

Beyond Copies

Digital Twins as Information Artefacts

Luca Biccheri^{1,*}, Roberta Ferrario^{1,†}

¹Laboratory for Applied Ontology, ISTC-CNR, Via alla Cascata 56/C 38123 Povo di Trento, Italy

Abstract

Digital Twins (DTs) are often described as virtual replicas or copies of physical systems. However, we argue that this view not only obscures key aspects of what DTs actually do but also fails to capture their underlying ontological rationale. In this paper, we offer a foundational account of DTs based on DOLCE, arguing that these technologies embed intentional representations at their core. We then qualify DTs as information artefacts shaped by stakeholders' goals and manifesting qualities, which we call 'differential qualities', that have causal relevance. These are essential for accomplishing tasks typically performed by DTs in engineering, such as diagnostics, optimisation, and simulation.

Keywords

Digital Twin, differential quality, information artefact, aboutness, epistemic causality, foundational ontology

1. Introduction

Digital Twins (DTs) are often informally presented as kinds of virtual or digital models of physical objects or systems, created to support activities like simulation, real-time monitoring, predictive analysis, visualisation, and operational control. While DTs have recently garnered attention from scholars working or interested in ontology at large (see, for instance, [1]), to the best of our knowledge, few and preliminary formal ontological accounts have been proposed that explicitly address what kind of entities DTs actually are (cfr. [2] and [3]).

In this paper, we approach the conceptualisation of Digital Twins (DTs) from a foundational ontology perspective, investigating the primitives required for their formalisation. The foundational ontology adopted in this work is DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering) [4, 5, 6], which has already been successfully applied to DT-related technological domains, such as artefacts and functions [7, 8, 9, 10]. The choice of leveraging DOLCE rather than other foundational ontologies is further motivated by the fact that in the latest years several works have been produced that were focused on intentionality, with a rich conceptualisation of the notion of aboutness. We shall gradually introduce the reader to the set of DOLCE entities necessary to characterise DTs in due course.

At first glance, as mentioned, one could be inclined to model DTs merely as copies or replicas of physical entities. However, as we will argue, although it is weaker than the logical identity [11], the notion of copy is overly rigid and ultimately insufficient: not only does it obscure key aspects of what DTs actually do, but it also fails to capture their underlying ontological structure.

Rather than viewing DTs as kinds of copies, we propose to understand them as specific representations of physical endurants, i.e., *intentional* representations, which allows us to identify them as a subclass of technical artefacts, namely information artefacts, whose features depend (in part) on the goals of various agents involved as stakeholders (e.g., designers, engineers, users, etc.).

This goal-oriented approach is further complemented by a causal view of DTs. Regardless of the physical entities related to DTs, we claim that there exists a subset of features that consistently matter

Proceedings of the Joint Ontology Workshops (JOWO) - Episode XI: The Sicilian Summer under the Etna, co-located with the 15th International Conference on Formal Ontology in Information Systems (FOIS 2025), September 8-9, 2025, Catania, Italy

*Corresponding author.

†These authors contributed equally.

✉ luca.biccheri@istc.cnr.it (L. Biccheri); roberta.ferrario@cnr.it (R. Ferrario)

ORCID 0000-0001-8002-7550 (L. Biccheri); https:0000-0002-9175-3096 (R. Ferrario)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

to identify them. We refer to these as ‘differential qualities’. These properties, by having causal import, *make a difference* in how an entity changes over time, and therefore are relevant for understanding, simulation, optimisation, and diagnostic tasks that DTs typically perform.

This perspective is particularly important in engineering contexts, where DTs can be used to evaluate and ensure the dependability of systems under a variety of conditions [12]. Although DTs’ applications span disparate domains, from agriculture to biology [13], here we restrict our analysis to engineering examples (without loss of ontological generality, so to speak).

A well-known distinction present in the literature is that between Digital Twin Prototype and Instance. While the former is a description of a prototypical physical artefact and is aimed at then producing the physical twin, the latter describes a specific, already existing, physical twin, to which it remains linked throughout its life¹.

However, although one can speak of DTs at the type level, thus addressing their conceptualisation to characterise the prototypical class of Digital Twins, in the current paper we prefer to limit our analysis to the token level. Beginning with individual tokens allows us to focus on practical examples and enables a step-by-step process of abstraction.

The paper is structured as follows: in **Section 2**, without aiming to disentangle the ontological complexity of the notion of copy [14, 11], we briefly highlight its shortcomings when applied to the analysis of Digital Twins. **Section 3** introduces the class of technical artefacts and argues that the primary relations between Digital Twins (DTs) and the physical entities they represent can be understood in terms of aboutness or intentionality in one direction (DT to Physical Twin – PT) and existential dependence in the other (PT to DT). This supports the view that DTs are best characterised as information artefacts. The notion of ‘differential quality’ is presented in **Section 4**, and is then used to provide an exploratory formal characterisation of Digital Twins in **Section 5**, building upon the theory of aboutness developed within a goal-oriented framework in [15], which is also employed here to provide a definition of ‘information artefact’. In **Section 6**, we provide some concluding remarks.

2. The Pitfalls of the ‘Copy’ Notion in Digital Twins

The label ‘Digital Twin’, which originated in the field of engineering, does not refer to a specific technology. Rather, it encompasses a wide range of computational disciplines, applications, and techniques, including computational logic, AI algorithms, ontology-based data modelling, the Internet of Things (IoT), and sensors [16].

