# Do Buzzwords Dream of Clearer Substance?

## The Journey of ‘Digital Twins’ Toward Becoming an Operational Concept

**Jérôme Vetillard**

VP R&D | Chief Product Officer | AI-Powered Healthcare & Life Sciences Products
Compliance by Design | PhD ENS Ulm | MIT Sloan | IE Business School & Brown University

## PART I — ANALYTICAL FRAMEWORK

## 1. Introduction: A Concept in Search of Definition

The concept of ‘Digital Twin’ has generated sustained and growing interest across industries, from manufacturing and aerospace to healthcare and urban planning. Google Trends data confirms a near-exponential rise in search volume over the past decade.

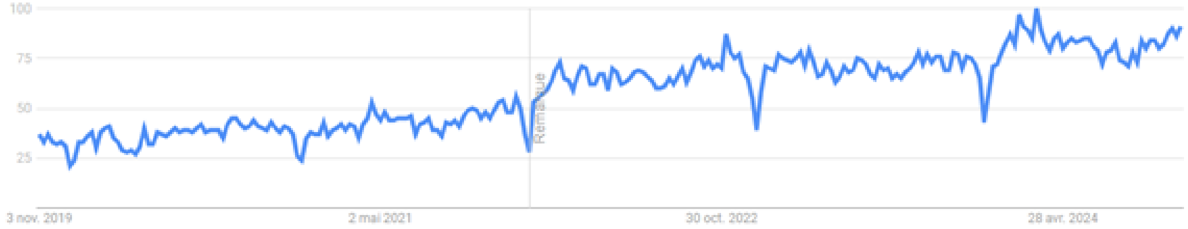

*Figure 1: Growing interest for the concept of “Digital Twins” (Google Trends)*

Yet this enthusiasm coexists with a persistent definitional ambiguity that risks relegating Digital Twins to the status of a buzzword rather than an operational concept.

The term itself is richly polysemic. ‘Digital’ can encompass IoT sensor streams, wearable devices, 5G connectivity, SCADA systems, knowledge graphs, or AI algorithms. ‘Twinning’ can refer to anything from a static 3D visualization to a fully bidirectional, real-time control loop. Each word carries multiple meanings depending on the industrial context and specific use cases, which is why the compound concept is so resistant to univocal definition.

For many, the concept remains strongly associated with the 3D representation of a real-world object, especially with the rise of augmented and virtual reality. As Sébastien Brasseur noted, a 2017 study by Negri et al. found that the most frequently occurring words in Digital Twin literature were “physical” and “product,” indicating a strong emphasis on a physical product-centered perspective. Nonetheless, the potential of Digital Twins should extend beyond tangible physical products, encompassing processes, systems, and even entire organizations.

In our discussions with customers, it has become clear that there is no universally accepted definition of what a Digital Twin truly is, nor a shared understanding of how it can drive ROI-proven initiatives at scale within existing value chains and organizational structures. This conceptual ambiguity risks Digital Twins being dismissed as mere hype. The purpose of this article is to cut through this ambiguity by proposing a structured analytical framework—grounded in existing standards and literature—and then examining how this framework translates into operational practice.

# 2. Reflecting on Recent History: The Cautionary Tale of Predictive Maintenance

Before constructing a definitional framework, it is instructive to examine a case where Digital Twin-adjacent technology generated enormous enthusiasm but ultimately underdelivered.

Predictive maintenance was once hailed as an obvious, high-ROI use case. Though not explicitly labeled as “Digital Twin-based,” it embodied the core concept: a digital representation of a complex physical object aimed at forecasting component failures so operators could preemptively replace parts and avoid costly production downtime.

However, the business case frequently faltered. Many ROI projections optimistically assumed a leap from no maintenance at all to a fully predictive landscape—overlooking the fact that “scheduled maintenance” already existed as a robust practice to reduce downtime. This oversight significantly inflated expectations and undermined the projected impact. The lesson is clear: Digital Twin initiatives must benchmark against existing operational practices, not against the absence of any solution. Realistic ROI calculation requires honest assessment of the incremental value over current processes, not hypothetical greenfield comparisons.

