



The ontology of the digital twin: a contemporary case for metaontological analysis

Orhan Onder^{1,2} 

Received: 5 January 2026 / Accepted: 7 April 2026
© The Author(s) 2026

Abstract

For decades, metaontological debates have been structured around paradigmatic examples: *Sherlock Holmes* for fictional objects and the *number 7* for abstract objects. These examples, while useful, are philosophically “safe” and historically entrenched. However, 21st-century technology creates entities that defy these clean categorizations. In this paper, I argue that the digital twin—a dynamic virtual model of a physical system—serves as a useful contemporary test case for metaontological reflection. It brings together the problematic features of both fiction and abstract objects while introducing a new, critical dimension: causal efficacy. This novel constellation of attributes forces refinements in debates between Quineanism, fictionalism, Meinongianism, and grounding-based approaches. Analysis shows that digital twins strengthen Quinean indispensability arguments, challenge fictionalist and Meinongian views due to their active role in the causal order, and are best elucidated as dependent entities through grounding relations. By applying these frameworks to digital twins, I aim not to declare any single framework victorious but to show where different theories must extend, qualify, or recalibrate their core commitments when confronted with this kind of entity. The goal is to demonstrate how ontological inquiry can illuminate emerging technologies—including their normative and institutional dimensions—and how these technologies can, in turn, reinvigorate metaontological theory.

Keywords Digital twins · Metaontology · Abstract artifacts · Fictionalism · Meinongianism · Quineanism · Grounding approach

✉ Orhan Onder
orhan.onder@marmara.edu.tr

¹ Department of History of Medicine and Ethics, Marmara University, Başlibüyük, Başlibüyük Yolu No: 9 D:2, Maltepe, İstanbul 34854, Türkiye

² Department of Philosophy, University of Vienna, Vienna, Austria

1 Introduction: Why we need a new philosophical example

For decades, the landscape of analytic ontology has been mapped using a familiar set of landmarks. When the debate turns to the status of non-existent objects, the quintessential example is *Sherlock Holmes*. When the focus shifts to abstract objects, the discussion invariably orbits around entities like *the number 7*. This familiarity, however, has a potential cost. These examples have become so well-established that they risk insulating ontological theories from the challenges posed by emerging, hybrid entities. This paper argues that the technological phenomenon of the *digital twin* provides a useful contemporary case through which to revisit classic debates between Quineanism, fictionalism, Meinongianism, and grounding-based approaches. By bringing together features that are often discussed separately, digital twins help show where different theories must clarify, refine, or qualify their commitments. The point is not that metaontological frameworks stand or fall with this case, but that the case makes certain pressures newly visible.

The utility of these classic examples is undeniable. *Sherlock Holmes* embodies “The Problem of Non-Existence” (Parsons, 1980). Holmes is the central figure in discussions about how we can refer to and make true claims about things that lack being (Crane, 2012). It is a subject of analysis of definite descriptions, a touchstone for Meinongian theories that grant non-existent objects a form of being-so (*Sosein*), and the prime example for fictionalists like Yablo (2001) who argue our discourse about him is “good without being true” (Berto, 2008; Berto & Plebani, 2015; Genot, 2021; Yablo, 2001). Holmes represents the challenge of intentionality—how our minds can be directed toward the non-actual—and the status of “created” objects whose properties are determined by authorial fiat (Goodman, 2003).

Conversely, *the number 7* is the archetype for “The Problem of Abstraction.” (Yablo, 2005) It is the central battleground in the dispute between Platonism, which holds that numbers are mind-independent abstract entities, and nominalism, which denies the existence of such abstracta (Berto & Plebani, 2015). The debate, tracing back to Frege (1884) and Russell (1919), hinges on the number’s seeming indispensability to our best scientific theories (Quine, 1948; Putnam, 2014) despite its causal inertness, lack of spatiotemporal location, and eternal, unchanging nature. The number 7 forces questions about epistemic access: how can we know anything about an entity with which we cannot interact causally?

Despite their foundational role, Holmes and the number 7 often get discussed within relatively stable problem-settings. This is not because they are philosophically simple—indeed, sophisticated accounts of ficta and abstracta already complicate familiar dichotomies—but because debates about them have generated comparatively well-developed literatures with recognizable fault lines. Digital twins are useful in this context not because they overturn those earlier cases, but because they combine several pressures at once: representation, particularity, dynamic updating, and causal integration into cyber-physical systems.

The digital twin presents a novel constellation of familiar problems. A digital twin is a dynamic, virtual representation of a specific physical system (e.g., Jet Engine #48, Patient X’s Heart) that is connected to it via real-time data flows (Jones et al.,

2020)¹. It is *abstract in its constitution*: like the number 7, it is comprised of data, algorithms, and models—intangible, non-spatiotemporal code. Yet, it is *concrete in its function and particularity*: like created representational entities more generally, it is intentionally constructed rather than naturally occurring, but unlike standard fictional cases it maintains a live, causal-informational link to a specific physical counterpart. Its state is constantly updated by sensory data from the physical world, and its simulations can inform decisions that act back upon that world (e.g., triggering a shutdown).

This fusion of attributes does not “shatter” our categories, but it does reveal their limits when confronted with a hybrid entity. Is it a useful fiction? A new kind of abstract object? A dependent particular? It is both *created* (like Holmes) and *objectively predictive* (like a mathematical model), possessing a utility that feels more substantive than literary pretense. This very ambiguity is what makes it such a productive philosophical tool.

I will argue that the digital twin is a supremely useful contemporary case for modern ontology. First, I will describe digital twins, highlighting their ontologically puzzling features, then use them to interrogate four key metaontological frameworks:

1. The Quinean Criterion: Does our best engineering science’s indispensable use of digital twins force us to include them in our ontology?
2. Fictionalism: Can we treat discourse about digital twins as a mere useful pretense, given their causal efficacy?
3. Meinongianism: Does the distinction between *Sein* (existence) and *Sosein* (being-so) adequately capture their status?
4. The Grounding Approach: Can we understand digital twins as fully grounded in their physical and formal constituents?