Ontologically, in a nutshell, one might say that Digital Twins (DTs) are technical artefacts² *representing* a physical entity. In this respect, they are not far from abstract models or dynamic simulations of the world, which have been widely employed in science and technology for decades [13]. However, the conceptualisation of DTs often presents them as ‘realistic’ or ‘high-fidelity’ representations of physical entities, thus departing from abstract models. Moreover, far from being just simulations, DTs are typically linked in real-time to physical entities from which they receive data, which are then used for monitoring, controlling, or decision-making activities that influence the behaviour of such entities.

Not surprisingly, given the technological complexity of the topic, we lack a universally accepted definition. Yet, the domain of discourse about DTs can be roughly understood as involving three main entities [13]: (1) a physical entity (2) a technical artefact (3) a relation between these entities.

Now, at a high level of abstraction, one might intuitively argue that, since each physical entity has a corresponding digital counterpart, that is, its ‘Digital Twin’, the very term itself suggests a kind of ‘copy’ relation is involved. However, this option is neither feasible nor desirable.

To start with, one might naively wonder what features and how many, should be identified in order to claim that a DT is a copy of a physical entity. Trivially, storing *every* single detail of the physical entity is often impractical, if not impossible (epistemologically speaking). This immediately underlines that we

¹It seems well-established in the literature that, regardless of whether it is created before or after its digital counterpart, the existence of a physical twin is a necessary condition for the existence of the digital twin.

²This notion will be explored in more detail in Section 3.

are not looking for a strict notion of copy or a *replica* (e.g. a PDF copy) [14]. Typically, we aim to capture only the most informative features of the physical entity. In this regard, the expression ‘high-fidelity representation’ used in DTs literature might relate to those aspects of the physical entity relevant to the intended task, not replicating every detail (e.g. a turbine’s Digital Twin for failure prevention may prioritise sensor data, including vibration, temperature, and pressure, over structural ones).

Furthermore, real-world social and biological systems are inherently emergent, requiring open and non-linear models that move beyond simple copying processes [17]. Though, in this paper, we limit our analysis to the engineering field, DTs are not confined to such domain, but can also be applied to entities in the Life Sciences and Earth Sciences, whose modelling typically involves working hypotheses and simplifications [13]. Moreover, as said, DTs are not only meant to mirror physical objects but also to undergo, for instance, virtual testing and optimisation; this means that some computational operation is performed on them in order to obtain improvements on the physical entities. If simulations reveal an enhanced design or better performance characteristics that the physical object lacks, the Digital Twin may serve as a blueprint for potential modifications to its physical counterpart. In this dynamic process, it becomes unclear what is truly a copy of what³. Perhaps, however, if one wishes to account for the relation between Digital Twins and their *Physical Twin* in terms of copying, a weakened notion of *copy* is required. To this end, alternative relations such as *approximate* and *inexact* copies have been introduced in the philosophical debate on artefacts. Nevertheless, these notions were specifically formulated to address logical properties among *computational* artefacts, not artefacts in general. Moreover, these are not meant to account for the relation between the digital and physical worlds [11]. To explore this relation, it is necessary to account not only for formal and, of course, physical properties, but above all for the intentional dimensions involved [7, 8, 14].

Unlike a physical replica (e.g. biologically identical twins), a DT is goal-designed, i.e. it is created with specific purposes in mind, typically: diagnostic, prediction, optimisation. This agent-dependency entails that the Digital Twin selectively represents *only* those features of the physical entity that are relevant for achieving the intended goals attributed to the related DT. Its identity is therefore determined by a goal-oriented process. This marks a crucial ontological distinction: DTs are not merely what they are, but what they are *for*, a dimension absent in standard notions of replication and, to some extent, also of copying. Regarding the latter, the intentional dimension is involved insofar as the act of copying entails an artefactual process of creation or modification, and not to the idea of copy itself. We claim therefore that while the copy relation may seem intuitive, it is too rigid to fully account for the complexity of Digital Twins.

Along these lines, Korenhof et al [2021] warn us that framing the conceptualisation of DTs solely in terms of ‘realistic’ representations of a physical entities, and indeed as a kind of copy, obscures the ‘prescriptive’ dimension of DTs, which are a-priori designed to monitor or improve the physical counterpart. This is not a neutral process, of course. As a result, a DT always contains something that is of a *different ontological kind* w.r.t. the physical entity it refers to. For example, the numerical models that simulate the growth of a grapefruit tree are not part of the actual grapefruit tree. These models are instead used to cope with a physical reality that is perceived, if not as faulty, at least as something in need of optimisation. Ultimately, this leads to a *reification* of the DT: it is treated as a full-fledged entity that, for all practical purposes, substitutes the physical object.

For instance, a farmer who uses a DT to monitor and automate the feeding of cows, thus optimising the process, engages with the DT as an interface. In the worst-case scenario, this might alienate the farmer from the actual physical activity, delegating much of the labour to an automated system driven by the DT [13]. This also affects human behaviour (as technology typically does) and introduces political

³Actually, the force of the statement depends on the notion of copy employed. For instance, if one assumes that, if x is a copy of y , then y serves as a kind of *model* for x , it is always possible to discern the ‘original’ from the ‘copy’ [14]. In this respect, the copy relation is neither symmetric nor reflexive, and thus it is not an equivalence relation, contrary to what Tzouvaras [1993] maintained [11]. However, even in this case, the notion of *copy* clashes with the peculiar relation between DTs and their physical counterparts. On the one hand, the physical entity serves as a model for designing the DT; on the other hand, the DT itself can be regarded as a sort of ideal model in which to experiment with improvements to be applied to the physical world.

and ethical matters (e.g., data ownership, transparency, explicability) related to different stakeholders' goals and needs, which ultimately affect the design of the DTs themselves.