# 3. Anatomy of a Digital Twin: Components and Properties

## 3.1 Core Components

Digital Twins comprise two primary components: a physical entity and a digital counterpart. Rather than focusing solely on the digital model, the Digital Twin paradigm integrates both the real-world system and its virtual representation, along with the interaction between them.

The physical component, known as the Physical Object (PO), represents the actual item or system—whether a device, product, hardware, or even a physical process operating in the real world. The digital component, or Logical Object (LO), is a virtual model of the physical system, generated through software that integrates data and algorithms. In essence, the Digital Twin links these two elements: the tangible system existing in the real world and its software-based counterpart, which can simulate and predict the system's performance and behavior.

## 3.2 Essential Properties

As described in the Handbook of Digital Twins (editor Zhyhan Lyu), Digital Twins should exhibit the following properties:

**Representativeness and Contextualization.** The LO should credibly mirror the PO, although capturing all facets can be complex and costly. DT models are therefore designed with specific objectives aligned to their application context, focusing on the essential properties, characteristics, and behaviors needed to qualify as an LO. This concept is often misunderstood, particularly in Western contexts influenced by analytical Cartesian thinking. The LO is not an "exact" replica of the PO but rather an abstraction. The level of abstraction must be carefully chosen, as higher complexity increases costs exponentially without a proportionate gain in quality or relevance.

**Reflection.** The PO's status, attributes, and behaviors—which may vary over time—are accurately represented in the LO, with each relevant value uniquely reflected.

**Entanglement.** Entanglement describes the communication link between the PO and LO, ensuring all necessary information flows to the LO for accurate representation. This link may involve a real-time aspect and can leverage IoT, IoMT, and wearable technologies.

**Replication.** The ability to create multiple instances of a PO across different virtual environments, enabling scalability of DT operations. The same DT model can be deployed across multiple data streams in various production lines to monitor and control production quality.

**Persistence.** While the PO may face physical limitations, the LO maintains the DT's continuity by overcoming these restrictions to ensure constant availability. Once the LO has proven reliable and accurate, it can independently generate synthetic data, even in the absence of the PO.

**Memorization.** The LO must retain and represent all relevant historical and current data of the DT, which is crucial for maintaining the DT throughout its lifecycle. When integrated with MLOps, this capability enables regular updates and facilitates model forking to create refined or specialized versions.

**Composability.** Multiple objects can be combined into a single composite structure, enabling observation and control of both the entire composite and its individual components.

This is essential for a multimodal DT approach, integrating additional data dimensions to enhance the DT's representation of the PO.

**Accountability/Management.** Ensures comprehensive management of the DT, enabling interaction between different LOs and the assembly of larger aggregates for DT construction.

**Servitization.** Provides users with services, functionalities, and data access related to the PO through DT's tools, software, and interfaces.

**Predictability.** Supports the simulation of the LO's behavior and interactions over time or within specific contexts to predict the PO's future performance.

**Programmability.** Offers APIs that enable the programming of DT functions and integration with external systems.

## 4. Maturity Models: From Digital Model to Collaborative Twin

Based on the "entanglement" property—the degree and directionality of data synchronization between PO and LO—Digital Twins can be classified into progressive maturity stages:

| Maturity Stage | Data Synchronization | Characteristics |
|---|---|---|
| Digital Model | No automated exchange between PO and LO | Static digital model created asynchronously from PO data. Changes in one do not affect the other. |
| Digital Shadow | One-way: PO → LO | Changes in PO reflected in LO, but not vice versa. E.g., real-time fraud detection monitoring financial data streams. |
| Digital Twin | Bidirectional: PO ↔ LO | Full interactive stage. LO functions as a virtual cockpit for real-time analysis and control commands to actuators. |
| Cognitive DT | Bidirectional + predictive | AI-driven anticipation of PO behavior. Can correct deviations without human intervention. |
| Collaborative DT | Bidirectional + advisory | Acts as decision-support assistant for human operators. Provides |

| | | |
|---|---|---|
| | | insights and recommendations without direct PO control. |

