management adds a two-layer adaptation loop in which "a feedback loop for architectural self-adaptation is used to reconfigure the twin's behavioral feedback loop," evaluated on one case study and a synthetic benchmark rather than in the field [128].

Tekinerdogan and Verdouw's architecture pattern catalogue is the sector's most transferable contribution, and it makes the integration axis structural. Nine patterns are named — Digital Model, Digital Generator, Digital Shadow, Digital Matching, Digital Proxy, Digital Restoration, Digital Monitor, Digital Control and Digital Autonomy — and the authors explicitly partition them: the catalogue covers patterns that "do not include two-way syn-

chronization… i.e., a Digital Model, Digital Generator, Digital Shadow, and Digital Proxy pattern," whereas "the 'real' digital twin patterns that we identified include a Digital Monitor, Digital Control, and Digital Autonomy" [129]. Their three agri-food cases cover arable within-field management zoning, dairy activity sensing for heat and health detection, and chain-integrated greenhouse production — but none was built: the cases "were prospective cases, that is, they included the system that was planned to be developed," and no case used the Digital Autonomy pattern, "indeed the most difficult pattern" [129].

Food and beverage processing supplies a rare fully bidirectional process twin. A pressurised beer-fermentation sampling rig is twinned explicitly as an interactive twin rather than a shadow, described as "the critical transformation from a digital shadow to an interactive Type 2 Digital Twin." Reported outcomes are operational rather than model-theoretic: manual sampling fell "from 45 minutes per batch per day to 4 minutes, representing a 91% reduction," while sampling frequency rose from eight manual readings per day to 17,280 automated data points, which "revealed fermentation patterns that were previously invisible to manual monitoring approaches," across "over 500 pressure cycles without incident." The authors' own framing of the lesson is generalisable: a "progressive implementation strategy evolving from passive monitoring to active control provided validation of foundational components before introducing bidirectional control complexity" [65]. Their stated limits are cost, scalability across multiple vessels, and sensor coverage that omits foam dynamics and yeast viability.

7.2 Environmental and Earth systems

Environmental twins scale the concept past any single owner. Work advancing a marine digital twin platform models a Mediterranean coastal lagoon *and its surrounding catchment basin*, integrating public sensor networks, buoys, satellite imagery and weather data, with machine-learning forecasts reported by normalised error across horizons — sub-hourly streamflow errors improving markedly once weather forecasts are included, and salinity and oxygen forecast errors below one percent at weekly horizons [130]. The authors' conceptual complaint generalises well beyond marine science: "most DT-related studies often focus on a single use case, with models and data formats that are often not interoperable," and "the methodological aspects of a DT approach to the study of the Earth system or its sub-components have not yet been established" [130]. Their stated limitations include unavailable public data for key nutrients, and — decisively for the integration axis — physical-system integration is future work. It is a sophisticated monitoring and forecasting system, not yet a bidirectional twin.

At the largest scale, the interTwin project co-designed and implemented "an interdisciplinary Digital Twin Engine, an open source platform that provides generic and domain-specific software components for modelling and simulation to integrate application-specific Digital Twins," built on a blueprint architecture guided by open standards, with co-design driven by use cases "from high energy physics, radio astronomy, astroparticle physics, climate research, and environmental monitoring" [131]. The scientific-computing framing is a genuine fourth category alongside product, process and system twins: the twinned entity is a natural phenomenon that no one owns, operates, or can actuate.

8. Cross-cutting concerns

8.1 Standards and interoperability

Every sector in this survey reports interoperability as a barrier, and each has responded with a different standard. The result is not an absence of standards but a surplus of mutually incompatible ones.

Schmidt and colleagues give the most concrete diagnosis. Comparing the Asset Administration Shell, the Digital Twins Definition Language, Web of Things descriptions, NGSI-LD and others, they find "these existing standards are incompatible with each other, i.e., they do not have the same (i) syntax, (ii) mechanisms for representing properties and behavior, (iii) communication mechanisms and languages or (iv) semantics of properties and behavior," adding that interoperability mechanisms "are typically proprietary within the standards" and that "it cannot be expected that in the future there will be only one dominant international standard" [132]. Their attempt to transform one representation into another surfaces failures that are structural rather than incidental: array and map datatypes "cannot be transformed as is, because they cannot be mapped into the AAS using a generic approach," since entries added at runtime "would lead to a structural change in the AAS." Semantic annotation is lost in transit, and round-tripping is broken outright — a reverse transformation "is currently not implemented," because "two successive transformations [do] not produce the same element" [132]. Two twins can conform to standards, exchange data, and still not mean the same thing by it. Work on rationalising asset-shell communication types addresses the same fragmentation from within one standard [19].

Manufacturing's ISO 23247 fares no better on adoption than on credibility. Ferko and colleagues audited the standard against 29 published architectures, drawn from 140 studies, plus practitioner survey responses and expert interviews, motivated by having "noticed that standards, and in particular the ISO 23247 standard, are not completely followed" [14]. Three of the standard's functional entities — plug and play, peer interface, and data assurance — "are not implemented by current DT architectures" at all. Three functionalities practitioners need are missing from the standard: data storage, twin versioning and continuous deployment, with data storage appearing in 69% of architectures that the standard gives them no element for. Their verdict is that adoption "is still in its embryonic stages," with an expert suggesting the standard "may have gone further compared to the current DT maturity stage" [14].

Beyond manufacturing, systematic interoperability frameworks argue that alignment must happen at the metamodel level — "when all systems share a common metamodel for encapsulating internal behaviors, capabilities, and purpose, they become inherently interoperable" — while conceding that the authors "may not have contemplated all permutations of system interoperability" [133]. Architecting studies over 140 primary studies find the layered pattern dominant at 35.9%, followed by service-oriented at 29.3%, with maintainability the most-cited quality attribute — and note that scalability, reconfigurability and extensibility matter to practitioners despite sitting outside the relevant quality standard [134]. Their observation about the field's incentives is worth carrying: a lack of consensus "may bring experts in keeping proposing novel solutions rather than working on consolidating existing ones" [134]. Unifying reference models [135], discipline-specific characterisation [136], semantic and cognitive modelling [137] and federation roadmaps [138] all attempt convergence, the last naming proprietary formats and absent governance frameworks as the

obstacles.

The pattern across sectors is consistent, and it is the same pattern Section 3.7 found in manufacturing alone. Standards exist for structure and communication; they do not exist for credibility. The National Academies state the consequence directly: "the absence of standardized quality assurance frameworks makes it difficult to compare and validate results across different organizations and systems," and "the lack of adopted standards in data generation hinders the interoperability of data required for digital twins" [34]. A practitioner can determine whether their twin conforms to an architecture. They cannot determine whether it should be believed.

8.2 Platforms, services and composition

The platform layer has matured faster than the validation layer. A survey of Digital Twin-as-a-Service platforms proposes a component-based taxonomy and identifies the three functions such platforms are expected to deliver: "real-time synchronisation between physical and digital entities," "emulation of system behaviours under varying conditions," and "closed-loop feedback mechanisms to control physical operations" [139]. Open-source tooling has been surveyed bottom-up, analysing "14 open-source DT frameworks in 10 different dimensions" grouped into six categories [140], with specific frameworks including compositional open-source platforms [141], [142]. An earlier Industry 4.0 reference model for twin-as-a-service builds one of its three axes from four integration levels derived from Kritzinger's, and observes that most studies in the field had focused only on manufacturing [143]. Reference architectures for twin software platforms address the industrial end [69], and evaluation frameworks for domain-specific platforms attempt comparability [144].

Composition and federation are the current frontier. Work on integration challenges for digital twin systems-of-systems identifies the difficulty of composing independently developed twins [145], with subsequent analyses of the challenges of integrating digital twins [146], roadmaps toward federation [138], and architectures for composite twins enabling collaborative ecosystems [147]. Semantic and cognitive approaches — knowledge graphs, ontologies and reasoning layers — are the main technical route being pursued [137], [148], [149].

8.3 Networks as twinned objects and as substrate

Telecommunications occupies a doubly interesting position, being both a sector that builds twins and the substrate on which other sectors' twins run. It also supplies a cautionary case in terminology, because "digital twin network" denotes at least three different objects in this corpus, and papers using the phrase frequently do not mean the same thing.

In the first sense the network *is* the twinned asset: standards bodies describe "twinning the physical network with a DT network," organised in application, twin and physical network layers [150]. In the second the network is the *substrate* for other domains' twins, and the research question is whether it can carry the traffic — "appropriate networking support is a key component to enable future DT development" [151], with the blunt corollary that "current communication networks, such as 4G and 5G, cannot meet the performance requirements of digital twin synchronization" [152]. In the third the phrase means a network *of* twins: "we define DTN as a many-to-many mapping network constructed by multiple one-to-one DTs," where "DT is suitable for reflecting a single independent object, while DTN

applies to model a group of objects" [153]. Some work combines senses, twinning an edge network whose twins are themselves interconnected [154], and visions for twin-enabled sixth-generation systems sit in the first [155].

The three senses have genuinely different referents — routers and links in the first, arbitrary assets in the third — and a reader encountering the acronym should establish which is meant before comparing results. That a subfield built around interoperability has not standardised its own central term is not merely ironic; it illustrates the definitional problem of Section 2.1 in miniature.

The substantive contribution of this literature is the argument for replay. Twinning a network enables "repeatability… and reproducibility by enabling replaying successions of events under different controlled variations of the network state" [150], and, as with the security testbeds of Section 8.4, the case rests on not experimenting live: "it is infeasible to attempt different options in a real network. Fortunately, with DITEN, we can process various communication operations in a virtual edge network, and then obtain the optimal operation parameters to feedback to the real network" [154]. The same literature supplies the field's clearest statement of fragmentation: "digital twins are harmed by fragmentation and heterogeneity since each model is developed from scratch, no common methods, models, or mechanisms are considered. There are no open and standardized open interfaces" [150]. Its own maturity assessment is candid — research on twins for fifth-generation networks and beyond "is still in its infancy," with twins acting as "a simple network simulation tool."

Two further cautions. Networking treatments are frank that full twinning may be unreachable: representing a physical system faithfully "involves representing the physical system with an extremely large number of state types and collecting and processing a prohibitively high volume of data," so that "the straightforward implementation of a fully functional DT of a complex PS operating within a complex interactive environment may be extremely hard or even impossible" [151]. And latency requirements from this literature should be read against Section 5.3's "right-time" argument: the network community optimises for synchronisation speed that several application domains do not need. One survey says so directly, calling it "a common misconception about the DT… that it should gather and process all of its data in almost real-time. However, these feats are not currently feasible, and certainly not always necessary" [35]. Industrial Internet of Things surveys bridge the twinned-object and substrate roles [152], as do treatments of the twin in the IoT context, whose property-based analysis yields the most deflating single sentence in this survey: "a fully fledged DT system does not exist yet, and it is not sure that one will be implemented soon" [156].