By applying these frameworks to digital twins, I will show that they offer a useful additional point of comparison for ontological theories, especially in light of technologically mediated entities that combine features often discussed separately in the literature.

A brief clarification about the aims of this inquiry is in order before proceeding. The claim defended here is not that a metaontological framework stands or falls based on its ability to accommodate digital twins. A Meinongian approach, for instance, might provide the best overall treatment of fictional characters even if it struggles with the causal efficacy of digital twins; a Quinean naturalism might be compelling in the philosophy of science even if it forces an uncomfortable expan-

¹ The philosophical and ethical implications of this linkage are particularly pronounced in medicine, where the digital twin of a patient’s organ or body raises fundamental questions about representation, identity, and extension (Braun, 2021, 2022; Nyholm, 2021). As Braun (2021) frames it, the core question is whether such a simulation can, and should, be treated as a ‘surrogate representing the body of the represented person.’ This has sparked a focused debate on whether a medical digital twin could be considered an extension of the patient’s body itself and what near-term ethical challenges, such as epistemic reliability and medical paternalism, this would entail (Mittelstadt, 2021; Nyholm, 2021). Iglesias et al. (2025) extend this discussion to the context of lifespan extension, asking whether digital doppelgängers might achieve some of the aims of prolonging personhood even absent biological continuity. (I return to these normative dimensions in Sect. 4.)

sion of our ontology. The goal, rather, is diagnostic. By examining how different frameworks fare when pressed to account for this unusual entity, we gain insight into their respective strengths, limitations, and points of internal pressure. A framework's capacity to illuminate the digital twin counts in its favor; its failure to do so reveals not a global verdict against it, but a specific structural limitation that any adequate account of hybrid, causally embedded entities must eventually address. In this sense, the digital twin functions less as a judge and more as a diagnostic lens—a way of revealing where each framework's conceptual resources run out, and what a more adequate ontology would need to supply.

2 What is a digital twin? A primer for ontologists

To subject digital twins to ontological scrutiny, a precise understanding of their nature is required. For the ontologist, a digital twin is not merely a piece of software; it is something whose nature can be illuminated only by examining a complex nexus of metaphysical problems. Technically, it is a *dynamic, virtual representation of a specific physical system*—such as Jet Engine #48, Building B, or Patient X's heart—that is inextricably linked to its physical counterpart through a continuous, bidirectional flow of data. This definition, while accurate, belies a series of ontologically puzzling features that make it a far more complex entity than a simple simulation or model. Its architecture synthesizes attributes traditionally assigned to separate ontological categories, forcing a re-examination of what it means to be abstract, concrete, particular, and causally efficacious.

The ontological interest of digital twins emerges from the tension between four core features:

Abstract constitution At its core, a digital twin is an abstract entity. It is composed of data structures, algorithms, and mathematical models; precisely the kinds of things philosophers classify as abstract objects. Abstract objects are typically characterized as non-spatiotemporal and causally inert. A digital twin's code exists as a pattern, not a physical thing; it is the software *qua* type, not its specific instantiation on a server. This is evidenced by its *multiple realizability*: the same digital twin (e.g., the model of Jet Engine #48) can be executed on different hardware platforms in different locations without being a different twin, much like the number 7 is referred to by different numerals but remains itself. In this sense, it shares a fundamental property with numbers and propositions.

This characterization, however, invites a potential confusion that must be addressed head-on: the distinction between type and token. To say a digital twin is multiply realizable might suggest it is a type—a universal or kind that can have multiple instances. This interpretation would be a mistake. The software program, the simulation engine, or the algorithm that defines the twin's behavior is indeed a type. It is the abstract structure that can be instantiated on countless servers. However, the digital twin of Jet Engine #48 is not this type; it is a token of that type. It is a particular, historical entity whose identity is fixed by its relationship to a specific physical counterpart.

When the simulation code for Engine #48 is executed on Server A in London and then migrated to Server B in Istanbul, it is the same digital twin—the same token—running on different hardware. What changes is its physical substrate, not its identity. This token, despite its particularity, retains an abstract constitution: its states are patterns of information, not physical stuff. Thus, the digital twin is best understood as a token-abstracta or a dynamic abstract particular: a non-physical entity whose spatio-temporal location is derivative, grounded in the physical systems that realize it, and which is nonetheless a single individual with a unique history and identity conditions.

Concrete functionality and particularity Despite its abstract constitution, a digital twin is not a universal. It is resolutely *particular*. It is not a simulation of *a* jet engine in general, but a representation of *Jet Engine #48*, a specific spatiotemporal particular with a unique history and location. This particularity is secured through its constant data linkage. Sensors on the physical engine feed real-time data (temperature, pressure, RPM) into the twin, anchoring its state to the precise, current state of its physical counterpart. This creates a direct informational correlation that transcends mere representation, establishing a unique relationship of mutual specification. This hybrid nature, as an abstract structure tied to a concrete particular, is its first major ontological novelty.

Dynamic temporal evolution This feature starkly contrasts digital twins with traditional philosophical examples of abstracta. The number 7 is static and eternal; its properties are timeless. Sherlock Holmes's properties, while extensible through later authors and fan fiction, are not updated by any mind-independent physical process. A digital twin, however, has a *lifecycle*. It evolves in real-time, mirroring the lifespan of its physical twin. Its properties are dynamic and state-dependent; the value of its variables changes from one moment to the next as new data arrives. This introduces a temporal dimension to its abstract core, challenging the notion that abstracta must be eternal and unchanging. It is an abstract particular with a history.