While the first three stages primarily relate to entanglement and data synchronization, the Cognitive and Collaborative stages focus on AI-driven insights and user experience, enhancing human collaboration and decision support. The distinction between these two advanced stages is significant: Cognitive DTs emphasize autonomous corrective action, while Collaborative DTs emphasize augmented human judgment—a design choice with profound implications for trust, liability, and regulatory compliance.

# 5. The Standardization Landscape: Situating Digital Twins in Normative Frameworks

Any operational definition of Digital Twins must be situated within the broader normative landscape. Several international standardization efforts provide essential context that shapes both the conceptual boundaries and the interoperability requirements of Digital Twin implementations.

## ISO 23247: Digital Twin Framework for Manufacturing

ISO 23247 provides a reference architecture for Digital Twins in manufacturing environments. It defines four core entities: the observable manufacturing element, the data collection and device control entity, the Digital Twin entity itself, and the user entity. The standard establishes requirements for data exchange, interoperability, and lifecycle management. While oriented toward discrete and process manufacturing, its architectural patterns are transferable to other domains, including healthcare.

## ISO/DIS 24576 and the Digital Twin Consortium

ISO/DIS 24576, currently under development, aims to establish a broader definitional framework for Digital Twins across industries. In parallel, the Digital Twin Consortium—an industry body with members from Microsoft, Dell, Ansys, and others—has published a capability periodic table and a set of technology requirements that emphasize composability, security, and trustworthiness as foundational properties.

## The Asset Administration Shell (AAS)

The Asset Administration Shell, developed within the Plattform Industrie 4.0 initiative, provides a concrete interoperability layer for Digital Twins. The AAS defines a standardized digital representation of an asset—including its properties, capabilities, and interfaces—that can be exchanged across organizational boundaries. It represents the most mature implementation of Digital Twin interoperability in the industrial context and provides a model that healthcare implementations can adapt, particularly for medical device management and supply chain traceability.

### A Working Operational Definition

Synthesizing the literature, the properties identified by Lyu, and the normative frameworks above, we propose the following operational definition:

*“A Digital Twin is a virtual representation of a system—whether a physical object, complex machine, organism, process, or organization—that mirrors its real-world characteristics and behaviors to a defined extent, with specified levels of reliability and complexity. By integrating real-time or asynchronous data and simulations, it allows for monitoring, analysis, predictions, and, in some cases, interactive feedback loops that enable users to influence and optimize the system’s performance. This interactive capability enhances decision-making and control across the system’s lifecycle.”*

This definition is deliberately broad in scope (accommodating diverse system types), explicit about the partiality of representation (avoiding the “exact replica” fallacy), and agnostic about maturity level (applicable from Digital Model to Collaborative DT). It is a working operational definition—not a normative standard—positioned to complement rather than compete with the ISO and AAS frameworks.

# 6. Ontological Intelligence: The Missing Layer in Digital Twin Architecture

A dimension that is frequently underexplored in Digital Twin literature is the role of formal ontologies in structuring the relationship between the Physical Object and the Logical Object. Yet this is arguably the most consequential architectural decision in any Digital Twin implementation: how is the level of abstraction determined, and how is the semantic structure of the Logical Object defined?

## 6.1 Ontology as Structural Blueprint

An ontology—in the computer science sense—is a formal representation of the concepts, relationships, and constraints within a specific domain. In the context of Digital Twins, the ontology serves as the structural blueprint that determines which aspects of the Physical Object are represented in the Logical Object, how they relate to each other, and what inference rules govern the model’s behavior.