8.4 Security and privacy

Security divides cleanly into two questions that the literature routinely conflates. The cleanest statement of the split: "the first one addresses the security of the DT itself… The second approach is about how the DT itself can provide security to its real twin" [35].

Securing the twin. Alcaraz and Lopez state that "the confluence of all these technologies and the implicit interaction with the physical counterpart of the DT in the real world generate multiple security threats that have not yet been sufficiently studied," and enumerate threats across the twin's functionality layers and both attack surfaces: "attackers may compromise the DT considering the physical attack surface… but physical assets may also be

at risk when the DT is attacked" [157]. Most of their catalogue — software attacks, privilege escalation, man-in-the-middle, denial of service — is generic infrastructure security instantiated in a new setting. A smaller set is genuinely twin-specific and deserves separate attention: tampering with the digital thread so that the two spaces silently desynchronise; tampering with the twin's *knowledge* and *representation* rather than its data; and rogue components masquerading as twins. Their impact analysis finds confidentiality most affected "because digital models represent an exact copy of the physical counterparts, thus requiring greater protection of intellectual property," and the sensing layer most exposed "due to the bidirectional link between spaces" [157].

The National Academies name two threats that follow specifically from twinning rather than from connectivity. The first is injection into the loop: "the close integration of physical and digital systems exposes an additional attack surface for the physical system. A malicious actor can inject an attack into the feedback loop (e.g., spoofing as the digital twin)." The second is reconnaissance — "a malicious actor could manipulate the digital twin to observe vulnerable traits or behaviors of the physical system," or "interrogate the digital twin to glean intellectual property data" [34]. A faithful twin is, by construction, an excellent reconnaissance target.

Formal-methods work adds a category the infrastructure taxonomies miss: because "the models underlying the DS/DT might also become attack vectors," the twin is vulnerable to *model malware* — corrupted model updates — mitigated by "model integrity checks and the setup of authorisation for model contributions" [158]. The same work proposes extending the standard confidentiality-integrity-availability triad with safety and accountability, and introduces "semantic integrity... a mechanism that can establish what the correct range of values is." Further reviews cover cybersecurity perspectives and open challenges, observing that "creating a DT of a system has the potential to increase the attack surface, as adversaries can target both the physical systems and their digital counterparts" [159], grade perceived risk by twin type and level from component to network-of-systems [160], and assess opportunities and challenges in industrial settings [161]. Experimental work injecting attacks and measuring the degradation of twin quality metrics finds that "cyber-attacks leave measurable fingerprints across DT metrics," though detection rates for man-in-the-middle and replay attacks remain low [162].

Using the twin for security. The inverse strand is more mature than its profile suggests. A review of 61 implemented systems quantifies the split: "security simulation constitutes the largest segment (32.8%)," followed by security testing at 19.7%, with intrusion detection and anomaly detection "represented at 18% and 16.4%, respectively," and security automation at 13.1% [163]. The applications are concrete — anomaly detection frameworks evaluated against public industrial control system datasets, smart-home frameworks that execute commands against the twin before the real home, twins feeding security operations centres, twin-based penetration testing, and forensic replay comparing twin and physical state. The argument for the testbed use is the strongest in this literature: "in the recent past, multiple incidents occurred due to penetration tests that were carried out on live systems, causing severe physical damage and business interruption," whereas with a twin "penetration testers can perform security tests virtually while preserving real systems" [161]. Divergence itself becomes the signal — where controller code is manipulated, "the behavior of the affected physical device will deviate from that of its corresponding DTs. This discrepancy serves as an indication of a potential intrusion" [161].

Both directions depend on the same precondition, and it is the one Section 8.5 shows is unmet. As one review puts it, "the primary challenge lies in ensuring that the DT accurately reflects the real system, as any discrepancies could lead to ineffective or even counterproductive responses" [159]. A twin too inaccurate to trust for engineering decisions is also too inaccurate to trust as an intrusion detector.

Privacy. The privacy dimension is most acute where the twinned entity is a person, and Section 4.5 records that this is where treatment is thinnest. The National Academies state the problem in a single line that ought to govern the whole human-twin literature: "a digital twin of a human or component of a human is inherently identifiable" [34]. Anonymisation is not available as a mitigation, because fidelity to an individual is the entire point.

8.5 Verification, validation and uncertainty quantification

This is the field's binding constraint, and the evidence is now distributed across every preceding section. What remains is to say precisely why the problem is hard, because "the field needs better validation" is a platitude and the actual difficulty is specific.

Validating a twin is not validating a model. Ali and colleagues state the difference most carefully. In conventional modelling and simulation, "a calibration attempt is performed first, after which the model is subjected to the validation procedure. Once positively validated, the model is generally also assumed 'finished'." A twin is never finished, and the reason is that "only the runtime data of the system in operation are available" [11]. Three consequences follow, none of which has a settled answer. First, there is no experiment: "with DTs however, data are streamed continuously... Therefore, one must define which part of this data stream can be considered an experiment for validation." Second, there are no replications: the analyst is "relegated to grouping or batching data from equal 'experiments'," with "a risk of averaging out any of the changes that we would want to observe." Third, and most limiting, the operating envelope shrinks: "with a DT, we are limited to the bounds that occur naturally from the system's routine/regular operation... usually only a subset of the entire range for which the used simulation models were validated at design time. As such, the range you can continuously validate against is limited" [11]. A twin can only be checked where its physical counterpart happens to go, which is rarely where failure lives.

Thelen and colleagues reach the same conclusion by a different route, and their formulation is the most compact: "validation of a digital twin is not the same as validation of a general purpose prediction model," because "the model updated at one time instant can only be validated with future data; thus validation is a lagging indicator of digital twin model quality" [33]. The National Academies add the physical-side version — "as the physical twin evolves over its lifetime, it is possible to enter system states that are far from the solution scenarios that were envisioned at initial verification" [34]. And composition adds a fourth difficulty: assembling twins compounds "uncertainty, fidelity, and assumptions that stem from the individual DTs," and "may break or render obsolete existing quality assurance techniques and artifacts" [146].

What the National Academies conclude. Their treatment is the field's most authoritative and their language is unusually direct. On the evidence base: "the sentiment expressed across multiple committee information-gathering sessions is that the publicity around digital twins and digital twin solutions currently outweighs the evidence base of success," so that "it is challenging to separate what is true from what is merely aspirational" [34].

On the technical gap: twin verification, validation and uncertainty quantification "must adapt to changes in the physical counterpart, digital twin virtual models, data, and the prediction/decision task at hand," and "a gap exists between the class of problems that has been considered in traditional modeling and simulation settings and the VVUQ problems that will arise for digital twins." On reporting: "there is a lack of standards in reporting VVUQ as well as a lack of consideration of confidence in modeling outputs." And on trust, stated flatly: "a digital twin without serious considerations of VVUQ is not trustworthy. However, a rigorous VVUQ approach across all elements of the digital twin may be difficult to achieve" [34]. Their recommendation is correspondingly structural: make verification, validation and uncertainty quantification an integral part of every new digital twin programme rather than a downstream activity. They also offer a corrective to how trust is usually discussed: "trust in a digital twin need not — and probably should not — be absolute. A digital twin cannot replace reality, but it might provide adequate insight to help a decision-maker."

Two runtime notions worth adopting. Because twins cannot be certified once, the useful work has moved to continuous runtime measures. Muñoz and colleagues measure fidelity by aligning traces of physical and digital snapshots, reporting the percentage of matched snapshots along with distance measures, and treating deviations semantically — "mismatches... can be interpreted as anomalies," while "gaps usually appear when there are delays in the behavior" [39]. Their most important caveat is easy to miss and generalises to every fidelity metric in this survey: the score depends on the observation window, since "the longer the interval duration, the less significant the drops in metrics become... A longer window masks the impact of a few mismatches." Measured fidelity is a property of the measurement, not of the twin alone.

Frasheri and colleagues address a failure mode that is not inaccuracy but lateness. Discretisation and network delay mean a twin's response "may be obsolete, and should not be considered by the PT," and without detection the twin "is oblivious to the current situation in which it cannot follow the PT" while "burdening the connection with meaningless messages" [164]. Their response is conceptually the most interesting move in this literature: when the time difference exceeds an application-specific threshold, they *demote* the system, degrading "the DT (allowing for bi-directional information exchange) to a digital shadow (DS), which only receives data from the PT but does not provide feedback," and promote it back on recovery. Kritzinger's integration levels, in other words, are not a fixed property of a system's architecture but a runtime state it can fall out of and return to. Any twin claiming bidirectionality is claiming it only for the intervals in which its timing holds. The authors are honest that the mechanism misfires — "even during a normal run the calculated threshold was exceeded on occasion" — and that in persistently degraded conditions "the delay will keep accumulating."

How immature, in numbers. Dalibor and colleagues' cross-domain mapping study of 356 selected publications supplies the field's most quotable maturity statistics, and they are sobering. Only 21.63% of studies "explicitly connect the Digital Twin with their real-world counterpart." Solution proposals make up 65.45% of the corpus, while validation papers account for 3.09% and evaluation reports 4.49%. By technology readiness, proof-of-concept work at levels one to three is 50.84%, while levels seven to nine — anything approaching deployment — is 3.65%. Contributions are 62.64% methods and 0.84% metrics. Quality assurance is considered by 14.32%, and "the number of publications considering the online verification of Digital Twins with their counterparts is vanishingly low (7, 1.97%)" [13]. They

also find that twins are usually built apart from the thing they twin: 55.06% are "developed in a separate development process, i.e., not in a joint engineering process with the actual system." Complementary evidence on architectures reports that validation is dominated by illustrative examples, "validated using simplified use cases or prototypes rather than real industrial systems" [134].

Supporting technical work addresses knowledge equivalence between twin and system, where "the analysis made by the digital twin is valid and reliable only when the model is equivalent to the physical world" [165], and modelling practice generally [166]. Systematic reporting frameworks have been proposed to make twin research reproducible [167], [168], and book-length engineering treatments consolidate practice [123], [169], [170].

9. Comparison across sectors

Table 1 sets the sectors of Sections 3 to 7 side by side on the three axes of Section 2, with a representative work and the stated limitations for each.

Table 1: Digital twin applications by sector: dominant twin type, typical scale, integration level typically achieved, a representative work and the stated limitations.