While representations are a crucial component in discussions about DTs, they coexist with goals, prescriptions, and actions enabled or required by such technologies. This ultimately leads Korenhof et al [2021] to qualify DTs as 'steering techniques', aligning them with the traditional cybernetic paradigm in which technical systems are used to steer or guide behaviour. In this view, DTs are employed to realise specific goals in the world through feedback mechanisms that provide information and influence action.

The informational dimension introduces an additional nuance to the concept of 'intentional', that must be distinguished from goal-oriented aspects or, in other words, from agents' intentions⁴. An intention or goal⁵, conceived as a mental state, is merely one among several such states, e.g. beliefs, desires, and perceptions, that exhibit the property of 'intentionality' or 'aboutness', which has to do with the very matter of meaning, whether mentally or linguistically conceived [20, 21]. DTs, in addition to embedding agents' intentions, also carry (albeit indirectly) intentionality. This peculiar characteristic, we argue, sets DTs apart from other kinds of technical artefacts. It is precisely this dimension that is at stake when we claim that a Digital Twin *represents* its physical counterpart.

From what has been said, we can refine our ontological understanding of DTs by drawing on two main assumptions:

i) First, we can view the representation (i.e. aboutness) relation as the fundamental link between a physical entity and its DT. DTs refer to their physical counterparts in a meaningful way and carry information about them.

ii) Second, alongside this informational aspect, we must adopt a goal-oriented perspective. This accounts for how and why DTs, as a specific type of technical artefact, are used as 'steering' technologies.

In the following sections, we will explore these assumptions in greater depth in order to provide a basic formal understanding of Digital Twins.

3. Introducing Information Artefacts

There is no doubt that DTs represent a complex technology, made up of multiple layers of material and code. Some of these layers are directly perceivable by humans, such as the output displayed on a screen, while most remain concealed beneath intricate arrangements of physical components, hardware engineering, and logical systems (from high-level programming languages all the way down to the electrical properties of copper and silicon) [22, 13].

However complex such technology may be, ontologically we can distinguish two main layers: the material and the artefactual, with the former constituting the latter. To account for this, we draw on the formal theory of artefacts developed in [7] and further contextualised in [8]. At its core, the theory emphasises that an artefact comes into being when an agent intentionally selects a physical object and attributes specific capacities to it. For example, a piece of wood becomes a bench because someone has intentionally (i.e. by adopting an intention) selected that object and attributed to it the capacity to support weight. Of course, in the case of DTs, the story is more complex. Here, one must not only select but also modify or arrange a specific amount of matter (e.g. silicon) based on some of its physical qualities (e.g., electrical conductivity), so that the resulting artefact can exhibit behaviours (i.e., computational operations enabled by the flow of electrical current through circuits) that the original physical matter alone does not possess. Typically, this transformation marks the creation of a technical artefact: including cars, bikes, laptops and, more generally, any devices designed to perform certain tasks [8, 23, 10]. However, fundamentally, artefacts are not defined merely by their physical capabilities,

⁴Unless explicitly stated, the twofold meaning of 'intentional' should be understandable from the context.

⁵Note that in the literature on BDI (Belief-Desire-Intention) models, intentions are conceived either as primitive mental states or as primitives used to define goals themselves. In the latter case, intentions are treated as synonymous with goals. Alternatively, an objectivist view considers goals as externally provided 'recipes' for agents to follow [19]. In this paper, we adopt the mentalist perspective.

but, above all, by the capacities attributed to them⁶. In the case of DTs, such attributed capacities might include a number of different dimensions, including reliable data transmission, crash simulation, diagnostic, etc.

The topic of capacities attributed to artefacts is well suited for our purpose, as it helps illuminate the steering role of DTs in achieving human goals. Moreover, by adopting a multiplicativist approach [7], according to which an artefact is a distinct entity from the physical object that constitutes it, the artefactual theory we embrace also accounts for why DTs' representations of physical entities introduce elements (e.g., numerical models) that cannot inhere to their physical counterparts. However, this analysis alone is not sufficient to fully characterise DTs as the specific kind of technical artefacts they are. In fact, as mentioned in the previous section, DTs are intended to represent their corresponding physical entities, and there is nothing inherent in the general ontological framework of artefacts we have just introduced that justifies such representational power.

To account for this, we need to appeal to the so called 'information entities'. This is a general label encompassing a variety of entities, ranging from semi-abstracta to meanings, ideas, and documentary entities. In order to highlight similarities and differences among these entities, Sanfilippo [2021] presents a comparative table with column indices reporting on some ontological dimensions, such as existence in time and space, criteria of unity, generic dependence, and aboutness. For our purposes, it is worth noting that all the theories examined in [24] appeal to the notion of aboutness to characterise information entities, though without committing to a particular stance or providing a formal account⁷. This lack of formal studies is not limited to the domain of information entities. In fact, the concept of intentionality or aboutness is primarily treated as a primitive and largely underanalysed in the applied field of Formal Ontology [26], despite being a fundamental notion across multiple domains.

Expressions such as 'reference', 'intentionality', 'aboutness' have considerable semantic overlap and are sometimes used interchangeably for all means and purposes. However, depending on whether the emphasis is on language or mind, we see a prevalence of terms like 'reference' and 'aboutness' in the former case, and 'intentionality' in the latter. For instance, according to Yablo [2014], aboutness "is the relation that meaningful items bear to whatever it is that they are on or of or that they address or concern.". Conversely, intentionality is, as mentioned, typically introduced as the property of mind, and more specifically a kind of power or ability, to be *about*, *represent* or *stand for* something [21]. Therefore, it is no surprise that the notion of aboutness is also at stake in the characterisation of information entities, insofar as an entity carries information precisely by virtue of being representational.