This ontological layer addresses the core challenge of representativeness identified in Section 3: the LO cannot and should not replicate every aspect of the PO. The ontology formalizes the abstraction decision—making explicit which dimensions are modeled, which are deliberately excluded, and why. Without this formalization, the abstraction is implicit and undocumented, making it difficult to validate, reproduce, or audit the Digital Twin's design choices.

### 6.2 Ontology-Driven Variable Selection

In practice, the ontology guides variable selection for the Digital Twin model. Consider a healthcare Digital Twin designed to predict cardiac decompensation risk. A naïve approach would ingest all available patient data—demographics, laboratory results, imaging, genomics, wearable sensor streams—and rely on machine learning to identify relevant features. This approach is computationally expensive, prone to overfitting, and produces models that are difficult to interpret clinically.

An ontology-driven approach begins instead with a formal model of the clinical domain: the pathophysiology of heart failure, the causal relationships between biomarkers and clinical outcomes, the temporal dynamics of decompensation. This ontological model constrains the feature space to clinically meaningful variables, guides the selection of appropriate modeling techniques, and provides a basis for clinical validation. The machine learning component then operates within this structured knowledge framework rather than in an unconstrained feature space.

### 6.3 Composability Through Shared Ontologies

The ontological layer also enables the composability property identified in Section 3. When multiple Digital Twins share a common ontological framework—or when their respective ontologies are formally aligned—they can be composed into higher-order Digital Twins that represent composite systems. For example, a patient-level Digital Twin (integrating cardiac, metabolic, and pharmacological sub-models) can be constructed from specialized component twins if these components share a common ontological structure that defines their interfaces and interaction rules.

This composability is not merely a technical convenience; it is the mechanism by which Digital Twins scale from isolated models to integrated systems capable of representing the complexity of real-world environments—whether a hospital ward, a manufacturing line, or an epidemiological surveillance system.

## 7. Cybersecurity and Data Governance: The Operational Prerequisites

The properties of entanglement, persistence, memorization, and servitization—identified in Section 3 as essential Digital Twin characteristics—each carry significant cybersecurity and data governance implications that must be addressed as foundational architectural concerns rather than afterthoughts.

## 7.1 Access Control in Synchronized Environments

The entanglement property—data synchronization between PO and LO—creates a data flow that must be governed by fine-grained access control. In traditional information systems, access control lists (ACLs) are relatively straightforward: users authenticate, receive role-based permissions, and access authorized resources. When a Digital Twin operates across multiple data sources with different sensitivity levels and ownership structures, the ACL problem becomes multi-dimensional.

Consider a Digital Twin deployed in a hospital setting that integrates electronic health records (EHRs), laboratory results, imaging data, and wearable sensor streams. Each data source has its own access control regime governed by the GDPR, the Medical Device Regulation (MDR), and local hospital policies. The Digital Twin’s LO must enforce access control that respects not only the authorization level of the requesting user but also the purpose of the access (clinical care vs. research vs. billing), the sensitivity classification of the specific data elements, and the regulatory context applicable to each data source.

A common architectural weakness is the use of a single service account for the Digital Twin, which effectively grants blanket access to all data sources regardless of the end user’s actual authorization level. This constitutes a de facto privilege escalation risk and violates the principle of minimum necessary access. Robust Digital Twin architectures must implement identity propagation—ensuring that the original user’s authorization context is maintained throughout the data processing chain—and purpose-based access control mechanisms that enforce regulatory constraints at the data element level.

## 7.2 Data Leakage and the Persistence/Memorization Paradox

The persistence and memorization properties create a paradox when confronted with data protection regulations. GDPR’s right to erasure (Article 17) requires that personal data be deleted upon request. Yet a Digital Twin that has “memorized” patient data across its lifecycle—and potentially generated synthetic data, model parameters, and derived insights from that data—poses a genuine challenge: what constitutes complete erasure when the original data has been transformed into model weights, statistical aggregates, and synthetic derivatives?