Sector	Dominant twin type	Typical scale	Integration level typically achieved	Representative work	Stated limitations
Discrete manufacturing	Process and product	Unit to system	Closed loop demonstrated but rare; about 1 in 20 on the strictest count, more only if alerts count as control	Dynamic scheduling [24]; unit-level review of 96 papers [30]	Virtual-to-physical feedback "rarely achieved" [30]; 2 of 42 predictive-maintenance studies implement control [56]
Process industries	Process and system	System to system-of-systems	Mostly shadow; legacy integration blocks closure	Oil and gas overview [62]; enablers and barriers [61]	Legacy equipment integration; data quality in continuous production [61]
Healthcare — patient	Human (product-analogue)	Unit (organ) to system (physiology)	Mostly model; real-time coupling contested	Glioma radiotherapy [79]; multi-omic classifier [80]	In silico cohorts [79]; no validation standards [76]

Sector	Dominant twin type	Typical scale	Integration level typically achieved	Representative work	Stated limitations
Healthcare — operations	Process	System	Shadow; manual data entry common	Operating room ecosystem [81]; ward bed allocation [82]	No quantitative evaluation; real-time acquisition not guaranteed [81]
Cities and society	System and system-of-systems	System-of-systems	Predominantly shadow or visualisation	Urban twin critique [37]; county twin governance [93]	Emergent properties resist bottom-up computation [37]; no standard framework [93]
Civil infrastructure	Product (asset) and system	Unit (bridge) to system (network)	Shadow; "right-time" rather than real-time	Bayesian structural twin [96]; railway viaduct campaign [98]	Synthetic validation data [96]; ready for isolated measurements, not integrated systems [99]
Construction and buildings	Product and process	Unit to system (campus)	Mostly digital model inherited from BIM	Model-driven AECO twins [17]; DT manager role [100]; campus energy twin [104]	BIM is static and design-time oriented [17]; severe sensor attrition in the field [105]
Mobility	System and human	System to system-of-systems	Bidirectional in-vehicle; conceptual at city scale	Mobility twin with vehicle trial [107]; federated urban mobility [16]	No universal definition or standardisation [107], [108]; state of the art "remained conceptual" [109]
Wind and energy	Product (asset)	Unit to system (farm)	Genuine, with full-scale validation	Floating turbine validated on full-scale prototype [115]; farm-level monitoring [116]	Diagnostic rather than prescriptive scope [31]

Sector	Dominant twin type	Typical scale	Integration level typically achieved	Representative work	Stated limitations
Aerospace	Product (asset)	Unit to system	Calibrated on real assets; decisions often simulated	UAV structural twin [118]; sector position paper [32]	Decision demonstrations synthetic [118]; configuration and validation burden [32]
Maritime and offshore	Product and system	Unit to system	Shadow to bidirectional	Shipboard crane study [122]; offshore monitoring review [124]	Twins are one method among several [124]
Agriculture and food	Process and living-entity	Unit (animal, plant) to system	Genuine closed loop in controlled environments	Animal twin framework [126]; greenhouse exemplar [127]	Biological variability; architecture patterns still emerging [129]
Environment and science	System-of-systems (unowned)	Planetary to regional	Monitoring and forecasting only	Coastal lagoon and catchment [130]; federated science engine [131]	No actuation path [130]; methodology not established [130]

10. Gap analysis: what this corpus does not cover, and what it disagrees about

10.1 Sectors genuinely absent

Repeated query reformulation — trying synonyms, narrower and broader terms, and adjacent concepts — established that several sectors are not merely thin in this corpus but effectively absent. Searches for retail and commerce, banking, finance and insurance, mining and extraction, apparel, and nuclear power returned only general digital twin surveys sharing vocabulary with the query, not sector studies. Textiles are thin rather than absent: the corpus holds one chemical-fibre shop-floor deployment (Section 3.8) [69]. Supply chain work exists but is conceptual rather than deployed [113]. These absences should be read as properties of this corpus, not necessarily of the world literature; a reader needing those sectors should treat this survey as silent rather than negative on them.

Two absences are more troubling because they are absences of *perspective* rather than of industry. First, searches for equity, low-income settings and global-south perspectives surfaced only tangential material, chiefly a single institutional chapter treating the digital divide within one wealthy US county [93] and an oncology paper making an accessibility argument for cheap modalities [80]. Digital twins are capital-intensive, data-intensive and instrumentation-intensive; a literature that barely discusses who can afford them is missing something structural. Second, education appears almost exclusively as a beneficiary of

digital twins rather than a subject of study, with the notable exception of Grieves' own treatment of their role in reengineering engineering education [171].

Water utilities deserve a separate note: reformulated searching yielded exactly one strong dedicated source, an operational municipal network twin [112]. That single paper is high quality, but one paper is not a literature, and utility twins are almost certainly underrepresented here relative to practice.

10.2 The validation gap, stated as a gap

The most important thing this corpus does not contain is a validated, real-data, closed-loop digital twin in most of the sectors it covers. The pattern is remarkably consistent once assembled:

- In civil infrastructure, the most methodologically complete twin is validated on synthetic data [96], and the most thoroughly instrumented field campaign has real measurements but no closed loop [98]. No paper in this cluster demonstrates all three of real data, real asset and closed loop.
- In healthcare, the strongest optimisation result uses an *in silico* cohort [79], the strongest predictive result is a classifier without dynamics [80], and the field's own reviewers report that no accepted validation standard exists [76].
- In aerospace, calibration uses real experimental data but decision-making is demonstrated synthetically [118].
- In mobility at urban scale, the state of the art "remained conceptual," with no reviewed study using real-time data [109].

The clearest exceptions, and therefore the most useful methodological exemplars in this survey, come from offshore wind: a physics-based turbine twin validated against a full-scale prototype with its full error distribution reported including ±50% outliers [115], and a diagnostic twin that found a real fault in a year of operational data which the asset's commercial monitoring system had missed [31]. That two papers stand out so sharply against the rest is itself the finding.

"Validated" also conceals a wide range of rigour, and the four strongest claims in the heavy-engineering cluster are not equivalent evidence. One reports errors across 576 samples and admits the tail [115]; one validates on a single turbine, a single year and a single confirmed fault [31]; one is deployed industrially across many real platforms but reports no error metric at all, asking of its own improvement "the question remains: by how much?" [125]; and one reports a single 0.2-degree roll offset from one day at sea [122]. A reader compiling evidence of digital twin maturity should not stack these.

Finally, there is a reason the gap persists that is not anyone's negligence. Validating a twin is a structurally different problem from validating a simulation model, because the twin changes: "the model updated at one time instant can only be validated with future data; thus validation is a lagging indicator of digital twin model quality" [33]. A self-updating artefact cannot be certified once. Until the field has a validation theory built for moving targets, the standards described in Section 8.1 will keep covering structure and communication while leaving credibility untouched.

10.3 Unresolved cross-source disagreement

Six disagreements are live in this corpus and are recorded rather than resolved.

Whether real-time synchronisation is constitutive. Healthcare splits internally, with one source devoting a section to the requirement [75] and another explicitly rejecting it as a criterion [74]. Civil engineering rejects it on domain grounds, arguing for "right-time" integration measured in days to years and identifying open stakeholder integration rather than latency as the real challenge [98]. Mobility and networking treat latency as the central engineering problem [107], [155]. This is a genuine domain-dependent split and a good organising axis for future work.

Whether mechanism is required. Medicine hosts the sharpest version: mechanistic models are argued to be constitutive because only they support virtual intervention [74], [75], against practice that names pure classifiers twins [80] under definitions broad enough to admit them [72]. The same tension is implicit wherever data-driven surrogates replace physics [172].

Whether more fidelity helps. Set out in Section 2.6, this is a named clash rather than a drift. The fit-for-purpose position holds that fidelity "is not necessarily the highest level of model fidelity feasible and is dependent on the use case" [11], with institutional backing [34] and empirical support from episodes where detailed simulation performed *worse* than aggregate models with data assimilation [37], [38]. The maximalist position scores twinning against becoming "indistinguishable from its physical counterpart to an observer" and seeks "physical realism without any compromises" [36]. Both are published, current and widely cited, and the choice between them determines what a project builds.

What a "digital twin network" is. Section 8.3 documents three incompatible referents for one term — the network as twinned asset [150], the network as substrate for twins [151], and a federation of many twins [153]. Unlike the disagreements above, this one appears to be unrecognised rather than contested: the papers do not argue with each other because they have not noticed that they differ.

How accurate a twin needs to be. A built-environment practitioner argues that "in the built environment the accuracy of data is not that critical as the sector has little knowledge of building performance. Anything is better that what is currently available" [106]. This is flatly opposed to the position of the structural and aerospace communities, where quantified uncertainty is the entire deliverable [33], [96], [118]. Both are coherent given their sectors' decision stakes; neither generalises.

Whether the field is mature. Bibliometric maturity is repeatedly mistaken for deployment maturity [10], against deployment assessments that are uniformly early-stage [38], [72], [106], [109]. The clearest instance is in the process industries, where one review concludes that the digital twin "is a technology that has reached maturity, thus supporting more widespread use in the near future" [61] while its own corpus is dominated by conceptual papers and contains only ten studies of the sector it assesses. Manufacturing's other reviewers say the opposite: implementations are "lab-scale developments in contrast to real-world applications" [30], "simulations only validate most of the practical realizations, and therefore DTs are not really assessed in real scenarios with a physical counterpart" [24], only about 5% of studies consider the full product lifecycle [42], and deployment in manufacturing "is still at an early stage" [47].

10.4 A caution about the evidence base itself

One structural observation should temper how this survey's societal sections are read. In the urban and built-environment literature specifically, the more enthusiastic assessments cluster in a single edited collection whose chapters repeatedly adopt one industry consortium's own capability framework as their evaluative instrument [6], [94], [95]. The deflationary assessments come from an independent journal comment [37], an independent review in an engineering journal [38], and practitioners speaking against their own commercial interest [106]. That asymmetry does not make the enthusiastic sources wrong, and one of them candidly reports that industry analysts place digital twin governance at the peak of inflated expectations [94]. But a reader weighing the sectoral evidence should know where it comes from.

A related caution applies to benefit figures throughout this survey. Quantified gains are reported per pilot, seldom with a counterfactual, and almost never re-audited. Where this survey reproduces a number, it reproduces the source's framing along with it.