In this respect, the notion of aboutness can be employed to capture the link between DTs and their physical counterparts. We will discuss this in Section 5, by leveraging the formalisation and the theory of aboutness offered in [15] and by adapting it to address the semantical aspect of DTs.

Before concluding this section, let us raise another important issue. Just as there is no single technology that fully identifies a DT, there is likewise no fixed set of tasks it is meant to perform. However, among the most common applications are diagnosis, simulation, and optimisation [13]. When representing the physical counterparts of DTs, not all of their properties (e.g. the colour of an object) are relevant for accomplishing these tasks. Presumably, only those qualities that affect or modify the behaviour of the object are of interest. We will refer to these kinds of properties as 'differential qualities'. We will introduce them in the next section and demonstrate its relevance to the ontological modelling of DTs in Section 5.

⁶In [7], capacities are categorised as DOLCE qualities. Two types of capacities are introduced: physical capacities and attributed capacities that, differently from the former, depend on an agent's intentions at the moment the artefact is created. The discourse on capacities has been further developed in recent work on DOLCE from an engineering perspective, although approached differently [10].

⁷In this respect, the only exception is the treatment on aboutness offered by Ceusters and Smith [2015].

4. On the Role of Differential Qualities within the Engineering Domain

Digital Twins (DTs) are deployed in many heterogeneous domains, but in this paper we shall mainly focus on engineering contexts, where causal and probabilistic aspects of artefacts' behaviour are central. In this section, we introduce the relation we dubbed 'isDifferentialQualityFor', specifically devised to capture such aspects, and developed in a paper currently under review [12]⁸. In the next section, we will then employ this relation for providing an exploratory formal characterisation of DTs.

To this aim, let us briefly contextualise the main theory of causality we appeal to, and the reasons behind such an endorsement in light of our objectives. First of all, we aim to articulate the intuitive idea of 'causal contribution' between properties, in alignment with the technological domain under investigation. In this respect, a probabilistic approach to causation appears best suited to account for the behaviour of DTs, e.g. simulating physical artefacts in complex environments often involves dealing with multiple sources of uncertainty (including stochastic processes such as sensor noise, unpredictable external influences, and incomplete or imprecise data about the system's components, the environment, etc).

Probabilistic accounts of causality belong to a broader family of approaches commonly referred to as 'difference-making' theories of causality. These are often contrasted with another group of approaches, known as 'mechanistic theories'. The distinction lies in their differing uses of causal claims: the difference-making approach serves an inferential purpose and applies to types of entities, while the mechanistic stance is aimed at providing explanations and is used for token entities [27]. The mechanistic stance focuses on identifying the physical processes that explain the link (i.e., the mechanism) between causes and effects in individual cases and through inductive reasoning. For example, by connecting the insurgence of John's lung cancer with his being a heavy smoker [28]. In contrast, the difference-making perspective emphasises the information required to infer effects from causes, e.g. by intervening on variables representing causes while controlling for variables representing effects, such as when constructing a Causal Bayesian Network from a dataset [29]. In brief, difference-making involves the idea of 'change', i.e. how is that the presence of causes brings about changes in their effects [cf. 28].

Hence, within the engineering context, both probabilistic and mechanistic factors are often required, as engineers are interested in both the predictive and explanatory uses of causal claims. From this standpoint, the epistemic theory of causality [29] is well-suited, as it unifies probabilistic and mechanistic evidence into a coherent framework of rational causal belief. Though being an epistemic theory (i.e. not committed to causality as a full-fledged entity), this view does not deny the existence of facts serving as truthmakers for causal beliefs [28, 30]. Rather, it openly acknowledges the evidential grounding of causal claims in the agent's perspective.

In this framework, the relation `isDifferentialQualityFor` helps formalise how certain qualities of artefacts contribute (both probabilistically and mechanistically) causally to changes observed in other qualities. More specifically, referring to DOLCE's categories⁹, it links two qualities q and q' (inhering in an

⁸Readers more interested in the philosophical aspects might recognize in the notion of differential qualities certain echoes of the metaphysics of dispositions. This is a delicate point that deserves clarification. First, when we speak of differential qualities, our interest lies primarily in the type of relationship that exists between two (or more) qualities, and not in the nature of those qualities themselves. The focus is on how such qualities behave within the dynamics of interaction between the endurants to which they belong. In [12], differential qualities are used, together with capabilities [10], to characterize powers. By definition, powers are relational qualities involving certain capabilities, which in turn are differential qualities with respect to the powers themselves. It is important to note, however, that powers are neither reducible nor limited to capabilities (i.e. powers may also depend on other differential qualities that are not capabilities), and that the definition of capability we adopt to define powers is modally characterized. Thus, the notion of power, as we use it, differs significantly from the standard notion of disposition in metaphysics. In the context of this work, however, we do not engage with the notion of power; our focus lies solely on leveraging the concept of differential quality for engineering modelling purposes.

⁹Differently from universals, qualities in DOLCE are *particulars* that *inhere* in other entities. The inherence relation is functional: a quality is specific to its bearer and existentially dependent on it. For instance, John's weight is a quality that inheres exclusively in John throughout his life and ceases to exist when John does. Hence, DOLCE qualifies such qualities as 'individual'.

endurant each¹⁰), where the bearer of q participates in an event x and the bearer of q' in an event y , with the former being the cause of the latter. Two further relations are provided in [12] to characterise the relation of `isDifferentialQualityFor`, i.e. the relation `isProbabilisticallyDependentOn`, holding between types of events, and the relation `isMechanismOf`, holding between particular events.