This is not a hypothetical concern. The concept of “machine unlearning”—removing the influence of specific training data from a model—is an active research area with no production-ready solutions for complex models. Digital Twin architectures must therefore implement data lineage tracking from ingestion through model training to output generation, so that the impact of any individual data point can be traced and, if necessary, excised.

The replication property compounds this challenge: if a Digital Twin model has been replicated across multiple environments (production, staging, research), each replica inherits the data governance obligations of the original, and erasure requests must propagate across all instances. Without automated data lineage and governance frameworks, this propagation is operationally infeasible at scale.

## 7.3 Attack Surface Expansion: When the Twin Becomes a Target

Digital Twins that reach the bidirectional synchronization stage (true Digital Twin and beyond) introduce a particularly concerning attack surface: the LO can send control commands to the PO. A compromised Digital Twin controlling industrial actuators, medical devices, or critical infrastructure components poses catastrophic risks.

The servitization property—exposing Digital Twin functionality through APIs and interfaces—further expands this attack surface. Each API endpoint represents a potential entry point for adversaries. In the context of LLM-based Cognitive Digital Twins, prompt injection attacks add a novel threat vector: adversarial content embedded in data streams processed by the AI component could manipulate the Digital Twin's reasoning and, consequently, its control commands.

Defensive architectures must implement defense-in-depth strategies: network segmentation between the LO and PO, cryptographic authentication of control commands, anomaly detection on the synchronization channel, rate limiting on actuator commands, and mandatory human-in-the-loop approval for safety-critical actions. For healthcare Digital Twins, these requirements align with the EU AI Act's obligations for high-risk AI systems and the MDR's requirements for software as a medical device (SaMD).

## 7.4 Auditability: Regulatory Compliance as Architectural Requirement

The EU AI Act's transparency requirements, the MDR's clinical evaluation obligations, and GDPR's accountability principle all converge on a common requirement: auditability. A Digital Twin deployed in a regulated environment must provide end-to-end traceability from data ingestion through model reasoning to output generation and, where applicable, control actions.

This requirement extends across the entire Digital Twin lifecycle. During development, data provenance, model architecture decisions, and validation protocols must be documented. During operation, every data access, model inference, and output must be logged with sufficient context to reconstruct the reasoning chain. During decommissioning, data retention and deletion obligations must be verifiable.

Current Digital Twin platforms rarely provide this level of auditability natively. Integrating audit trail capabilities must be a first-order architectural concern, not a compliance afterthought.

# PART II — OPERATIONAL APPLICATION

Having established an analytical framework grounded in existing standards, the role of ontological intelligence, and the cybersecurity prerequisites, we now examine how these concepts translate into operational practice through TweenMe.

# 8. TweenMe: Operationalizing the Framework

## 8.1 Positioning in the Maturity Model

TweenMe is positioned as a universal generator of Digital Models and Digital Shadows, with the architectural capacity to support Cognitive and Collaborative Digital Twins:

- **Asynchronous Data Integration:** TweenMe integrates datasets asynchronously within an automated data pipeline, enabling the creation of flexible digital representations of physical objects or systems. Additionally, TweenMe includes advanced ETL capabilities to ingest data in real time from points of service when required.
- **Digital Model Generation with MLOps:** TweenMe produces Digital Models leveraging MLOps, supporting various use cases including regression, clustering, segmentation, classification, dimensionality reduction, and reinforcement learning. In predictive use cases, TweenMe generates Cognitive Digital Twins capable of forecasting the behavior of the PO represented by the LO.
- **User Experience and Collaborative Digital Twins:** TweenMe packages the Digital Model into an application that optimizes User Experience. For strategic decision-making, TweenMe can produce a Collaborative Digital Twin to explore “What if?” scenarios, aiding users in evaluating potential outcomes and guiding decisions.
- **Future Integration with Bidirectional Synchronization:** TweenMe’s outputs can be integrated into broader systems that enable bidirectional data synchronization between PO and LO. However, such integration depends on specific use cases, existing infrastructure, and data architecture, and therefore cannot be fully automated.