The most authoritative statement of this problem is worth ending the gap analysis on, because it is not a sceptic's characterisation of the field but the field's own assessment of itself. The National Academies committee, whose study was funded by the US Department of Defense, Department of Energy, National Institutes of Health and National Science Foundation, records that "the sentiment expressed across multiple committee information-gathering sessions is that the publicity around digital twins and digital twin solutions currently outweighs the evidence base of success," and concludes that "it is challenging to separate what is true from what is merely aspirational, due to a lack of agreement across domains and sectors as well as misinformation" [34]. Every quantitative finding assembled in this survey — 21.63% of studies connecting to a real counterpart, 3.65% at deployment-grade readiness, 1.97% performing online verification [13], 17% implementing control and actuation [14] — is a measurement of that same gap from a different angle.

10.5 Where the concept is being stretched

Three developments are extending the concept beyond anything in the original formulation, and each is currently under-theorised. Twins of software systems rather than physical ones — the cyber-cyber twin [83] — break the physical-virtual pairing that every definition in Section 2.1 assumes. Twins of unowned natural systems [130], [131] have no actuation path and therefore cannot satisfy the bidirectionality criterion even in principle, raising the question of whether a different word is needed. And generative approaches [173] together with large-language-model interfaces [174] introduce components whose behaviour cannot be validated by the methods the field currently possesses. Each of these is a candidate for the next serious taxonomic revision.

11. Conclusion

Reading the digital twin application literature across sectors and types yields a picture more coherent than any single sector suggests, and less flattering.

The concept works best, and closest to its own claims, where three conditions hold together: the physics is well understood, the asset is individually valuable enough to justify per-instance instrumentation, and the decision horizon is short with well-posed objectives.

Wind turbines, aircraft structures, machine tools and production schedules satisfy all three, and it is no accident that manufacturing and heavy engineering produced both the concept and its best-validated instances. Bettencourt's scope condition — that successful applications share "a relatively short time horizon, the absence of significant behavioral change, and relatively well posed objectives" [37] — turns out to describe not just cities' difficulty but the whole field's success pattern.

Where any one condition fails, the failure is predictable from which condition it is. Medicine fails on physics: theory is weaker, bodies cannot be instrumented at will, and the result is a literature rich in models and poor in twins [76], [77]. Cities fail on well-posedness: their valuable properties are emergent and their objectives are contested, so disaggregation does not buy explanation [37]. Construction fails on instrumentation economics and inherits static design-time models instead [17], [106]. Environmental science fails on actuation: one can twin a lagoon but not act on it in the way the concept assumes [130].

Two things follow for a researcher entering the field.

The first is diagnostic. The integration axis [12] remains the most efficient tool for reading any application paper, and applying it honestly disqualifies a large share of what is published as a digital twin. That is not one author's scepticism. Four independent criteria — Kritzinger's bidirectional data flow, the aerospace requirement that "to have a Digital Twin requires a physical asset" [32], Thelen's five dimensions [33], and the National Academies' bidirectionality clause [34] — reach the same verdict from different disciplines within six years of one another, and the quantitative studies agree: 21.63% of publications connect a twin to a real counterpart [13], 17% of manufacturing architectures implement control [14]. A useful habit for a new reader is to ask of any reported twin what it would take to demote it — and to notice how often the answer is nothing, because it was never a twin.

The second is agenda-setting. The field's binding constraint is not modelling capability, sensing or compute, but the absence of any accepted way to decide whether a given twin should be believed. Manufacturing has a framework standard that explicitly excludes credibility [67]; healthcare has no validation standard at all [76]; the most developed credibility proposal describes itself as immature [74]; online verification appears in under 2% of the literature [13]. And this is not simple neglect. A twin that updates itself cannot be certified once, because "validation is a lagging indicator of digital twin model quality" [33] and because the range available for continual validation is only "a subset of the entire range for which the used simulation models were validated at design time" [11]. The discipline needs a validation theory built for artefacts that move, and it does not have one.

The most promising work in this survey points at what that would look like. Fidelity measured continuously from aligned traces rather than asserted at design time [39]; integration level treated as a runtime state a system can fall out of and recover, rather than a fixed architectural property [164]; credibility attached to data as well as models, with new inputs entering unqualified until earned [74]; and error distributions reported with their tails rather than their means [115]. None of these is a modelling advance. All of them are ways of being honest at runtime about how much a twin currently deserves to be trusted — which, on the evidence assembled here, is the question the field has been least willing to ask and most needs to answer.

References

[1] "The rise of digital twins," *Nature*. Mar. 2024. Accessed: May 27, 2024. [Online]. Available: https://www.nature.com/collections/feicjiideh

[2] "The increasing potential and challenges of digital twins," *Nature Computational Science*, vol. 4, no. 3, pp. 145–146, Mar. 2024, doi: 10.1038/s43588-024-00617-4.

[3] E. VanDerHorn and S. Mahadevan, "Digital Twin: Generalization, characterization and implementation," *Decision Support Systems*, vol. 145, p. 113524, Jun. 2021, doi: 10.1016/j.dss.2021.113524.

[4] M. Grieves and J. Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems," in *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches*, F.-J. Kahlen, S. Flumerfelt, and A. Alves, Eds., Cham: Springer International Publishing, 2017, pp. 85–113. doi: 10.1007/978-3-319-38756-7_4.

[5] "Foundational Research Gaps and Future Directions for Digital Twins \textbar National Academies." Accessed: Apr. 23, 2024. [Online]. Available: https://www.nationalacademies.org/our-work/foundational-research-gaps-and-future-directions-for-digital-twins

[6] S. Sabri, K. Alexandridis, and N. Lee, "Introduction to Digital Twins," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 1–12. doi: 10.1007/978-3-031-67778-6_1.

[7] Y. Jiang, S. Yin, K. Li, H. Luo, and O. Kaynak, "Industrial applications of digital twins," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 379, no. 2207, p. 20200360, Aug. 2021, doi: 10.1098/rsta.2020.0360.

[8] C. Semeraro, M. Lezoche, H. Panetto, and M. Dassisti, "Digital twin paradigm: A systematic literature review," *Computers in Industry*, vol. 130, p. 103469, Sep. 2021, doi: 10.1016/j.compind.2021.103469.

[9] R. Eramo, F. Bordeleau, B. Combemale, M. van den Brand, M. Wimmer, and A. Wortmann, "Conceptualizing Digital Twins," *IEEE Software*, vol. 39, no. 2, pp. 39–46, Mar. 2022, doi: 10.1109/MS.2021.3130755.

[10] M.-E. Iliuţă, M.-A. Moisescu, E. Pop, A.-D. Ionita, S.-I. Caramihai, and T.-C. Mitulescu, "Digital Twin—A Review of the Evolution from Concept to Technology and Its Analytical Perspectives on Applications in Various Fields," *Applied Sciences*, vol. 14, no. 13, p. 5454, Jan. 2024, doi: 10.3390/app14135454.

[11] Z. Ali, R. Biglari, J. Denil, J. Mertens, M. Poursoltan, and M. K. Traoré, "From modeling and simulation to Digital Twin: Evolution or revolution?" *SIMULATION*, vol. 100, no. 7, pp. 751–769, Jul. 2024, doi: 10.1177/00375497241234680.

[12] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, "Digital Twin in manufacturing: A categorical literature review and classification," *Ifac-PapersOnline*, vol. 51, no. 11, pp. 1016–1022, 2018, Accessed: Dec. 11, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896318316021

[13] M. Dalibor *et al.*, "A Cross-Domain Systematic Mapping Study on Software Engineering for Digital Twins," *Journal of Systems and Software*, vol. 193, p. 111361, Nov. 2022, doi: 10.1016/j.jss.2022.111361.

[14] E. Ferko, A. Bucaioni, P. Pelliccione, and M. Behnam, "Standardisation in Digital Twin Architectures in Manufacturing," in *2023 IEEE 20th International Conference on Software Architecture (ICSA)*, Mar. 2023, pp. 70–81. doi: 10.1109/ICSA56044.2023.00015.

[15] P. Zech, C. Nardin, S. Ristov, M. Flora, and R. Breu, "Digital-Twins-as-a-Service in Construction Engineering," in *2024 IEEE 20th International Conference on Automation Science and Engineering (CASE)*, Aug. 2024, pp. 3004–3010. doi: 10.1109/CASE59546.2024.10711409.

[16] J. Li, J. Grübel, A. Nadi, M. Snelder, B. van Arem, and J. Gao, "Digital Twin Federation for Urban Mobility Assessment: A Functional Architecture for Low-Car Transformations in the Netherlands," 2025, Accessed: Oct. 25, 2025. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5370630

[17] P. Zech, E. Goldin, S. Senoner, J. Michael, and S. Hammes, "Model-driven Digital Twins for AECO," Grand Rapids, USA, Jul. 2025. doi: 10.5283/epub.77666.

[18] T. Bergs, S. Gierlings, T. Auerbach, A. Klink, D. Schraknepper, and T. Augspurger, "The concept of digital twin and digital shadow in manufacturing," *Procedia CIRP*, vol. 101, pp. 81–84, 2021, Accessed: Jan. 21, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2212827121006612

[19] C. Ellwein *et al.*, "Rethinking Asset Administration Shell Communication Types: A Systematic Mapping Study and Portfolio-Based Classification," *Production Engineering*, vol. 20, Dec. 2025, doi: 10.1007/s11740-025-01378-3.

[20] M. Heithoff, M. Trinh, J. Michael, B. Rumpe, and C. Brecher, "A Digital Shadow for Accurate Robot Motion Control: Integrating Data with Friction Models," Grand Rapids, USA, Jul. 2025. doi: 10.5283/epub.77667.

[21] G. Shao and M. Helu, "Framework for a digital twin in manufacturing: Scope and requirements," *Manufacturing Letters*, vol. 24, pp. 105–107, Apr. 2020, doi: 10.1016/j.mfglet.2020.04.004.

[22] F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, and F. Sui, "Digital twin-driven product design, manufacturing and service with big data," *The International Journal of Advanced Manufacturing Technology*, vol. 94, no. 9–12, pp. 3563–3576, Feb. 2018, doi: 10.1007/s00170-017-0233-1.

[23] K. Y. H. Lim, P. Zheng, and C.-H. Chen, "A state-of-the-art survey of Digital Twin: Techniques, engineering product lifecycle management and business innovation perspectives," *Journal of Intelligent Manufacturing*, vol. 31, no. 6, pp. 1313–1337, Aug. 2020, doi: 10.1007/s10845-019-01512-w.

[24] A. Villalonga *et al.*, "A decision-making framework for dynamic scheduling of cyber-physical production systems based on digital twins," *Annual Reviews in Control*, vol. 51, pp. 357–373, Jan. 2021, doi: 10.1016/j.arcontrol.2021.04.008.