This temporal dimension might further suggest that digital twins are best understood as abstract continuants—entities that persist through time by having different temporal parts or stages at different times (see, e.g., Kaplan, 1990; Sainsbury and Tye, 2012 on related notions). On a stages view, the twin at time $t1$ (when the engine is cold) and the twin at time $t2$ (when the engine is hot) would be distinct but related temporal counterparts. While this captures the twin's dynamic evolution, it misidentifies the basis of its unity. The digital twin's identity across time is not secured by the immanent causality of a physical substance, but by its continuous, diachronic link to its physical counterpart. It is the same twin because it is the informational mirror of the same engine, and that engine provides a unifying thread of identity. Furthermore, the same twin can be migrated across hardware platforms, meaning it is not constituted by its physical stages. Its identity is tied to its historical and functional relationship with a specific physical particular, not to the series of computational events that realize it. This makes it a continuant of a unique sort—an informational continuant—whose persistence conditions are relational rather than substantial.

This profile resonates with recent work on the ontology of artifacts. Irmak (2012), for instance, argues that software is best understood as an abstract artifact: non-spatial yet temporal, brought into existence by intentional acts, and capable of surviving changes in its underlying algorithm (Irmak, 2012). Irmak (2018) extends this framework to words, showing that the category of abstract artifacts encompasses entities whose identity is secured by historical continuity (Irmak, 2018). Irmak (2023) further develops this approach, arguing that social and cultural types are dependent objects whose existence hinges on human activities rather than eternal properties—a framework that illuminates the ontologically hybrid character of digital twins. Digital twins share some of these features—they are created, abstract, and historically individuated—but they also introduce a novel element: a dynamic, causal-informational tie to a specific physical counterpart that itself lacks the creative intentions typical of artifact makers. A further difference concerns spatiality: while Irmak’s abstract artifacts are simply non-spatial, a digital twin’s spatiotemporal location is derivative, grounded in the physical systems that realize it. In this respect, they both invite comparison with and resist full assimilation to the abstract artifact paradigm.

Epistemic and causal efficacy This is the most potent feature from an ontological perspective. Digital twins are not passive descriptors; they are active, epistemically powerful tools. Their simulations generate predictions and insights. Crucially, these outputs can *directly cause actions in the physical world*. For instance, if a digital twin predicts an imminent fault in its physical engine, this prediction can automatically trigger a shutdown procedure. This blurs, or perhaps erases, the traditional line between abstract representation and concrete cause. Events are often understood as concrete particulars that are changes in objects and are located in the causal network (Bennett, 2002; Casati & Varzi, 2015). The digital twin’s prediction is an event *within the abstract model* that precipitates a physical event (the shutdown). This grants the abstract entity a form of *derived causal efficacy* that is deeply problematic for standard ontologies. It is no longer merely a tool for human thought but an autonomous node in a cyber-physical causal network.

This claim of causal efficacy, however, immediately invites a serious philosophical puzzle, one familiar from debates in the philosophy of mind. If the digital twin’s prediction is a cause of the physical shutdown, how does this cause relate to the physical processes occurring in the computer hardware? On the one hand, if the physical state of the server (the electrons, the circuits) is also a cause of the shutdown, the event appears to have two distinct causes—one abstract, one physical—raising the specter of causal overdetermination. On the other hand, if the twin’s prediction is the sole cause, we seem to violate the principle of the causal closure of the physical world, which holds that every physical event has a sufficient physical cause. This dilemma—overdetermination or closure violation—will prove to be a crucial test for the metaontological frameworks examined in Sect. 3. As we shall see, some frameworks (like fictionalism and Meinongianism) struggle to accommodate any causal role for the twin at all, while others (like the grounding approach) offer resources to show how the twin’s causal power can be derived from or realized by the physical processes that implement it, thereby avoiding the dilemma. For now, it is enough

to note that the digital twin's causal entanglement with the physical world is not a simple feature, but a source of profound ontological pressure.

At this point, a comparison with recent work on the metaphysics of virtual reality is instructive. David Chalmers (2017, 2022) has argued that virtual objects—such as the sword in a video game or the avatar in a virtual world—are real digital objects. They are constituted by data structures and have genuine properties within their virtual environments. Crucially, for Chalmers, these objects can be experienced perceptually and can enter into causal relations (a virtual sword can cause a virtual dragon to lose health points). While this framework resonates with the present account, a key difference must be noted. For Chalmers, virtual objects are typically environment-bound; their causal powers, while they can affect us psychologically through experience (e.g., via a VR-headset), do not extend to direct causal intervention in mind-independent physical systems in the way digital twins do. A digital twin, by contrast, is defined by its trans-boundary causal role. Its predictions are not merely events within a simulation; they are explicitly designed to reach out of the virtual environment and directly cause changes in the physical world. The twin's causal efficacy straddles the digital-physical divide in a way that Chalmers' paradigmatic virtual objects do not. This reinforces the sense in which digital twins constitute a distinct and particularly challenging ontological category, even within a robust realism about the digital domain.

In summary, a digital twin is an entity that is: abstractly constituted yet concretely functional; a particular yet non-physical (though spatiotemporally realized); dynamically temporal yet not a substance; and causally efficacious yet not tangible. This combination of attributes makes it a uniquely challenging and fertile object for ontological inquiry, as it refuses to be neatly placed into any of the categories carved out by philosophers centered on examples like Sherlock Holmes or the number 7.

3 How digital twins challenge established frameworks

3.1 The Quinean challenge: What does our best science quantify over?

The Quinean metaontological framework provides a seemingly straightforward methodology for ontology: we look to our best scientific theories, regiment them into first-order logic, and see what objects their bound variables range over. “To be is to be the value of a variable” (Berto & Plebani, 2015, p. 77; Boolos, 1984). This naturalistic approach defers ontological questions to the practices of the sciences. The digital twin provides a useful case for this criterion, since it raises the question whether our best contemporary engineering theories commit us to entities of this kind.