## 8.2 The Ontological Engine: What Makes TweenMe Universal

A Digital Twin is a model or abstraction that must be thoughtfully designed to represent the Physical Object with a level of detail appropriate to the specific business problem. What makes TweenMe a “universal” generator is not that it provides a single model for all use cases, but that its ontological engine provides the structured framework within which domain-specific Digital Twins are constructed.

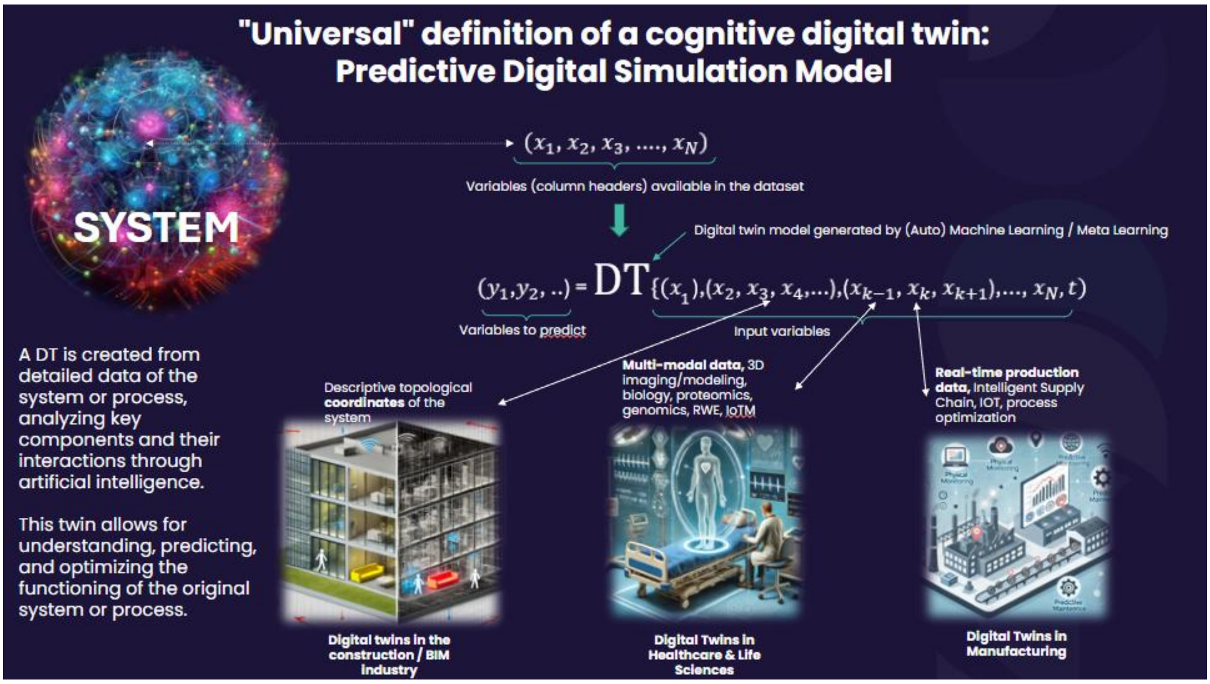

**"Universal" definition of a cognitive digital twin:**
**Predictive Digital Simulation Model**

SYSTEM

$(x_1, x_2, x_3, ...., x_N)$

Variables (column headers) available in the dataset

Digital twin model generated by (Auto) Machine Learning / Meta Learning

$$\underbrace{(y_1, y_2, ..)}_{\text{Variables to predict}} = \mathrm{DT}_{\underbrace{\{(x_1),(x_2, x_3, x_4,...),(x_{k-1}, x_k, x_{k+1}),..., x_N, t)}_{\text{Input variables}}}$$

A DT is created from detailed data of the system or process, analyzing key components and their interactions through artificial intelligence.