[25] T. Müller *et al.*, "Reconfiguration management in manufacturing: A systematic literature review," *at - Automatisierungstechnik*, vol. 71, pp. 330–350, May 2023, doi: 10.1515/auto-2022-0139.

[26] B. Caesar, K. Barton, D. Tilbury, and A. Fay, "Digital Twin Framework for Reconfiguration Management: Concept & Evaluation," *IEEE Access*, vol. PP, pp. 1–1, Jan. 2023, doi: 10.1109/ACCESS.2023.3331221.

[27] F. Tao *et al.*, "Five-dimension digital twin model and its ten applications," *Comput. Integr. Manuf. Syst*, vol. 25, no. 1, pp. 1–18, 2019.

[28] N. Li *et al.*, "Six-dimensional digital twin modeling and software platform design for complex industrial systems," *Journal of Intelligent Manufacturing*, Jan. 2025, doi: 10.1007/s10845-025-02567-8.

[29] W. Jia, W. Wang, and Z. Zhang, "From simple digital twin to complex digital twin Part I: A novel modeling method for multi-scale and multi-scenario digital twin," *Advanced Engineering Informatics*, vol. 53, p. 101706, Aug. 2022, doi: 10.1016/j.aei.2022.101706.

[30] T. Böttjer *et al.*, "A review of unit level digital twin applications in the manufacturing industry," *CIRP Journal of Manufacturing Science and Technology*, vol. 45, pp. 162–189, Oct. 2023, doi: 10.1016/j.cirpj.2023.06.011.

[31] F. Stadtmann and A. Rasheed, "Diagnostic Digital Twin for Anomaly Detection in Floating Offshore Wind Energy." arXiv, Jun. 2024. doi: 10.48550/arXiv.2406.02775.

[32] "Digital Twin: Definition & Value – An AIAA and AIA Position Paper," *Aerospace Industries Association*. Accessed: Apr. 23, 2024. [Online]. Available: https://www.aia-aerospace.org/publications/digital-twin-definition-value-an-aiaa-and-aia-position-paper/

[33] A. Thelen *et al.*, "A comprehensive review of digital twin—part 2: Roles of uncertainty quantification and optimization, a battery digital twin, and perspectives," *Structural and Multidisciplinary Optimization*, vol. 66, no. 1, p. 1, Dec. 2022, doi: 10.1007/s00158-022-03410-x.

[34] *Foundational Research Gaps and Future Directions for Digital Twins*. Washington, D.C.: National Academies Press, 2024. doi: 10.17226/26894.

[35] S. Mihai *et al.*, "Digital Twins: A Survey on Enabling Technologies, Challenges, Trends and Future Prospects," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 4, pp. 2255–2291, 2022, doi: 10.1109/COMST.2022.3208773.

[36] A. Rasheed, O. San, and T. Kvamsdal, "Digital Twin: Values, Challenges and Enablers From a Modeling Perspective," *IEEE Access*, vol. 8, pp. 21980–22012, 2020, doi: 10.1109/ACCESS.2020.2970143.

[37] L. Bettencourt, "Recent achievements and conceptual challenges for urban digital twins," *Nature Computational Science*, vol. 4, pp. 150–153, Mar. 2024, doi: 10.1038/s43588-024-00604-9.

[38] H. Naderi and A. Shojaei, "Digital twinning of civil infrastructures: Current state of model architectures, interoperability solutions, and future prospects," *Automation in Construction*, vol. 149, p. 104785, May 2023, doi: 10.1016/j.autcon.2023.104785.

[39] P. Muñoz, J. Troya, and A. Vallecillo, "Towards Measuring Digital Twins Fidelity at Runtime," in *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, Linz Austria: ACM, Sep. 2024, pp. 507–512. doi: 10.1145/3652620.3688267.

[40] A. Thelen *et al.*, "A comprehensive review of digital twin — part 1: Modeling and twinning enabling technologies," *Structural and Multidisciplinary Optimization*, vol. 65, no. 12, p. 354, Nov. 2022, doi: 10.1007/s00158-022-03425-4.

[41] A. Khoshkenar and H. Nassereddine, *Exploring Digital Twin platforms across industries*. 2024. doi: 10.22260/ISARC2024/0119.

[42] M. Liu, S. Fang, H. Dong, and C. Xu, "Review of digital twin about concepts, technologies, and industrial applications," *Journal of Manufacturing Systems*, vol. 58, pp. 346–361, Jan. 2021, doi: 10.1016/j.jmsy.2020.06.017.

[43] J. Leng, D. Wang, W. Shen, X. Li, Q. Liu, and X. Chen, "Digital twins-based smart manufacturing system design in Industry 4.0: A review," *Journal of Manufacturing Systems*, vol. 60, pp. 119–137, Jul. 2021, doi: 10.1016/j.jmsy.2021.05.011.

[44] Q. Qi *et al.*, "Enabling technologies and tools for digital twin," *Journal of Manufacturing Systems*, vol. 58, pp. 3–21, Jan. 2021, doi: 10.1016/j.jmsy.2019.10.001.

[45] S. Liu, P. Zheng, and B. Jinsong, "Digital Twin-based manufacturing system: A survey based on a novel reference model," *Journal of Intelligent Manufacturing*, vol. 35, pp. 1–30, Jul. 2023, doi: 10.1007/s10845-023-02172-7.

[46] N. Anwer, R. Stark, F. Tao, and J. Erkoyuncu, "Developing and leveraging digital twins in engineering design," *CIRP Annals*, vol. 2025, Jun. 2025, doi: 10.1016/j.cirp.2025.05.002.

[47] "Digital Twins along the product lifecycle: A systematic literature ..." *Digital Twin \textbar Open Access Publishing Platform*. Mar. 2023. doi: 10.12688/digitaltwin.17807.2.

[48] Z. Chen *et al.*, "Service oriented digital twin for additive manufacturing process," *Journal of Manufacturing Systems*, vol. 74, pp. 762–776, Jun. 2024, doi: 10.1016/j.jmsy.2024.04.015.

[49] K. Gleich, S. Behrendt, M. Hörger, M. Benfer, and G. Lanza, "An Asset Administration Shell-Based Digital Product Passport as a Gaia-X Service," *Procedia CIRP*, vol. 127, pp. 224–229, Jan. 2024, doi: 10.1016/j.procir.2024.07.039.

[50] C. Zhang, W. Xu, J. Liu, Z. Liu, Z. Zhou, and D. T. Pham, "Digital twin-enabled reconfigurable modeling for smart manufacturing systems," *International Journal of Computer Integrated Manufacturing*, vol. 34, no. 7–8, pp. 709–733, Aug. 2021, doi: 10.1080/0951192X.2019.1699256.

[51] Y. Fan *et al.*, "A digital-twin visualized architecture for Flexible Manufacturing System," *Journal of Manufacturing Systems*, vol. 60, pp. 176–201, Jul. 2021, doi: 10.1016/j.jmsy.2021.05.010.

[52] Q. Qi, F. Tao, Y. Zuo, and D. Zhao, "Digital Twin Service towards Smart Manufacturing," *Procedia CIRP*, vol. 72, pp. 237–242, Jan. 2018, doi: 10.1016/j.procir.2018.03.103.

[53] M. Martinelli, J. Zhang, A.-K. Splettstoßer, M. Picone, M. Lippi, and A. Wortmann, "Hierarchical Digital Twin Ecosystem for Industrial Manufacturing Scenarios," in *2024 50th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, Aug. 2024, pp. 56–63. doi: 10.1109/SEAA64295.2024.00018.

[54] L. Esterle, M. Frasheri, and P. G. Larsen, "Autonomous Reconfiguration Enabled by Digital Twins," in *The Engineering of Digital Twins*, J. Fitzgerald, C. Gomes, and P. G. Larsen, Eds., Cham: Springer International Publishing, 2024, pp. 345–362. doi: 10.1007/978-3-031-66719-0_14.

[55] A. Wooley, J. Bitencourt, and D. Silva, "Bridging the gap between discrete event simulation and digital twin: A manufacturing case study☆☆," *Manufacturing Letters*, vol. 44, pp. 1274–1284, Aug. 2025, doi: 10.1016/j.mfglet.2025.06.147.

[56] R. van Dinter, B. Tekinerdogan, and C. Catal, "Predictive maintenance using digital twins: A systematic literature review," *Information and Software Technology*, vol. 151, p. 107008, Nov. 2022, doi: 10.1016/j.infsof.2022.107008.

[57] R. van Dinter, B. Tekinerdogan, and C. Catal, "Reference architecture for digital twin-based predictive maintenance systems," *Computers & Industrial Engineering*, vol. 177, p. 109099, Mar. 2023, doi: 10.1016/j.cie.2023.109099.

[58] R. van Dinter, G. Ekmekci, S. Rieken, B. Tekinerdogan, and C. Catal, *Architecting a Digital Twin-Based Predictive Maintenance System for Modelling Cable Joint Degradation*. 2023. doi: 10.36001/phmap.2023.v4i1.3753.

[59] D. Zhong, Z. Xia, Y. Zhu, and J. Duan, "Overview of predictive maintenance based on digital twin technology," *Heliyon*, vol. 9, no. 4, p. e14534, Apr. 2023, doi: 10.1016/j.heliyon.2023.e14534.

[60] R. Liu *et al.*, "Survey on Foundation Models for Prognostics and Health Management in Industrial Cyber-Physical Systems," *IEEE Transactions on Industrial Cyber-Physical Systems*, vol. 2, pp. 264–280, 2024, doi: 10.1109/TICPS.2024.3425326.

[61] M. Perno, L. Hvam, and A. Haug, "Implementation of digital twins in the process industry: A systematic literature review of enablers and barriers," *Computers in Industry*, vol. 134, p. 103558, Jan. 2022, doi: 10.1016/j.compind.2021.103558.

[62] T. R. Wanasinghe *et al.*, "Digital Twin for the Oil and Gas Industry: Overview, Research Trends, Opportunities, and Challenges," *IEEE Access*, vol. 8, pp. 104175–104197, 2020, doi: 10.1109/ACCESS.2020.2998723.

[63] E. Altamiranda, "A System of Systems Foundation for Digital Asset Lifecycle Management," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 59–87. doi: 10.1007/978-3-031-67778-6_4.

[64] P. Singh, N. Nidhi, J. Karpavice, M. Beliatis, and M. Presser, *Digital Dataspace and Business Ecosystem Growth for Industrial Roll-to-Roll Label Printing Manufacturing: A Case Study*. 2023.

[65] P.-E. Goffi, R. Tremblay, and B. Oakes, "Engineering a Digital Twin for the Monitoring and Control of Beer Fermentation Sampling." arXiv, Aug. 2025. doi: 10.48550/arXiv.2508.18452.