Ontologically, this setup asserts that when two events are taken as evidence for a causal claim, the time-dependent value change in the qualities of the participating endurants is also causally related. That is, a change in quality q at time t implies, within some probabilistic constraint and against some relevant mechanistic explanation, a change in q' at time t' . For example, increased humidity (a quality q of the environment) participating in a *vaporising* event can reduce the accuracy (q') of a sensor participating in a subsequent *error reading* event [32].

With the introduction of differential qualities, we are now ready to ontologically qualify Digital Twins.

5. Foundations of Digital Twins

In this section, we introduce a real-world example in the engineering domain to help understand the formalisation we will provide. We already mentioned that DTs are rooted in real-time data and grounded in physics-based simulations, making them powerful tools for early fault detection and decision-making.

The example is taken from Siemens' use of DTs in monitoring large rotors in generators and compressors¹¹. When one of such rotors fails, downtime can cost up to five million euros per day. Even scheduled maintenance results in weeks of inactivity (mainly due to safety reasons) and requires expert analysis.

In contrast to this traditional maintenance scenario, the Digital Twin-based approach offers considerably greater flexibility. Every rotor has a Digital Twin created at the design stage that runs in parallel with the actual machine, constantly simulating ideal sensor readings (e.g., vibration data). When discrepancies arise between simulated and real data, the DT provides diagnostic. By simulating various fault scenarios (like imbalance forces or dust buildup), adjusting parameters (force location, magnitude), and minimising the difference between real and simulated data, the DT identifies the cause and location of the fault and allows enough time to schedule repairs.

Let us now introduce the DOLCE framework that will enable us to ontologically account for the practical applications of Digital Twins, including the industrial scenario just described. To start with, we recall that, of course, Digital Twins are artefacts, as they are the product of human intentional activities. So far, in DOLCE only *material artefact* has been formalised in [7], but successive works have specialised such notion into that of *technical artefact* [8].

In our intuition, Digital Twins are still more specific kinds of artefacts, namely, able to carry information. We shall call this kind 'information artefact' and look for an ontological characterisation.

To this aim, we build on previous work, presented in [15]. We consider the notion of *aboutness* as a complex relation composed of two simpler relations: *aboutness*₁ and *aboutness*₂. The former is a ternary relation involving an agent's (private) mental state x , a concept y , and a time t . Concepts in DOLCE are to be intended as reifications of properties, those that must necessarily be possessed by an entity to be classified under those concepts, the latter being not private but defined by a community of

In this respect, they resemble tropes that extend in time and undergo change. We shall therefore refer to *dependability* as a kind of individual quality. An individual quality is distinct from its value, called a *quale*: this is the position of an individual quality in a structured *quality space*. Quality spaces are inspired by Gärdenfors' [2004] conceptual spaces. Qualities of the same kind take values within the same space, enabling their comparison via qualia. For instance, if the colour of a car and the colour of a book cover exhibit the same shade of blue, their respective qualities occupy the same position (quale) in the quality space of colour. Thus, they are distinct qualities sharing the same quale, i.e. value.

¹⁰In the more general case, q and q' inhere in different endurants, but it could also be the case that they inhere in the same endurant. In such case, the change in one of its qualities causes the change in another of its qualities, as when a swan baby, becoming adult, turns from gray to white).

¹¹<https://www.siemens.com/global/en/company/stories/research-technologies/digitaltwin/error-diagnosis-digital-twin.html>

agents. Concepts can classify entities at a certain time¹², but they can also exist without classifying anything. For those concepts that succeed in classifying some entities at a given time, we use in our framework also the relation of *aboutness*₂, which is a ternary relation involving an agent’s mental state x , an entity (instance) z , and a time (moment or period) t at which such an entity is classified by that very concept.

So, the overall picture is that when agents are in a certain mental state (like belief or intention), such mental state is always directed towards a concept. The fact that a concept is always involved captures the common-sense idea that thinking is always *thinking about something*. In our framework, the relation of being directed towards a concept is called *aboutness*₁ and, while this relation holds, the concept may be called the ‘intentional content’ of the mental state. Now, in some cases the concept does not classify any entity or, in other words, the content does not denote any entity, such as for fictional entities. In such cases, *aboutness*₁ is the only relation that holds; in all other cases, there will also be a relation of *aboutness*₂ between an agent’s mental state and an individual entity in the ‘world’ or ‘domain’ of interest, that we shall call ‘intentional object’, which is then classified by the concept towards which the mental state is directed.

It is worth noticing the two senses of aboutness formalised in [15] are slightly different from those needed to tackle the sense in which an information artefact may be ‘about’ something. For one thing, while in [15] the focus was on intentional content and object, here what is at stake are information content and object; in a sense, provided that one can clearly distinguish between the two, we are moving from the mental to the semantic or linguistic realm.

Now we can relate the theory of aboutness to the characterisation of information artefacts. The basic idea is that an information artefact is a technical artefact that *encodes* information. However, for such information to be meaningful, it must be ‘cognitively graspable’ by an intentional agent, therefore here is where the notion of *aboutness*₁ comes into play. This is tantamount to say that the ‘aboutness’ is not so much intrinsic to the concepts encoded by the informational artefact per se, but rather arises indirectly through the agent’s cognitive intervention [34]. The encoding is possible insofar as the physical object (or part of it) that constitutes the artefact is arranged in a way that makes it cognitively accessible; in other words, insofar as it functions as a symbol (e.g., ink patterns in books, voltage states in computer gates). We stress that, generally speaking, many information artefacts only need to appeal to *aboutness*₁, that is, a relation to content; e.g. in the case of books about fictional characters. Likewise, there are cases in which such content, through the mediation of an intentional agent who entertains some mental attitudes towards it, can refer to something in the world. Those are the cases in which *aboutness*₂ is at stake¹³. Examples of such information artefacts are medical records, dashboards in industrial settings and Digital Twins as well.