This twin allows for understanding, predicting, and optimizing the functioning of the original system or process.

Descriptive topological **coordinates** of the system

**Multi-modal data,** 3D imaging/modeling, biology, proteomics, genomics, RWE, IoTM

**Real-time production data,** Intelligent Supply Chain, IOT, process optimization

**Digital twins in the construction / BIM industry**

**Digital Twins in Healthcare & Life Sciences**

**Digital Twins in Manufacturing**

*Figure 2: Every use case is based on data which might be specific to an industry but still be exposed by a data store (SGBD, Graph Database)*

TweenMe's ontological layer formalizes the abstraction decision for each Digital Twin: which variables are modeled, how they relate to each other, what constraints govern the model's behavior, and what inference rules are applicable. This ontological structure serves as both the design blueprint and the validation framework. It guides variable selection (reducing dimensionality to clinically or operationally meaningful features), constrains the model space (ensuring that outputs respect domain knowledge), enables composability (allowing specialized sub-models to be integrated through shared ontological interfaces), and provides auditability (documenting why certain modeling choices were made).

While TweenMe significantly reduces the expertise required and lowers the marginal cost of creating new Digital Twins, aligning the ontological model with business goals remains essential for success. The ontology is the intellectual investment; the automated pipeline is the operational efficiency.

# 9. Illustrative Examples

## 9.1 Modeling the Clinical Benefit of Paracetamol: The Abstraction Problem

From a clinical standpoint, one could develop a straightforward model to predict the effects of acetaminophen (paracetamol) as a pain reliever and fever reducer. The clinical question is: given a certain dose of acetaminophen, what level of pain relief and fever reduction can be expected? A basic dataset would include dosage, patient-reported pain relief outcomes, and thermometer-recorded fever reduction.

However, an analytically minded researcher might explore what happens once acetaminophen enters the bloodstream: weak inhibition of COX enzymes (especially in the brain), modulation of the endocannabinoid system via the metabolite AM404, potential activation of TRPV1 receptors, and influence on serotonergic pathways. One could further model molecular binding through quantum chemistry simulations to estimate binding forces.

While such detailed modeling might be intellectually fascinating, it would not directly address the initial clinical question. This underscores the importance of defining the right question—and the role of the ontological layer in formalizing this choice. The ontology for a clinical efficacy Digital Twin would model dose-response relationships, patient covariates, and outcome measures; it would deliberately exclude molecular binding dynamics as outside the scope of the clinical question, documenting this exclusion as an explicit design decision rather than an implicit omission.

## 9.2 A Multimodal Model for Senology: Methodology and Validation

We developed a multimodal Digital Twin model beginning with an open dataset containing features extracted by a computer vision algorithm from 569 mammograms, each labeled as either benign or malignant. Our first step was to refine the dataset by identifying clusters of tumors with differing growth speeds, categorizing them as slowly growing or rapidly growing. We then created additional clusters based on tumor surface characteristics, classifying tumors as “smooth” or “rugged.”

Next, we built covariance matrices to integrate synthetic genomic and proteomic data, focusing on tumor attributes like growth speed, texture, and size. We applied a range of analytical techniques including Principal Component Analysis (PCA), Multivariate Linear Regression, Partial Least Squares (PLS) Regression, and Machine Learning models such as Support Vector Machines (SVM).

By examining these covariance patterns, we defined two new output variables—“tumor aggressiveness” and “treatment response”—which we linked to proteomic markers such as ER, PR, and HER2. These derived variables provide insights into tumor behavior and potential therapeutic outcomes, enhancing the model’s clinical utility beyond binary benign/malignant classification.