[66] G. Shao, "Use Case Scenarios for Digital Twin Implementation Based on ISO 23247," *NIST*, May 2021, Accessed: Mar. 06, 2024. [Online]. Available: https://www.nist.gov/publications/use-case-scenarios-digital-twin-implementation-based-iso-23247

[67] G. Shao, S. Frechette, and V. Srinivasan, "An Analysis of the New ISO 23247 Series of Standards on Digital Twin Framework for Manufacturing," American Society of Mechanical Engineers Digital Collection, Sep. 2023. doi: 10.1115/MSEC2023-101127.

[68] M. Heithoff, N. Jansen, J. Michael, F. Rademacher, and B. Rumpe, "Model-Based Engineering of Multi-Purpose Digital Twins in Manufacturing," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 89–126. doi: 10.1007/978-3-031-67778-6_5.

[69] F. Tao *et al.*, "makeTwin: A reference architecture for digital twin software platform," *Chinese Journal of Aeronautics*, vol. 37, no. 1, pp. 1–18, Jan. 2024, doi: 10.1016/j.cja.2023.05.002.

[70] "Case studies of digitalization for creating digital twins." Change2Twin project. Available: https://www.change2twin.eu/about/deliverables/

[71] D. Parle, G. Sharma, N. Anand, N. Padgaonkar, D. Stoddart, and D. Malley, "A Comparative Analysis for Harnessing Digital Twin Platforms for Net-Zero Manufacturing," *IEEE Access*, vol. PP, Aug. 2024, doi: 10.1109/ACCESS.2024.3447475.

[72] E. Katsoulakis *et al.*, "Digital twins for health: A scoping review," *npj Digital Medicine*, vol. 7, no. 1, pp. 1–11, Mar. 2024, doi: 10.1038/s41746-024-01073-0.

[73] R. Laubenbacher *et al.*, "Building digital twins of the human immune system: Toward a roadmap," *npj Digital Medicine*, vol. 5, no. 1, pp. 1–5, May 2022, doi: 10.1038/s41746-022-00610-z.

[74] M. Viceconti, M. De Vos, S. Mellone, and L. Geris, "Position Paper From the Digital Twins in Healthcare to the Virtual Human Twin: A Moon-Shot Project for Digital Health Research," *IEEE Journal of Biomedical and Health Informatics*, vol. 28, no. 1, pp. 491–501, Jan. 2024, doi: 10.1109/JBHI.2023.3323688.

[75] F. Emmert-Streib, S. Parkkila, R. Laubenbacher, A. Mannermaa, L. Hood, and O. Yli-Harja, "The role of digital twins in P4 medicine: A paradigm for modern healthcare," *npj Digital Medicine*, vol. 8, no. 1, p. 735, Dec. 2025, doi: 10.1038/s41746-025-02115-x.

[76] S. A. Niederer, M. S. Sacks, M. Girolami, and K. Willcox, "Scaling digital twins from the artisanal to the industrial," *Nature Computational Science*, vol. 1, no. 5, pp. 313–320, May 2021, doi: 10.1038/s43588-021-00072-5.

[77] F. Emmert-Streib, H. Cherifi, K. Kaski, S. Kauffman, and O. Yli-Harja, "Complexity data science: A spin-off from digital twins," *PNAS Nexus*, vol. 3, no. 11, p. pgae456, Nov. 2024, doi: 10.1093/pnasnexus/pgae456.

[78] Y. Liu *et al.*, "A Novel Cloud-Based Framework for the Elderly Healthcare Services Using Digital Twin," *IEEE Access*, vol. 7, pp. 49088–49101, 2019, doi: 10.1109/ACCESS.2019.2909828.

[79] A. Chaudhuri *et al.*, "Predictive digital twin for optimizing patient-specific radiotherapy regimens under uncertainty in high-grade gliomas," *Frontiers in Artificial Intelligence*, vol. 6, Oct. 2023, doi: 10.3389/frai.2023.1222612.

[80] A. Osipov *et al.*, "The Molecular Twin artificial-intelligence platform integrates multi-omic data to predict outcomes for pancreatic adenocarcinoma patients," *Nature Cancer*, vol. 5, no. 2, pp. 299–314, Feb. 2024, doi: 10.1038/s43018-023-00697-7.

[81] S. Burattini *et al.*, "An Ecosystem of Digital Twins for Operating Room Management," in *2023 IEEE 36th International Symposium on Computer-Based Medical Systems (CBMS)*, Jun. 2023, pp. 770–775. doi: 10.1109/CBMS58004.2023.00317.

[82] R. Sieve, P. Kobialka, L. Slaughter, R. Schlatte, E. B. Johnsen, and S. L. T. Tarifa, "BedreFlyt: Improving Patient Flows through Hospital Wards with Digital Twins," *Electronic Proceedings in Theoretical Computer Science*, vol. 418, pp. 1–15, May 2025, doi: 10.4204/EPTCS.418.1.

[83] C. Lu, Q. Xu, T. Yue, S. Ali, T. Schwitalla, and J. F. Nygård, "EvoCLINICAL: Evolving Cyber-Cyber Digital Twin with Active Transfer Learning for Automated Cancer Registry System," in *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, Nov. 2023, pp. 1973–1984. doi: 10.1145/3611643.3613897.

[84] "Model-based digital twins of medicine dispensers for healthcare IoT applications - Sartaj - Software: Practice and Experience - Wiley Online Library." Accessed: Apr. 05, 2024. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/spe.3311

[85] V. Villani, M. Picone, M. Mamei, and L. Sabattini, "A Digital Twin Driven Human-Centric Ecosystem for Industry 5.0," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 11291–11303, 2025, doi: 10.1109/TASE.2024.3410703.

[86] R. Azevedo *et al.*, "Human Digital Twins to Support Nurse Practitioners' Clinical Decision-Making Using Multimodal Data: A Theoretical, Methodological, and Analytical Framework," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 149–172. doi: 10.1007/978-3-031-67778-6_7.

[87] A. Karnoub, D. Stuber, J. Mertes, M. Klar, and L. Yi, "How to understand the role of human in digital twins: A digital triplet concept," *Digital Engineering*, vol. 6, p. 100052, Sep. 2025, doi: 10.1016/j.dte.2025.100052.

[88] J. M. Rožanec *et al.*, "Human-centric artificial intelligence architecture for industry 5.0 applications," *International Journal of Production Research*, vol. 61, no. 20, pp. 6847–6872, Oct. 2023, doi: 10.1080/00207543.2022.2138611.

[89] D. Shangguan, L. Chen, C. Su, J. Ding, and C. Liu, "A Triple Human-Digital Twin Architecture for Cyber-Physical Systems," *Computer Modeling in Engineering & Sciences*, vol. 131, no. 3, pp. 1557–1578, 2022, doi: 10.32604/cmes.2022.018979.

[90] T. L. Leirmo, "Digital Twins for Industry 5.0: Unlocking the Human Potential," *Procedia CIRP*, vol. 130, pp. 761–766, Jan. 2024, doi: 10.1016/j.procir.2024.10.161.

[91] A. Tóth, L. Nagy, R. Kennedy, B. Bohuš, J. Abonyi, and T. Ruppert, "The human-centric Industry 5.0 collaboration architecture," *MethodsX*, vol. 11, p. 102260, Dec. 2023, doi: 10.1016/j.mex.2023.102260.

[92] J. Barata and I. Kayser, "How will the digital twin shape the future of industry 5.0?" *Technovation*, vol. 134, p. 103025, Jun. 2024, doi: 10.1016/j.technovation.2024.103025.

[93] K. Alexandridis and M. LaFontaine, "Social and Human Dimensions of Digital Twin Technologies in Formal and Informal Institutional Settings," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 211–236. doi: 10.1007/978-3-031-67778-6_10.

[94] O. J. Pinon Fischer, S. Sabri, and Y. Chen, "Fundamentals of Digital Twins, Modeling Approaches, and Governance," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 13–29. doi: 10.1007/978-3-031-67778-6_2.

[95] D. McKee and D. Dokter, "DISCS: An Approach for Accelerating the Development of Digital Twins for Smart Cities," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 31–58. doi: 10.1007/978-3-031-67778-6_3.

[96] M. Torzoni, M. Tezzele, S. Mariani, A. Manzoni, and K. E. Willcox, "A digital twin framework for civil engineering structures," *Computer Methods in Applied Mechanics and Engineering*, vol. 418, p. 116584, Jan. 2024, doi: 10.1016/j.cma.2023.116584.

[97] M. Torzoni, A. Manzoni, and S. Mariani, "A Deep Neural Network, Multi-fidelity Surrogate Model Approach for Bayesian Model Updating in SHM," in *European Workshop on Structural Health Monitoring*, P. Rizzo and A. Milazzo, Eds., Cham: Springer International Publishing, 2023, pp. 1076–1086. doi: 10.1007/978-3-031-07258-1_108.

[98] R. Chacón *et al.*, "Digital twinning during load tests of railway bridges - case study: The high-speed railway network, Extremadura, Spain," *Structure and Infrastructure Engineering*, vol. 20, no. 7–8, pp. 1105–1119, Aug. 2024, doi: 10.1080/15732479.2023.2264840.

[99] R. Chacón, J. R. Casas, C. Ramonell, H. Posada, I. Stipanovic, and S. Škarić, "Requirements and challenges for infusion of SHM systems within Digital Twin platforms," *Structure and Infrastructure Engineering*, vol. 21, no. 4, pp. 599–615, Apr. 2025, doi: 10.1080/15732479.2023.2225486.

[100] H. Posada, R. Chacón, and C. Ramonell, "Job roles for digital twinning building construction processes. Introducing the digital twin manager position," *International Journal of Construction Management*, vol. 25, no. 12, pp. 1415–1431, Sep. 2025, doi: 10.1080/15623599.2024.2417632.

[101] H. Posada, R. Chacón, and C. Ramonell, "MATCHFEM: A COMPUTATIONAL DESIGN ASSISTANT TOOL FOR DIGITAL TWINS OF BUILDINGS AND INFRASTRUCTURE," 2025, Accessed: Oct. 25, 2025. [Online]. Available: https://ec-3.org/publications/conferences/EC32025/papers/EC32025_255.pdf

[102] I. Yitmen, S. Alizadehsalehi, İ. Akıner, and M. E. Akıner, "An Adapted Model of Cognitive Digital Twins for Building Lifecycle Management," *Applied Sciences*, vol. 11, no. 9, p. 4276, Jan. 2021, doi: 10.3390/app11094276.