For Quine, ontology is not an autonomous, *a priori* philosophical discipline. It is a direct consequence of scientific theorizing. Quine's criterion of ontological commitment is the thesis that “a theory is committed to those and only those entities to which the bound variables of the theory must be capable of referring in order that the affirmations made in the theory be true” (Quine, 1948). Ontological inquiry thus proceeds by presupposing our best scientific theories, paraphrasing their claims into the canonical notation of first-order logic, and then reading off their ontological com-

mitments from those entities over which the regimented theory must existentially quantify in order to be true (Berto & Plebani, 2015, p. 32; Quine, 1964). This methodology naturalizes ontology, making it continuous with science. The existence of electrons is settled by physics, not philosophy. The key question becomes: what does our *current* best science say?

The rise of cyber-physical systems and Industry 4.0 has given birth to a new engineering paradigm where digital twins are not optional tools but *indispensable components* of the theory itself. Our best theories for managing complex systems—modern aerospace engineering, smart power grids, advanced medical prosthetics—are no longer expressed solely in terms of physical mechanics and material science. They are fundamentally *informational* theories that *routinely and indispensably quantify over digital twins*.

Consider the formalized claim within an automated maintenance system:

‘ $\exists x$ (x is a digital twin of Turbine #101 & x predicts(stress > yield_point))’.

This is not a mere *façon de parler*; it is a proposition whose truth value is essential for the system’s operation. The theory’s predictive power, its ability to explain system states, and its capacity to prescribe actions are all contingent upon the existence of this entity ‘x’. To deny the existence of ‘x’ would be to render a vast body of successful engineering practice either false or meaningless. Following the Quinean mandate, a strict naturalism seems to *force us to admit digital twins into our ontology*. They are no less indispensable than the theoretical entities of physics, such as quarks or gravitational fields, which we accept on similar grounds of their utility in our best theories (Berto & Plebani, 2015; Putnam, 2014).

The classic Quinean escape route from unwanted ontological commitments is paraphrase. Could we paraphrase away all reference to digital twins, reducing talk of them to talk of their physical counterparts? A revolutionary ontological minimalist might attempt an eliminative paraphrase strategy akin to those (Linnebo, 2012). For instance, they might argue that “The digital twin of Turbine #101 predicts a fault” is just a misleading way of saying “Turbine #101 is arranged *twin-wise*” or “The data stream from Turbine #101, when processed by algorithm A, outputs ‘fault’.”

However, this paraphrase strategy faces a formidable challenge. The *autonomous, counterfactual behavior* of the digital twin makes reduction to purely physical predicates hopelessly complex and arguably incomplete. The digital twin can run “what-if” scenarios: simulations of conditions the physical turbine has *never* encountered. The claim “The digital twin shows that if the inlet temperature were 50°C higher, the rotor would fail” cannot be paraphrased as a claim solely about the actual, historical state of the physical turbine, for that state never obtained. The twin possesses *emergent properties*—predictive capacities and states—that are not reducible to a simple conjunction of physical states of its target system or the hardware it runs on. It is a distinct, albeit abstract, particular within the explanatory framework.

Digital twins powerfully strengthen the indispensability argument for a new class of entities that are abstract yet particular, and causally relevant yet non-physical. They force the Quinean naturalist to confront the fact that our “best scientific theories” now include those from advanced engineering and computer science, which are

deeply committed to the existence of such virtual objects. To reject digital twins on a priori philosophical grounds would be to violate the very naturalistic principle that animates the Quinean project. Consequently, digital twins act as a concrete example of how scientific practice itself can drive ontological revision, pushing Quineanism toward a more expansive and technologically engaged ontology that acknowledges the reality of the virtual objects indispensable to our understanding and control of the physical world.

3.2 A fictionalist interpretation: “Useful without being true”?

If the Quinean path leads to an ontological commitment to digital twins, the fictionalist offers an attractive escape route. Fictionalism provides a metaontological strategy for retaining the practical utility of a discourse while avoiding commitment to the controversial entities it appears to describe (Rosen, 1990). Applying this framework to digital twins reveals both its initial plausibility and its ultimate limitations, highlighting how digital twins pose a sterner challenge than traditional examples like Sherlock Holmes.

Fictionalism is the view that discourse apparently about a certain domain (e.g., mathematics, morality, possible worlds) should not be taken at face value. While we speak *as if* certain entities exist, these statements are not literally true; they are part of a “useful fiction” or a pretense. As Hartry Field famously argued regarding mathematics, the claim is that we can have all the benefits of mathematical discourse without the ontological costs—it can be “good without being true” (Field, 1980, 1989). When we utter such statements, we are not genuinely asserting their literal content but rather “quasi-asserting” within the rules of a make-believe game (Walton, 1990; Yablo, 2001). This allows us to bypass the Quinean criterion by denying that the indispensable statements of a theory require a literal interpretation.

A compelling case can be made for a hermeneutic fictionalist interpretation of digital twin discourse. Drawing on Kendall Walton’s theory of prop-oriented make-believe, one could argue that engineers engage in a sophisticated pretense (Walton, 1990). The physical data stream from Turbine #101 acts as a *prop*. The rule of the game, or principle of generation, is: “Let’s pretend that this incoming data stream *is* the turbine itself, and that this software model *is* a real, dynamic copy of it.” Within this pretense, statements like “The digital twin is overheating” are true. The real, valuable content of this quasi-assertion, however, is not about an abstract entity but about the *physical dispositions and states of the actual turbine*. The fiction is a powerful, indirect way to describe and predict the behaviour of the concrete world.