To provide the definition of information artefact, we introduce the primitive binary relation Encodes, which holds between the technical artefacts and an ‘informative content’(concept) at a certain time¹⁴. We have therefore the following:

Axiom 1: $\text{Encodes}(x, y, t) \rightarrow \text{TechArt}(x) \wedge C(y) \wedge T(t)$

Definition 1: $\text{InfoArt}(x) \leftrightarrow \text{Encodes}(x, y, t) \wedge \exists z(\text{MS}(z) \wedge \text{About}_1(z, y, t))$

¹²In the seminal work on concepts [33], these were introduced to classify endurants. However, in [15] the classification relation is extended to perdurants as well.

¹³There is no consensus in the philosophical literature regarding the ontology of the entities that serve as the referents of intentional mental states, commonly labeled as intentional objects [19]. Depending on the philosophical assumptions, these entities may range from particulars to properties, events, or even complex entities such as states of affairs. In our case, for practical modelling purposes and assuming DOLCE as the background ontology, the physical twin, *qua* intentional object, could be conceived in various ways: as the ‘whole’ composed of the endurant and the perdurant tied by the participation relation (e.g., the robotic arm performing a pick action), as just the arm (or more precisely, some of its qualities, such as weight), or even as the event itself (e.g., picking up a bottle at a given time). Ultimately, much depends on the modelling focus and purpose. However, the key point is that our proposal always allows for the identification of an intentional object corresponding to the physical counterpart of the digital twin.

¹⁴The information that an agent grasps while interacting with a technical artefact (like reading a book) may change in time.

Axiom 1 minimally constrains the arguments of the relation Encodes, while Definition 1 states that x is an information artefact iff x is a technical artefact that encodes a concept y at a certain time t and there exists an agent's mental state z that is About₁ concept y at t ¹⁵.

The following step is that of characterising Digital Twins (DTs) as a subclass of information artefact. We mentioned (Section 2) that DTs are used in goal-oriented contexts, underlining their being 'steering technologies'. In addition, we introduced a class of qualities (Section 4), called 'differential qualities' helpful to identify features of the physical entity useful for DTs' tasks (e.g. simulation, optimisation, etc). Now let us resume the example about the Siemens' rotor to contextualise the ontological scenario.

Consider a system in which the DT of a rotor includes a parameter representing the rotor's orientation angle. This parameter is modeled as a quality q inhering in the DT. The DT participates in a simulation event e , during which it predicts friction due to dust accumulation. As a result, it updates the value of the quality q' (i.e., the angle) that inheres in the actual rotor. This update causally triggers an adjustment in the physical rotor's orientation, mitigating dust build-up. Thus, the quality q of the DT can be said to be a *differential quality* for q' related to the Physical Twin. In addition, the simulation event e serves as a satisfier¹⁶ of an *instrumental goal* g . The occurrence of e brings about another event e' , i.e., a decision-making event that, in turn, satisfies an *end goal* g' . This complex event ultimately results in scheduling a cleaning intervention without requiring system shutdown, thereby achieving the end goal of minimising operational costs.

Now, in order to formalise the causal dimension, thanks to which the DT modifies a quality of the actual rotor, we leverage the primitive relation of 'causal contribution' introduced in [10] holding between events, by constraining it with the relation isDifferentialQualityFor [12]. Hence, we have the following:

Axiom 2: $\text{causalContr}(e, e') \rightarrow \text{EV}(e) \wedge \text{EV}(e') \wedge \exists q, q' (\text{Q}(q) \wedge \text{Q}(q') \wedge \text{isDifferentialQualityFor}(q, q') \wedge \text{ED}(x) \wedge \text{ED}(x') \wedge \text{I}(q, x) \wedge \text{I}(q', x') \wedge \text{PC}(x, e, t) \wedge \text{PC}(x', e', t'))$

Axiom 2 states that if e causally contributes to e' , then both e and e' are events, and there exist qualities q and q' such that q is a differential quality of q' (a change in the value of q brings about a change in the value of q'). These qualities inhere in the endurants x and x' , respectively, with x participating in e and x' participating in e' . Now we introduce the goal-oriented dimension which, together with the relation of causal contribution, allows us to qualify what a Digital Twin is:

Axiom 3: $\text{dgTwin}(x) \rightarrow \text{InfoArt}(x) \wedge \text{PED}(w) \wedge \text{SD}(x, w) \wedge \text{goal}(z, y, t) \wedge \exists y' t' (\text{goal}(z, y', t') \wedge t \prec t' \wedge (\exists e. \text{sat}(e, y, t) \wedge \text{PC}(x, e, t) \rightarrow \exists e'. \text{sat}(e', y', t') \wedge \text{PC}(w, e', t') \wedge \text{causalContr}(e, e')))$

Axiom 3 characterises a Digital Twin x as an information artefact that is specifically dependent on a physical endurant w (i.e., any material entity, including other artefacts) for its existence. Moreover, the DT x participates in an event e which satisfies an agent's (instrumental) goal z concerning a concept y at time t . Event e causally contributes to event e' , which satisfies the agent's (end) goal related to a concept y' at time t' , where t precedes t' . In addition to qualifying Digital Twins as steering technologies that exert a causal impact on their physical counterparts, Axiom 3 also emphasises their existential dependence on the corresponding Physical Twins. Finally, we also provide a characterisation of the notion of a Physical Twin:

Axiom 4: $\text{phyTwin}(x) \rightarrow \text{PED}(x) \wedge \exists y, z, w, t (\text{dgTwin}(y) \wedge \text{Encodes}(y, z, t) \wedge \text{About}_1(w, z, t) \wedge \text{About}_2(w, x, t))$

¹⁵Notice that t must be the same in the two relations in Definition 1, because in the periods in which one of the two conditions fails, the technical artefact still exists, but it does not convey information, thus the information artefact ceases to exist.