[103] S. Aissat, J. Beaulieu, É. Poirier, A. Motamedi, J. Gascon-Samson, and F. Bordeleau, "A devops framework for the systematic engineering and evolution of digital twins for built assets," *Software and Systems Modeling*, Oct. 2025, doi: 10.1007/s10270-025-01330-0.

[104] C. Lin, T.-Y. Dai, A. D. Dilsiz, D. Crawley, D. Niyogi, and Z. Nagy, "UTwin: A digital twin of the UT Austin campus," in *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, Istanbul Turkey: ACM, Nov. 2023, pp. 282–283. doi: 10.1145/3600100.3626261.

[105] S.-W. Yang, Y. Lee, and S.-A. Kim, "Design and Validation of a Real-Time Maintenance Monitoring System Using BIM and Digital Twin Integration," *Buildings*, vol. 15, p. 1312, Apr. 2025, doi: 10.3390/buildings15081312.

[106] J. Fitzgerald *et al.*, "The Potential of Digital Twins: Four Industry Perspectives," in *The Engineering of Digital Twins*, J. Fitzgerald, C. Gomes, and P. G. Larsen, Eds., Cham: Springer International Publishing, 2024, pp. 19–43. doi: 10.1007/978-3-031-66719-0_2.

[107] Z. Wang *et al.*, "Mobility digital twin: Concept, architecture, case study, and future challenges," *IEEE Internet of Things Journal*, vol. 9, no. 18, pp. 17452–17467, 2022, Accessed: Jan. 08, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9724183/

[108] J. Pfeiffer *et al.*, "Modeling Languages for Automotive Digital Twins: A Survey Among the German Automotive Industry," in *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, in MODELS '24. New York, NY, USA: Association for Computing Machinery, Sep. 2024, pp. 92–103. doi: 10.1145/3640310.3674100.

[109] M. Aghaabbasi, S. Sabri, and A. Rajabifard, "Digital Twins for Assessing the Impact of Autonomous Vehicles on Built-Environment Changes," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 127–147. doi: 10.1007/978-3-031-67778-6_6.

[110] "Edge-Cloud Architectures for Urban Mobility ... `\textbar` Open Research Europe." Accessed: Feb. 06, 2026. [Online]. Available: https://open-research-europe.ec.europa.eu/articles/6-14/v1?src=rss

[111] L. Bedogni, M. Mamei, M. Picone, M. Pietri, and F. Zambonelli, "Fluid Computing & Digital Twins for intelligent interoperability in the IoT ecosystem," *Future Generation Computer Systems*, vol. 171, p. 107855, Oct. 2025, doi: 10.1016/j.future.2025.107855.

[112] P. Conejos Fuertes, F. Martínez Alzamora, M. Hervás Carot, and J. C. Alonso Campos, "Building and exploiting a Digital Twin for the management of drinking water distribution networks," *Urban Water Journal*, vol. 17, no. 8, pp. 704–713, Sep. 2020, doi: 10.1080/1573062X.2020.1771382.

[113] R. Bhandal, "Conceptualising the Application of Digital Twins in Supply Chain Management: A Path Towards Supply Chain Resilience," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 173–189. doi: 10.1007/978-3-031-67778-6_8.

[114] J. Reynolds, S. Sabri, and B. Lee, "Digital Twins for Creating Value Through 'Buildings as Batteries' Using a Mass Customization Network," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 191–209. doi: 10.1007/978-3-031-67778-6_9.

[115] E. Branlard, J. Jonkman, C. Brown, and J. Zhang, "A digital twin solution for floating offshore wind turbines validated using a full-scale prototype," *Wind Energy Science*, vol. 9, no. 1, pp. 1–24, Jan. 2024, doi: 10.5194/wes-9-1-2024.

[116] L. A. Bull *et al.*, "Data-Centric Monitoring of Wind Farms: Combining Sources of Information," in *Data Driven Methods for Civil Structural Health Monitoring and Resilience*, CRC Press, 2023.

[117] P. U. Haselbach *et al.*, "Blade research and demonstration platform," *Journal of Physics: Conference Series*, vol. 1618, no. 5, p. 052073, Sep. 2020, doi: 10.1088/1742-6596/1618/5/052073.

[118] M. G. Kapteyn, J. V. R. Pretorius, and K. E. Willcox, "A probabilistic graphical model foundation for enabling predictive digital twins at scale," *Nature Computational Science*, vol. 1, no. 5, pp. 337–347, May 2021, doi: 10.1038/s43588-021-00069-0.

[119] M. G. Kapteyn, D. J. Knezevic, and K. Willcox, "Toward predictive digital twins via component-based reduced-order models and interpretable machine learning," in *AIAA Scitech 2020 Forum*, in AIAA SciTech Forum., American Institute of Aeronautics; Astronautics, 2020. doi: 10.2514/6.2020-0418.

[120] A. Thelen, X. Huan, N. Paulson, S. Onori, Z. Hu, and C. Hu, "Probabilistic machine learning for battery health diagnostics and prognostics—review and perspectives," *npj Materials Sustainability*, vol. 2, no. 1, pp. 1–33, Jun. 2024, doi: 10.1038/s44296-024-00011-1.

[121] A. Thelen, M. Li, C. Hu, E. Bekyarova, S. Kalinin, and M. Sanghadasa, "Augmented model-based framework for battery remaining useful life prediction," *Applied Energy*, vol. 324, p. 119624, Oct. 2022, doi: 10.1016/j.apenergy.2022.119624.

[122] Z. Liu, Y. Chu, G. Li, H. P. Hildre, and H. Zhang, "Shipboard crane digital twin: An empirical study on R/V Gunnerus," *Ocean Engineering*, vol. 302, p. 117675, Jun. 2024, doi: 10.1016/j.oceaneng.2024.117675.

[123] B. J. Oakes *et al.*, "Case Studies in Digital Twins," in *The Engineering of Digital Twins*, J. Fitzgerald, C. Gomes, and P. G. Larsen, Eds., Cham: Springer International Publishing, 2024, pp. 257–310. doi: 10.1007/978-3-031-66719-0_12.

[124] H. Pezeshki, H. Adeli, D. Pavlou, and S. C. Siriwardane, "State of the art in structural health monitoring of offshore and marine structures," *Proceedings of the Institution of Civil Engineers - Maritime Engineering*, vol. 176, no. 2, pp. 89–108, Apr. 2023, doi: 10.1680/jmaen.2022.027.

[125] U. T. Tygesen, K. Worden, T. Rogers, G. Manson, and E. J. Cross, "State-of-the-Art and Future Directions for Predictive Modelling of Offshore Structure Dynamics Using Machine Learning," in *Dynamics of Civil Structures, Volume 2*, S. Pakzad, Ed., Cham: Springer International Publishing, 2019, pp. 223–233. doi: 10.1007/978-3-319-74421-6_30.

[126] A. Youssef *et al.*, "IUMENTA: A generic framework for animal digital twins within the Open Digital Twin Platform." arXiv, Nov. 2024. doi: 10.48550/arXiv.2411.10466.

[127] E. Kamburjan *et al.*, "GreenhouseDT: An Exemplar for Digital Twins," in *Proceedings of the 19th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, in SEAMS '24. New York, NY, USA: Association for Computing Machinery, Jun. 2024, pp. 175–181. doi: 10.1145/3643915.3644108.

[128] E. Kamburjan, N. Bencomo, S. L. Tapia Tarifa, and E. B. Johnsen, "Declarative Lifecycle Management in Digital Twins," in *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, Linz Austria: ACM, Sep. 2024, pp. 353–363. doi: 10.1145/3652620.3688248.

[129] B. Tekinerdogan and C. Verdouw, "Systems Architecture Design Pattern Catalog for Developing Digital Twins," *Sensors*, vol. 20, no. 18, p. 5103, Jan. 2020, doi: 10.3390/s20185103.

[130] Y. Ye, A. González-Vidal, A. Cisterna-García, A. Pérez-Ruzafa, M. A. Z. Izquierdo, and A. F. Skarmeta, "Advancing Towards a Marine Digital Twin Platform: Modeling the Mar Menor Coastal Lagoon Ecosystem in the South Western Mediterranean." arXiv, Sep. 2024. doi: 10.48550/arXiv.2409.10134.

[131] "interTwin: Advancing Scientific Digital Twins through AI, Federated Computing and Data."

[132] C. Schmidt, F. Volz, L. Stojanovic, and G. Sutschet, "Increasing Interoperability between Digital Twin Standards and Specifications: Transformation of DTDL to AAS," *Sensors*, vol. 23, p. 7742, Sep. 2023, doi: 10.3390/s23187742.

[133] A. Budiardjo and D. Migliori, "Digital twin system interoperability framework," Tech. rep. Digital Twin Consortium, East Lansing, Michigan, 2021. Accessed: Feb. 25, 2025. [Online]. Available: https://www.digitaltwinconsortium.org/pdf/Digital-Twin-System-Interoperability-Framework-12072021.pdf

[134] E. Ferko, A. Bucaioni, and M. Behnam, "Architecting Digital Twins," *IEEE Access*, vol. 10, pp. 50335–50350, 2022, doi: 10.1109/ACCESS.2022.3172964.

[135] J. Pfeiffer *et al.*, "Towards a Unifying Reference Model for Digital Twins of Cyber-Physical Systems." arXiv, Jul. 2025. doi: 10.48550/arXiv.2507.04871.

[136] R. Richstein and K.-U. Schröder, "Characterizing the Digital Twin in Structural Mechanics," *Designs*, vol. 8, no. 1, p. 8, Feb. 2024, doi: 10.3390/designs8010008.

[137] L. Jinzhi, Y. Zhaorui, X. Zheng, W. Jian, and K. Dimitris, "Exploring the concept of Cognitive Digital Twin from model-based systems engineering perspective," *The International Journal of Advanced Manufacturing Technology*, vol. 121, Aug. 2022, doi: 10.1007/s00170-022-09610-5.

[138] H. Marah and M. Challenger, "(Re-)Engineering Digital Twins Towards Federation: Vision and Roadmap," in *Leveraging Applications of Formal Methods, Verification and Validation. Software Engineering Methodologies*, T. Margaria and B. Steffen, Eds., Cham: Springer Nature Switzerland, 2025, pp. 60–81. doi: 10.1007/978-3-031-75387-9_5.

[139] K. Duran *et al.*, "Toward Digital Twin-as-a-Service (DTaaS) Platforms: A Survey on Architecture, Design Requirements, and Performance Metrics," *IEEE Communications Surveys & Tutorials*, vol. 28, pp. 1845–1878, 2026, doi: 10.1109/COMST.2025.3635582.