This interpretation passes a version of the “Oracle Test”. If an oracle were to reveal that digital twins do not *really* exist as entities, it is plausible that engineers would shrug and continue using them. Their primary concern is the accurate functioning of the physical system, not the ontological status of the software model. This pragmatic attitude—focusing on predictive success rather than metaphysical truth—aligns perfectly with a hermeneutic fictionalist stance, which holds that speakers are not, and need not be, committed to the literal truth of their indispensable discourse (Fraassen, 1980).

However, the fictionalist interpretation runs aground on the rock of *causal efficacy*. This is where digital twins fundamentally diverge from paradigmatic fictions like Sherlock Holmes and present a far more serious challenge than mathematical objects. On an interventionist account of causation, such as Woodward's, the digital twin's output qualifies as a genuine cause: intervening on the computational state of the twin would, *ceteris paribus*, change the downstream physical behavior of the turbine (Woodward, 2003).

The decisive objection is this: *Fictions do not cause changes in the physical world*. Sherlock Holmes does not arrest anyone; the number 7 does not physically calculate anything. Their utility is entirely mediated by human cognition. In contrast, a digital twin's output can *autonomously and directly trigger actions in the physical world*. In a fully automated system, the statement 'Twin(x).predicts(fault)' can be the literal, algorithmic cause of a shutdown command being sent to the physical turbine. This command is not issued by a human engineer interpreting the twin's output "as if" it were true; it is executed by the system itself.

This causal efficacy lends digital twins a reality that pure fictions lack. While their *constitution* is abstract, their *function* is concretely causal. This makes them a uniquely difficult case for fictionalism, straddling the line between useful pretense and real-world actor in a way that neither Sherlock Holmes nor the number 7 ever could. The fictionalist can perhaps maintain that we speak *as if* the twin causes the shutdown, but the physical reality of the event's causal history makes this stance feel profoundly inadequate, revealing the limits of fictionalism when confronted with technologies that blur the line between the abstract and the physically effective.

One might object, however, that this critique targets only naive versions of fictionalism, and that more sophisticated accounts of fictional objects fare considerably better. The most important such account is Artefactualism, defended most prominently by Thomasson (1999) and along with closely related abstract creationist views developed in subsequent work (e.g., Friedell, 2016, 2020). On this view, fictional characters are neither non-existent Meinongian items nor mere linguistic pretenses. They are created abstract artefacts: entities that come into existence through deliberate intentional acts, persist as long as their originating practices are sustained, and can acquire new properties over time. Crucially, on Friedell's account, Holmes is not frozen at the moment of Doyle's writing; he can acquire further properties through post-Doylean continuations and the ongoing practices that govern the development of the fiction. This makes artefactualist *ficta* dynamic, contingent, and historically situated—a profile that appears, at first glance, remarkably close to that of digital twins. Both are created; both are contingent abstract particulars; both evolve over time. If the parallels run this deep, one might wonder whether digital twins are simply a technologically updated instance of the artefactualist category, posing no fundamentally new ontological challenge.

This conclusion, however, is premature. Two differences resist the assimilation. The first concerns the source of dynamism. For artefactualist *ficta*, change is always mediated by human intentionality: Holmes acquires new properties because new authors write new stories, or because communities of readers reinterpret existing texts. The engine of change is psychological and cultural. A digital twin's state, by contrast, updates automatically in response to a mind-independent physical system.

When Jet Engine #48's temperature rises, the twin's variables change without any interpretive act, without any human creative input. The twin's dynamism is not the product of changing minds but of changing physical facts. The second difference is more decisive: causal efficacy. Even on the most liberal artefactualist accounts, fictional characters possess only indirect causal powers, mediated by human psychology. Holmes can cause readers to feel emotions or visit Baker Street, but these effects require minds as intermediaries. A digital twin's outputs, by contrast, can directly trigger physical events in fully automated systems, with no intentional agent in the loop. This is not causal efficacy *à la* remove; it is automated, embedded causal participation in a cyber-physical network.

These differences do not consign artefactualism to irrelevance. On the contrary, it provides valuable resources for understanding the creation and historical identity of digital twins. But when it comes to their operational profile—their mind-independent dynamism and direct causal integration—artefactualism requires significant extension. Digital twins thus exert pressure not only on Quinean, fictionalist, and Meinongian frameworks, but on the most sophisticated available accounts of created abstract objects as well.

3.3 A Meinongian interpretation: objects beyond existence

If Quineanism demands ontological commitment and Fictionalism seems inadequate due to causal efficacy, Meinongianism offers a third path. This framework, which radically decouples quantification from existence, appears uniquely positioned to capture the peculiar status of digital twins. However, upon closer inspection, the very feature that makes digital twins ontologically puzzling—their causal power—also exposes a critical weakness in the Meinongian approach, pushing it toward a concession it seeks to avoid.

Meinongianism, as articulated by Alexius Meinong and developed by modern proponents like Routley (1980) and Priest (2005), challenges the core Quinean equation of being with existence (Priest, 2016; Routley, 1980). It introduces a fundamental distinction between *Sein* (existence) and *Sosein* (being-so) (Deniz, 2018). According to the Principle of Independence, an object can have a definite set of properties (*Sosein*)—it can be thus-and-so—independently of whether it exists (*Sein*) (Meinong, 1988). This allows for a domain of objects that are “beyond being and non-being”; we can truly quantify over and refer to them, even though they lack existence. As Meinong famously put it: “There are objects of which it is true that there are no such objects” (Meinong, 1904/1960). This framework neatly accommodates Sherlock Holmes: he has the property of being a detective (his *Sosein*) but does not exist (lacks *Sein*).

At first glance, Meinongianism is a perfect fit for digital twins. It provides a coherent framework for describing our intuitive dealings with them without forcing a full-blooded existential commitment.