¹⁶In [15] a satisfier is defined as an (existing) event that satisfies the content of the agent's goal.

Axiom 4 says that if a physical endurant is a Physical Twin, it is the reference of an information content encoded by the Digital Twin. This should render the fact that an entity is a Physical Twin only when there is at least an agent holding a mental state directed to an information content encoded by the Digital Twin that classifies (is made concrete by) such entity. In other terms, together with the existential dependence of the Physical Twin from the Digital Twin, it is also the case, roughly speaking, that an entity becomes (plays the role of) a Physical Twin, only when it has a Digital Twin.

6. Conclusions

In this paper we argue against the use of the concept of copy to account for the ontology of Digital Twins. While intuitively appealing, the notion of copy proves too rigid and obscures the prescriptive nature of these technologies. Instead, we propose a goal-oriented approach which, supported by a class of qualities we dubbed *differential qualities*, highlights how DTs function as paradigms for modifying, controlling, and ultimately guiding the behaviour of their physical counterparts. Moreover, although a relation of existential dependence holds between the DT and the physical entity (see Axiom 3), we contend that *aboutness* is the main relation linking the two. It is through its representational character that the DT carries information about its corresponding Physical Twin. This, in turn, allows us to characterise DTs as a particular kind of technical artefact, namely, *information artefacts*.

Unlike some common information artefacts, such as a science fiction book, in which the encoded concepts need not refer to something actual, DTs belong to a specific class of information artefacts that are always connected with a material entity, in the specific sense of *aboutness₂* introduced in Section 5. This is not surprising, given that DTs are designed to causally affect the physical world: there must always be a physical counterpart upon which they act, either directly (e.g., when the DT autonomously performs an action) or indirectly (e.g., through human intervention). Our analysis ultimately reveals that, among all types of information artefacts, DTs are arguably among the very few that also exert a causal impact on the world they represent. Finally, a further interesting feature of DTs – and of information artefacts in general – that we have highlighted in our framework is that not only the intentions of the designer are relevant for their characterisation, but also of the agents that can grasp the information that the artefact encodes.

In this paper, we have focused specifically on the concept of the DTs at the token level. However, a comprehensive ontological account cannot disregard the type level, that is, the prototypical dimension that underpins the design, modelling, and implementation of Digital Twins as a class of artefacts. In this respect, the present work should be considered as a primer that serves as a base for a future type-oriented level analysis.