[140] S. Gil, P. H. Mikkelsen, C. Gomes, and P. G. Larsen, "Survey on open-source digital twin frameworks–A case study approach," *Software: Practice and Experience*, vol. 54, no. 6, pp. 929–960, Jun. 2024, doi: 10.1002/spe.3305.

[141] J. Robles, C. Martín, and M. Díaz, "OpenTwins: An open-source framework for the development of next-gen compositional digital twins," *Computers in Industry*, vol. 152, p. 104007, Aug. 2023, doi: 10.1016/j.compind.2023.104007.

[142] S. Infante *et al.*, "Integrating FMI and ML/AI models on the open-source digital twin framework OpenTwins," *Software Practice and Experience*, Mar. 2024, doi: 10.1002/spe.3322.

[143] S. Aheleroff, X. Xu, R. Y. Zhong, and Y. Lu, "Digital Twin as a Service (DTaaS) in Industry 4.0: An Architecture Reference Model," *Advanced Engineering Informatics*, vol. 47, p. 101225, Jan. 2021, doi: 10.1016/j.aei.2020.101225.

[144] Z. Tang, D. Zhuang, and J. Zhang, "Evaluation framework for domain-specific digital twin platforms," *Scientific Reports*, vol. 15, no. 1, p. 10544, Mar. 2025, doi: 10.1038/s41598-024-82154-8.

[145] J. Michael, J. Pfeiffer, B. Rumpe, and A. Wortmann, "Integration challenges for digital twin systems-of-systems," in *Proceedings of the 10th IEEE/ACM International Workshop on Software Engineering for Systems-of-Systems and Software Ecosystems*, in SESoS '22. New York, NY, USA: Association for Computing Machinery, Nov. 2022, pp. 9–12. doi: 10.1145/3528229.3529384.

[146] B. Combemale *et al.*, "On the Challenges of Integrating Digital Twins," in *2nd International Conference on Engineering Digital Twins (EDTconf 2025)*, 2025. Accessed: Sep. 29, 2025. [Online]. Available: https://inria.hal.science/hal-05221809/

[147] P. Kuruppuarachchi, S. Rea, and A. McGibney, "An architecture for composite digital twin enabling collaborative digital ecosystems," in *2022 IEEE 25th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, IEEE, 2022, pp. 980–985. Accessed: Feb. 25, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9776073/

[148] E. Kamburjan, A. Pferscher, R. Schlatte, R. Sieve, S. L. T. Tarifa, and E. B. Johnsen, "Semantic Reflection and Digital Twins: A Comprehensive Overview," in *The Combined Power of Research, Education, and Dissemination*, vol. 15240, M. Hinchey and B. Steffen, Eds., Cham: Springer Nature Switzerland, 2025, pp. 129–145. doi: 10.1007/978-3-031-73887-6_11.

[149] A. Ricci, A. Croatti, S. Mariani, S. Montagna, and M. Picone, "Web of Digital Twins," *ACM Transactions on Internet Technology*, vol. 22, no. 4, pp. 1–30, Nov. 2022, doi: 10.1145/3507909.

[150] A. Hakiri, A. Gokhale, S. B. Yahia, and N. Mellouli, "A comprehensive survey on digital twin for future networks and emerging Internet of Things industry," *Computer Networks*, vol. 244, p. 110350, May 2024, doi: 10.1016/j.comnet.2024.110350.

[151] M. Vaezi *et al.*, "Digital twins from a networking perspective," *IEEE Internet of Things Journal*, vol. 9, no. 23, pp. 23525–23544, 2022, Accessed: Jan. 08, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9863238/

[152] H. Xu, J. Wu, Q. Pan, X. Guan, and M. Guizani, "A Survey on Digital Twin for Industrial Internet of Things: Applications, Technologies and Tools," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 4, pp. 2569–2598, 2023, doi: 10.1109/COMST.2023.3297395.

[153] Y. Wu, K. Zhang, and Y. Zhang, "Digital Twin Networks: A Survey," *IEEE Internet of Things Journal*, vol. 8, no. 18, pp. 13789–13804, Sep. 2021, doi: 10.1109/JIOT.2021.3079510.

[154] F. Tang, X. Chen, T. K. Rodrigues, M. Zhao, and N. Kato, "Survey on Digital Twin Edge Networks (DITEN) Toward 6G," *IEEE Open Journal of the Communications Society*, vol. 3, pp. 1360–1381, 2022, doi: 10.1109/OJCOMS.2022.3197811.

[155] L. U. Khan, W. Saad, D. Niyato, Z. Han, and C. S. Hong, "Digital-Twin-Enabled 6G: Vision, Architectural Trends, and Future Directions," *IEEE Communications Magazine*, vol. 60, no. 1, pp. 74–80, Jan. 2022, doi: 10.1109/MCOM.001.21143.

[156] R. Minerva, G. M. Lee, and N. Crespi, "Digital twin in the IoT context: A survey on technical features, scenarios, and architectural models," *Proceedings of the IEEE*, vol. 108, no. 10, pp. 1785–1824, 2020, Accessed: Feb. 25, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9120192/

[157] C. Alcaraz and J. Lopez, "Digital Twin: A Comprehensive Survey of Security Threats," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 3, pp. 1475–1503, 2022, doi: 10.1109/COMST.2022.3171465.

[158] T. Kulik, Z. Kazemi, and P. G. Larsen, "Security and Privacy-related Issues in a Digital Twin Context," in *The Engineering of Digital Twins*, J. Fitzgerald, C. Gomes, and P. G. Larsen, Eds., Cham: Springer International Publishing, 2024, pp. 313–344. doi: 10.1007/978-3-031-66719-0_13.

[159] A. Jaber, I. Koufos, and M. Christopoulou, "A Comprehensive State-of-the-Art Review for Digital Twin: Cybersecurity Perspectives and Open Challenges," in *Advances on P2P, Parallel, Grid, Cloud and Internet Computing*, L. Barolli, Ed., Cham: Springer Nature Switzerland, 2025, pp. 78–98. doi: 10.1007/978-3-031-76462-2_8.

[160] N. Alhamam, M. M. Hafizur Rahman, and A. Aljughaiman, "A Comprehensive Review on Cybersecurity of Digital Twins Issues, Challenges, and Future Research Directions," *IEEE Access*, vol. 13, pp. 45106–45124, 2025, doi: 10.1109/ACCESS.2025.3545004.

[161] A. J. G. de Azambuja, T. Giese, K. Schützer, R. Anderl, B. Schleich, and V. R. Almeida, "Digital Twins in Industry 4.0 – Opportunities and challenges related to Cyber Security," *Procedia CIRP*, vol. 121, pp. 25–30, Jan. 2024, doi: 10.1016/j.procir.2023.09.225.

[162] M. Picone, E. Bouquillon, M. Pietri, and M. Mamei, "Assessing the Impact of Cybersecurity Attacks on Digital Twin Metrics: An Experimental Study," in *2026 IEEE 23rd Consumer Communications & Networking Conference (CCNC)*, IEEE, 2026, pp. 1–4. Accessed: Mar. 27, 2026. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/11366535/

[163] A. R. Qureshi, A. Asensio, M. Imran, J. Garcia, and X. Masip-Bruin, "A survey on security enhancing Digital Twins: Models, applications and tools," *Computer Communications*, vol. 238, p. 108158, Jun. 2025, doi: 10.1016/j.comcom.2025.108158.

[164] M. Frasheri *et al.*, "Addressing time discrepancy between digital and physical twins," *Robotics and Autonomous Systems*, vol. 161, p. 104347, Mar. 2023, doi: 10.1016/j.robot.2022.104347.

[165] N. Zhang, R. Bahsoon, N. Tziritas, and G. Theodoropoulos, "Knowledge Equivalence in Digital Twins of Intelligent Systems," *ACM Transactions on Modeling and Computer Simulation*, vol. 34, no. 1, pp. 1–37, Jan. 2024, doi: 10.1145/3635306.

[166] G. Abbiati, C. Gomes, M. Sandberg, Z. Kazemi, S. T. Hansen, and P. G. Larsen, "Modelling for Digital Twins," in *The Engineering of Digital Twins*, J. Fitzgerald, C. Gomes, and P. G. Larsen, Eds., Cham: Springer International Publishing, 2024, pp. 89–127. doi: 10.1007/978-3-031-66719-0_5.

[167] S. Gil, B. Oakes, C. Gomes, M. Frasheri, and P. Larsen, "Toward a systematic reporting framework for Digital Twins: A cooperative robotics case study," *SIMULATION*, Aug. 2024, doi: 10.1177/00375497241261406.

[168] A. Barbie and W. Hasselbring, "Toward Reproducibility of Digital Twin Research: Exemplified with the PiCar-X." arXiv, Aug. 2024. doi: 10.48550/arXiv.2408.13866.

[169] P. G. Larsen, J. Fitzgerald, and C. Gomes, "Engineering Digital Twins for Cyber-Physical Systems," in *The Engineering of Digital Twins*, J. Fitzgerald, C. Gomes, and P. G. Larsen, Eds., Cham: Springer International Publishing, 2024, pp. 3–17. doi: 10.1007/978-3-031-66719-0_1.

[170] J. Fitzgerald, C. Gomes, and P. G. Larsen, Eds., *The Engineering of Digital Twins*. Cham: Springer International Publishing, 2024. doi: 10.1007/978-3-031-66719-0.

[171] M. Grieves, "Digital Twins and Their Role in Reengineering Engineering Education," in *Digital Twin: Fundamentals and Applications*, S. Sabri, K. Alexandridis, and N. Lee, Eds., Cham: Springer Nature Switzerland, 2024, pp. 237–261. doi: 10.1007/978-3-031-67778-6_11.

[172] X. Liu and I. David, "AI simulation by digital twins: Systematic survey, reference framework, and mapping to a standardized architecture," *Software and Systems Modeling*, Aug. 2025, doi: 10.1007/s10270-025-01306-0.

[173] C. Savaglio, V. Barbuto, F. Mangione, and G. Fortino, "Generative Digital Twins: A Novel Approach in the IoT Edge-Cloud Continuum," *IEEE Internet of Things Magazine*, vol. 8, no. 1, pp. 42–48, Jan. 2025, doi: 10.1109/IOTM.001.2400035.

[174] F. Franco, L. Lamazzi, M. Picone, M. Savarese, C. A. Grazia, and L. Bedogni, "Customizing Human Machine Interfaces leveraging Digital Twins and Large Language Models," in *2026 IEEE 23rd Consumer Communications & Networking Conference (CCNC)*, Jan. 2026, pp. 1–6. doi: 10.1109/CCNC65079.2026.11366401.