1. True Predication and Quantification: We can make true claims about digital twins (“The digital twin of Turbine #101 is more computationally complex than the

original CAD² model”) and quantify over them (“Some digital twins in the fleet are more accurate than others”). For the Meinongian, this is unproblematic: these statements are about the *Sosein* of the twins, about the properties they genuinely possess. We are quantifying over objects that have being, without thereby claiming they exist. This captures the intuition that we are talking about *something* real (the model, its properties) without having to claim that this something is a physical object or a Platonic form.

2. The Ambiguous Status: Digital twins seem to be prime examples of objects that defy traditional existential categories. They are not concrete physical particulars, nor are they eternal abstracta. They occupy a shadowy realm in between, seemingly possessing a rich *Sosein* (complexity, accuracy, predictive power) without a clear or traditional *Sein*. The Meinongian framework is explicitly designed to accommodate such ambiguous entities, offering a logical space for objects that are “there” to be referred to and have properties, yet are not existent in the ordinary sense.

However, the Meinongian solution unravels when confronted with the same feature that troubled the fictionalist: *causal interaction*. Meinong’s “jungle” of non-existent objects is populated by famously inert entities: the golden mountain, the round square, and Sherlock Holmes. These objects are causally isolated; they do not process information or effect changes in the world. This inertness is a direct consequence of their non-existence.

The digital twin, however, does not share this inertness. It *actively processes real-world data* from its physical counterpart and *outputs commands that physically change the world*. This active causal role creates a severe problem for the Meinongian:

How can a non-existent object be a causal agent? Causation is typically considered a relation that holds between existent states of affairs. If a digital twin’s prediction is the proximate cause of a turbine shutdown, it seems we must grant it some robust form of reality. To claim that a non-existent object can be a relatum in the causal network is a radical and counterintuitive proposition that stretches the Meinongian framework to its breaking point.

This problem forces a dilemma upon the Meinongian:

1. *Deny the Causation*: Argue that the causation is not *really* performed by the twin itself, but by the physical computer hardware executing the code. This, however, is a reductionist move that abandons the core Meinongian insight and aligns with the eliminativist paraphrase strategies they typically reject. It also fails to capture the explanatory practice of engineers who attribute efficacy to the twin as a functional unit.
2. *Grant a Form of Existence*: Concede that digital twins must have a “watered-down” form of being or a special mode of existence (Berto & Plebani, 2015, p. 101). This is the “modes-of-being Meinongianism” mentioned by Eklund (2006),

² Computer-Aided Design.

where being is not univocal (Eklund, 2006). Perhaps digital twins exist *virtually* or *digitally*—a new ontological category distinct from physical and abstract existence.

While the second option is more promising, it constitutes a significant retreat. It effectively admits that the pure Meinongian distinction between *Sein* and *Sosein* is insufficient for digital twins. These entities demand a richer ontological status than Sherlock Holmes. They require us to define a new kind of *Sein*—e.g. *virtual existence*—that accommodates their causal and dynamic nature. Thus, rather than simply being accommodated by Meinongianism, digital twins *pressure the theory to evolve*, forcing it to clarify and expand its conception of what it means to exist beyond the merely physical. They demonstrate that non-existence is not a single, monolithic category but may itself have a structure that includes causally potent virtual entities.

3.4 The grounding approach: mapping the dependency relations

The preceding analyses reveal the shortcomings of applying traditional metaontological frameworks directly to digital twins. Quineanism seems over-committed, Fictionalism is evasive, and Meinongianism struggles with causality. The grounding approach, however, offers a more nuanced and productive path forward. It shifts the question from the blunt “Do they exist?” to the more revealing “In *what manner* do they exist, and what explains that existence?” This framework provides the conceptual tools to map the unique metaphysical structure of digital twins without forcing them into ill-fitting categories.

Given that a digital twin’s spatiotemporal location is derivative (see Sect. 2), grounding offers a way to articulate this dependence. Grounding is a relation of non-causal, metaphysical determination. It is the relation that holds between a more fundamental reality and a less fundamental, dependent reality that it explains or ontologically supports (Rosen, 2010; Schaffer, 2009). As Berto and Plebani note, this approach reframes ontology. The key question is no longer the Quinean “What is there?” but rather “What grounds what?” or “What is fundamental?” (Berto & Plebani, 2015, p. 113). For instance, a statue is grounded in the clay that constitutes it; the moral wrongness of an act might be grounded in its natural properties. Grounding provides a hierarchical structure to reality, explaining the being of certain entities in terms of the being of others.

The grounding framework has certain advantages for capturing the ontology of digital twins. A digital twin is not a primitive or fundamental entity; its entire being is *grounded in* two distinct, more fundamental realms:

1. The Physical Particular: The digital twin’s *existence* and its *core identifying properties* are grounded in the existence of its specific physical counterpart. There is no digital twin of Turbine #101 without Turbine #101 itself. The twin’s identity, its history, and its real-time state are metaphysically determined by the state of this physical entity. The physical turbine is the ground for the digital twin.
2. The Abstract Model/Data: The digital twin’s *structure*, *form*, and *predictive behaviour* are grounded in the realm of abstracta. Its essence is defined by

the mathematical models, algorithms, and data structures that constitute it. Its dynamic evolution is grounded in the incoming stream of data from its physical source. The abstract model provides the formal cause, while the data stream provides the specific content that realizes that form at any given moment.

The digital twin, therefore, exists at the intersection of these two grounding chains. It is a *hybrid, dependent entity* whose entire nature is exhaustively explained by the combination of a specific physical particular and a specific abstract model fed by a data stream.