### Methodological Note: Synthetic Data and Validation

The use of synthetic genomic and proteomic data in this example warrants explicit discussion. Synthetic data generation rests on assumptions about the statistical distributions and inter-variable correlations of the real-world data it seeks to emulate. The properties of the resulting covariance matrices—and consequently the validity of downstream models—depend entirely on the fidelity of these assumptions.

Specifically, synthetic data generators may underestimate rare correlations, smooth over non-linear relationships, or introduce distributional artifacts that do not exist in the real-world data. When synthetic data is used to construct covariance matrices that in turn drive variable selection and model architecture, these artifacts can propagate silently into the final model. For this reason, any Digital Twin model built on synthetic data must undergo rigorous clinical validation against real-world patient cohorts before deployment. The synthetic data phase is a development accelerator—not a substitute for clinical evidence.

In TweenMe’s methodology, the ontological layer plays a critical role in mitigating these risks: by constraining the synthetic data generation to clinically plausible parameter spaces defined by domain expertise, and by requiring that derived variables (such as “tumor aggressiveness”) map to established clinical constructs with known biomarker correlations, the ontology acts as a guardrail against statistically valid but clinically meaningless model artifacts.

Beyond serving as a Digital Twin generator, TweenMe also creates and stores data models (ontologies) and insightful databases, which can ultimately be used to train cognitive digital models for integration into digital shadows or digital twins.

## 10. Conclusion: From Buzzword to Operational Discipline

The journey of Digital Twins from buzzword to operational concept requires clarity on multiple fronts. This article has argued that five elements are essential to this maturation.

**Definitional precision.** The concept of Digital Twin encompasses a spectrum from static models to fully interactive, AI-driven systems. A working definition must acknowledge this spectrum while being specific enough to guide implementation decisions. Situating this definition within the ISO and AAS normative landscape provides the interoperability context that isolated definitions lack.

**Ontological rigor.** The level of abstraction—what is modeled and what is deliberately excluded—is the most consequential design decision in any Digital Twin implementation. Formal ontologies provide the mechanism to make this decision explicit, auditable, and composable across systems.

**Cybersecurity as foundation.** The properties that make Digital Twins powerful—entanglement, persistence, memorization, servitization—are precisely the properties that create cybersecurity risks. Access control, data lineage, attack surface management, and auditability must be architectural foundations, not compliance afterthoughts.

**Honest ROI assessment.** The cautionary tale of predictive maintenance reminds us that Digital Twin business cases must benchmark against existing practices, not against their absence. Incremental value over current processes is the honest metric.

**Methodological transparency.** Where synthetic data, derived variables, or complex modeling pipelines are employed, their limitations and validation requirements must be explicitly stated. Clinical or operational deployment requires evidence that goes beyond statistical performance on development datasets.

The Digital Twin concept has the potential to move from trendy buzzword to operational discipline—but only if practitioners, vendors, and regulators approach it with the definitional precision, architectural rigor, and intellectual honesty that a genuinely transformative technology demands.

# References

Negri, E., Fumagalli, L., & Macchi, M. (2017). A review of the roles of Digital Twin in CPS-based production systems. Procedia Manufacturing, 11, 939–948.

Lyu, Z. (Ed.). (2024). Handbook of Digital Twins. CRC Press.

ISO 23247:2021. Automation systems and integration — Digital twin framework for manufacturing.

ISO/DIS 24576. Digital twin — Framework and requirements (under development).

Digital Twin Consortium (2022). Capabilities Periodic Table. Object Management Group.

Plattform Industrie 4.0 (2023). Asset Administration Shell Specification, Part 1–3.

Grieves, M. (2014). Digital Twin: Manufacturing excellence through virtual factory replication. White Paper, Florida Institute of Technology.

Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405–2415.

European Parliament (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act).

European Parliament (2017). Regulation (EU) 2017/745 on medical devices (MDR).

European Parliament (2016). Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (GDPR).

Bourtoule, L., et al. (2021). Machine unlearning. IEEE Symposium on Security and Privacy (SP), 141–159.