Acknowledgments

This research has been funded by the Project PRIN 2022 SMARTTEST (Project nr. 202223E8Y4X), awarded by the Italian Ministry of University and Research (MUR). Roberta Ferrario has also been supported by the Project PRIN2020 BRIO – Bias, Risk and Opacity in AI (2020SSKZ7R), funded by MUR and by the project FAIR (Project nr. PE 00000013), funded by MUR.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: grammar and spelling check, paraphrase and reword. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] E. Karabulut, S. F. Pileggi, P. Groth, V. Degeler, Ontologies in digital twins: A systematic literature review, *Future Generation Computer Systems* 153 (2024) 442–456. URL: <https://www.sciencedirect.com/science/article/pii/S0167739X23004739>. doi:10.1016/j.future.2023.10.014.
- [2] C. Barros, R. Salles, E. Ogasawara, G. Guizzardi, F. Porto, Requirements for an ontology of digital twins, in: *CEUR Workshop Proceedings*, volume 2941, 2021. URL: <http://ceur-ws.org/Vol-2941/>, publisher Copyright: © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
- [3] F. Wilson, R. Hurley, D. Maxwell, J. McLellan, J. Beverley, Foundations for digital twins, <https://arxiv.org/abs/2405.00960>, 2024. Preprint available at arXiv.
- [4] C. Masolo, A. Oltramari, A. Gangemi, N. Guarino, L. Vieu, Wonderweb deliverable d18: Ontology library, Laboratory for Applied Ontology, ISTC-CNR (2003).
- [5] S. Borgo, C. Masolo, Foundational Choices in DOLCE, in: S. Staab, R. Studer (Eds.), *Handbook on Ontologies*, 2nd ed., Springer Verlag, 2009, pp. 361–381.
- [6] S. Borgo, R. Ferrario, A. Gangemi, N. Guarino, C. Masolo, D. Porello, E. M. Sanfilippo, L. Vieu, Dolce: A descriptive ontology for linguistic and cognitive engineering, *Applied Ontology* 17 (2022) 45–69.
- [7] S. Borgo, L. Vieu, Artefacts in formal ontology, in: A. Meijers (Ed.), *Philosophy of Technology and Engineering Sciences*, North-Holland, Amsterdam, 2009, pp. 273–307.
- [8] S. Borgo, M. Franssen, P. Garbacz, Y. Kitamura, R. Mizoguchi, P. E. Vermaas, Technical artifacts: An integrated perspective, *Applied Ontology* 9 (2014) 217–235.
- [9] F. Compagno, S. Borgo, Functional modelling in engineering: A first comparison of approaches on two problems, in: E. M. S. et al. (Ed.), *12th International Workshop on Formal Ontologies Meet Industry (FOMI)*, volume 3240 of *CEUR Workshop Proceedings*, 2022.
- [10] F. Compagno, S. Borgo, Towards a formal ontology of engineering functions, behaviours, and capabilities, *Semantic Web* 15 (2024) 285–318. URL: <https://journals.sagepub.com/doi/abs/10.3233/SW-223188>. doi:10.3233/SW-223188. arXiv: <https://journals.sagepub.com/doi/pdf/10.3233/SW-223188>.
- [11] N. Angius, G. Primiero, The logic of identity and copy for computational artefacts, *Journal of Logic and Computation* 28 (2018) 1293–1322. doi:10.1093/logcom/exy012.
- [12] L. Biccheri, F. Compagno, S. Borgo, Dependability in engineering: An ontological stance, 2025. Manuscript submitted for publication.
- [13] P. Korenhof, V. Blok, S. Kloppenburg, Steering representations—towards a critical understanding of digital twins, *Philosophy & Technology* 34 (2021) 1751–1773. doi:10.1007/s13347-021-00472-0.
- [14] M. Carrara, M. Soavi, Copies, replicas, and counterfeits of artworks and artefacts, *The Monist* 93 (2010) 414–432. doi:10.5840/monist201093325.
- [15] L. Biccheri, R. Ferrario, D. Porello, Needs and intentionality: An ontological analysis and an application to public services, in: B. Brodaric, F. Neuhaus (Eds.), *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2020.
- [16] M. Liu, S. Fang, H. Dong, C. Xu, Review of digital twin about concepts, technologies, and industrial applications, *Journal of Manufacturing Systems* 58 (2020) 346–361. doi:10.1016/j.jmsy.2020.06.017.
- [17] D. J. Wagg, C. Burr, J. Shepherd, Z. X. Conti, M. Enzer, S. Niederer, The philosophical foundations of digital twinning, *Data-Centric Engineering* 6 (2025) e12. doi:10.1017/dce.2025.12.
- [18] A. Tzouvaras, Significant parts and identity of artifacts, *Notre Dame Journal of Formal Logic* 34 (1993) 445–452.
- [19] L. Biccheri, Needs as Mental Attitudes: An Ontological Study for PA’s Service Design, Phd thesis, University of Urbino, 2021.
- [20] S. Yablo, *Aboutness*, volume 3, Princeton University Press, Princeton, NJ, 2014.
- [21] P. Jacob, Intentionality, in: E. N. Zalta, U. Nodelman (Eds.), *The Stanford Encyclopedia of*

- Philosophy, spring 2023 ed., Metaphysics Research Lab, Stanford University, 2023. <https://plato.stanford.edu/archives/spr2023/entries/intentionality/>.
- [22] Y. Hui, What is a digital object?, in: H. Halpin, A. Monnin (Eds.), *Philosophical Engineering: Toward a Philosophy of the Web*, John Wiley & Sons, 2013, pp. 52–67.
- [23] C. Masolo, E. M. Sanfilippo, Technical artefact theories: A comparative study and a new empirical approach, *Review of Philosophy and Psychology* 11 (2020) 831–858. URL: <https://doi.org/10.1007/s13164-020-00475-9>. doi:10.1007/s13164-020-00475-9.
- [24] E. M. Sanfilippo, Ontologies for information entities: State of the art and open challenges, *Applied Ontology* 16 (2021) 111–135.
- [25] W. Ceusters, B. Smith, Aboutness: Towards foundations for the information artifact ontology, in: *Proceedings of ICBO 2015*, volume 1515, 2015, pp. 1–5.
- [26] F. Toyoshima, A. Barton, O. Grenier, Foundations for an ontology of belief, desire and intention, in: *Formal Ontology in Information Systems*, 2020, pp. 140–154.
- [27] J. Williamson, Causal pluralism versus epistemic causality, *Philosophica* 77 (2006).
- [28] F. Russo, J. Williamson, Interpreting causality in the health sciences, *International Studies in the Philosophy of Science* 21 (2007) 157–170.
- [29] J. Williamson, *Bayesian Nets and Causality: Philosophical and Computational Foundations*, Oxford University Press, Oxford, 2004.
- [30] M. Wilde, J. Williamson, Evidence and epistemic causality, in: *Statistics and Causality: Methods for Applied Empirical Research*, Wiley, 2016, pp. 31–41.
- [31] P. Gärdenfors, *Conceptual Spaces: The Geometry of Thought*, MIT Press, Cambridge, MA, 2004.
- [32] Y. Zou, J. D. Clark, A. A. May, A systematic investigation on the effects of temperature and relative humidity on the performance of eight low-cost particle sensors and devices, *Journal of Aerosol Science* 152 (2021) 105715.
- [33] C. Masolo, L. Vieu, E. Bottazzi, C. Catenacci, R. Ferrario, A. Gangemi, N. Guarino, Social roles and their descriptions, in: D. Dubois, C. Welty, M.-A. Williams (Eds.), *Proceedings of the 9th International Conference on the Principles of Knowledge Representation and Reasoning (KR)*, AAAI Press, Menlo Park, CA, 2004, pp. 267–277.
- [34] J. R. Searle, Intrinsic intentionality, *Behavioral and Brain Sciences* 3 (1980) 450–457.