One advantage of the grounding approach is that it avoids the pitfalls of the other frameworks. In contrast to blunt Quineanism, it does not require treating digital twins as fundamental entities: their existence can be acknowledged without placing them on the same ontological level as the physical particles or abstract mathematical objects that ground them, since their existence is derivative and explained. At the same time, the grounding approach avoids the shortcomings of Fictionalism and Meinongianism by fully acknowledging that digital twins really exist as dependent entities. There is no need to regard discourse about them as a mere useful pretense or to claim that they lack existence altogether. Instead, they possess a genuine, albeit grounded, mode of existence.

This framework also provides the resources to resolve the causation puzzle. Recall the dilemma: if a digital twin's prediction causes a physical shutdown, does this overdetermine the event or violate causal closure? The grounding approach dissolves this dilemma by distinguishing between different senses of "cause" and clarifying the relation between the twin's abstract state and its physical realization. The twin's state predicts(fault) is not a fundamental or autonomous feature of reality; it is grounded in the physical hardware's computational states. The pattern of electrons in the server's memory, the execution of machine code, the operation of the CPU—these physical processes realize or implement the abstract state of the twin. The relation between the abstract prediction and the physical computation is one of metaphysical determination: the physical facts determine the abstract facts.

With this grounding relation in place, we can provide a precise account of the twin's causal role. The proximate or efficient cause of the shutdown is a physical event: the command signal transmitted from the server to the turbine's control system. This signal has a complete physical explanation in terms of the prior physical states of the computing hardware. The causal closure of the physical domain is thus respected; every link in the physical causal chain is filled by physical events. This means that attributing causal relevance to the digital twin does not introduce a second, rival cause alongside the physical realizer. The twin's predictive state and the realizing computational process are not two independent sources of efficacy, but the same causally efficacious process described at different explanatory levels: physically, as a sequence of hardware events; informationally, as the twin's prediction that guides system behavior.

What, then, is the role of the abstract prediction? It is best understood as a formal or informational cause. The prediction specifies what the shutdown is for, why it is occurring, and what information triggered it. It provides the content or the pattern that the physical process instantiates. But this formal cause does not compete with

the efficient cause; rather, it characterizes it at a different level of description. The shutdown event is both a physical process (the efficient cause) and the realization of an informational state (the formal cause). There is no overdetermination here because there are not two independent causal powers at work. The twin's causal efficacy is entirely realized by the physical processes that ground it. The abstract state does not reach down and add an extra push to the physical world; rather, it is the way we describe the informational organization of that physical push. As Kim (1998, pp. 72–87) himself suggested in his later work, the problem of causal exclusion can be dissolved if the higher-level property is not conceived as a distinct causal power but as a way of characterizing the lower-level realizer (Kim, 1998).

This analysis turns the apparent problem into a feature of the twin's dependent existence. The digital twin's causal efficacy is precisely a function of its being grounded in physical reality. It is because the twin is not a free-floating abstractum, but is tethered to the physical world through computational implementation, that it can participate in causal processes. Far from threatening naturalism, the twin's causal role exemplifies how abstract information structures can be seamlessly integrated into the physical causal order.

It is worth noting, however, that the above resolution—compelling as it may be on its own terms—presupposes a traditional conception of “physical” in the statement of the causal closure principle, one built around paradigm cases like particles, organisms, and tangible objects. A more radical reading is available. If digital twins are genuinely hybrid entities with a *sui generis* spatiotemporal profile, as argued above, then the dilemma itself (overdetermination or closure violation) may reflect a conceptual framework that was not designed with such entities in mind. On this alternative view, the proper response to the causation puzzle is not to show that digital twins fit within existing causal categories, but to acknowledge that understanding them may require revisiting what “physical” means in the closure principle itself. Whether this stronger, revisionary conclusion is warranted goes beyond the diagnostic aims of this paper. But it is a direction that the grounding framework, more than any other examined here, makes visible—and it is one that future work on the metaphysics of hybrid cyber-physical entities may need to pursue.

For present purposes, then, the grounding framework offers a way of respecting the twin's reality while explaining its derivative status. It provides a metaphysical map: the digital twin exists at the intersection of two grounding chains—one descending to the physical turbine that anchors its identity, another descending to the abstract models and data structures that constitute its form. This map explains the twin's ambiguous character: it feels abstract because its form is grounded in abstracta, and it feels concretely efficacious because its states are realized by physical processes. The grounding approach does not eliminate the puzzle of the digital twin, but it provides the sophisticated conceptual machinery needed to describe its unique ontological structure accurately.

4 Why digital twins matter for metaontological discussion

In this section, I present three reasons why digital twins matter for metaontological discussion alongside traditional examples such as Sherlock Holmes or the number 7. These reasons focus on the hybrid nature of digital twins, their causal and epistemic entanglement, and their deep connection to normative and institutional realities.

First of all, digital twins place familiar ontological distinctions under unusual pressure and thereby motivate a closer examination of our ontological categories. They embody a *hybrid* profile, merging attributes that would typically be classified as abstract or concrete, representation or participation, and static or dynamic. On one hand, digital twins are abstract entities: they consist of code, algorithms, and models, which are intangible and typically understood as non-physical, much like other abstract objects. On the other hand, they are deeply connected to concrete systems, as they mirror and interact with real-world counterparts in real time. This relationship challenges the classical distinction between representation and participation. While a traditional representation simply refers to an object without influencing its physical counterpart, a digital twin participates actively in the system it models, influencing it directly.

Additionally, the temporally dynamic coupling between the digital twin and its physical counterpart further complicates the traditional metaphysical division between static abstractions and concrete particulars. Unlike eternal (or platonic, or traditionally-conceived) abstract objects, which remain unchanged over time, digital twins evolve and adapt in real-time, continuously updating in response to data from their physical counterparts. This dynamic relationship complicates the familiar contrast between fictional and abstract objects by concentrating, within one case, features that are often treated separately in more traditional examples.

Secondly, unlike traditional representational entities, digital twins are embedded both *causally* and *epistemically* in the systems they model. Their outputs are not just passive depictions of reality; they actively alter the physical systems they represent. For example, when a digital twin predicts a failure in a machine and triggers an automated response, this action directly affects the physical world, often without human intervention. This causal involvement of digital twins makes them ideal for testing how ontology handles derived or implemented causal powers—powers that are not intrinsic to the entity itself but emerge from its relationship with other systems.

This characteristic also pressures ontologies that treat representational entities as lacking the kind of direct, automated integration into physical control loops that digital twins can possess. In contrast, digital twins blur the lines between representation and action, illustrating a new category of objects that possess causal efficacy. As such, they present a compelling challenge to both fictionalism, which denies the real existence of objects like digital twins, and the traditional debates over abstract objects, which assume a lack of causal involvement in the world.

Thirdly, digital twins are not only technically sophisticated; they also exist within a broader network of *normative* and *institutional realities*. Their identity, persistence, and utility are governed by conventions, regulations, and dependencies, such as certification processes, responsibility assignments, and data provenance protocols. These aspects of their existence are shaped by social, ethical, legal, and technical

frameworks, not just by intrinsic metaphysical essences. In the medical context for instance, this embedding is starkly evident. A digital twin's reliability as a health technology is not a given but is contingent on frameworks of epistemic responsibility for clinicians and developers, raising questions about the standards of evidence required to justify acting on a twin's predictions (Mittelstadt, 2021; Onder & Ilkilic, 2026). Its use also intersects with existing frameworks of consent and privacy, potentially introducing new forms of medical paternalism if the twin's 'objective' output is privileged over the patient's own testimony (Braun, 2021, 2022; Mittelstadt, 2021). This institutional embedding reveals that digital twins cannot be understood merely as isolated metaphysical entities but must be considered in relation to the frameworks that sustain their existence.

This institutional and normative embedding matters not because it dictates a particular ontological account, but because it further complicates the conditions under which digital twins persist, function, and are treated as entities within explanatory practices. Unlike traditional philosophical examples, digital twins are sustained not only by technical implementation but also by regulatory, legal, and organizational frameworks that shape their identity and use over time. This adds an additional dimension to their role as a test case, insofar as it highlights how ontological categories interact with broader structures of practice and recognition.

Taken together, these three considerations suggest that digital twins occupy a distinctive position within contemporary ontological reflection. Their hybrid constitution unsettles familiar abstract–concrete distinctions; their causal and epistemic entanglement challenges standard assumptions about representation and efficacy; and their institutional embedding connects ontological questions to normative and social structures in an unusually direct way. For these reasons, digital twins appear well suited to function as a useful additional case for metaontological reflection alongside the more stable and isolated examples that have traditionally structured the debate.

At the same time, however, precisely because this claim concerns the role digital twins play in metaontological inquiry, rather than their ultimate ontological classification, it is reasonable to ask how robust the test-case proposal itself is. If digital twins are to serve as a productive point of comparison for competing metaontological approaches, it must be shown that their philosophical significance does not rest on exaggeration, technological novelty alone, or an implicit commitment to any particular framework. The following remarks therefore do not aim to defend a specific ontological view, but rather to clarify the sense in which digital twins genuinely raise metaontological pressure, even when viewed from different theoretical perspectives.

One possible reservation is that digital twins may not be fundamentally new, but instead represent technologically updated instances of familiar ontological categories—such as models, artefacts, or simulations. On this view, they would not introduce novel philosophical pressure, but merely reinstantiate existing problems in a more complex technical setting. While this reaction is understandable, it arguably underestimates the way in which digital twins combine features that are typically examined in isolation. Traditional test cases tend to separate abstraction from causal relevance, representation from participation, and temporal stability from dynamic evolution. Digital twins are philosophically salient not because each of these features is unprecedented, but because they are systematically intertwined in a single kind of

entity in ways that traditional examples rarely are. It is this conjunction, rather than any individual property, that makes them difficult to accommodate within familiar ontological distinctions.

A second reservation is that the challenges posed by digital twins are primarily pragmatic or epistemic, rather than ontological. One might argue that they complicate explanation, prediction, or control, but do not substantially affect questions about what there is. Yet this response presupposes a sharp boundary between epistemic practice and ontological commitment, a boundary that metaontology itself has repeatedly called into question. Digital twins are not merely useful representational tools; they are often indispensable within theories that successfully guide intervention in physical systems. Once this indispensability is acknowledged, familiar metaontological questions naturally re-emerge: whether such entities ought to be taken at face value, paraphrased away, or assigned a derivative status. In this respect, their epistemic role is precisely what renders them ontologically interesting.

A third concern is that digital twins may fail as a genuine test case because they can, in principle, be accommodated by a wide range of metaontological frameworks. If Quinean, fictionalist, Meinongian, and dependency-oriented approaches can all offer coherent accounts, one might doubt that digital twins meaningfully discriminate between them. However, the aim of a test case need not be to refute theories outright. What digital twins appear to do is reveal where different frameworks must qualify, extend, or recalibrate their core commitments—whether concerning abstraction, existence, causation, or dependence. The philosophical value of the example lies in the comparative pressure it generates, rather than in its capacity to settle disputes.

A final point concerns the historical contingency of digital twins as artefacts of contemporary engineering practice. But historical contingency does not by itself diminish ontological interest: metaontology has often relied on concrete and historically situated examples, including fictional characters and scientific models, to probe broader distinctions. The interest of digital twins does not depend on the permanence of any particular technology, but on the structural possibility they instantiate: entities that are abstractly constituted, dynamically updated, and causally integrated into physical systems. It is this structural profile, rather than its technological origin, that carries philosophical significance.

Taken together, these considerations suggest a revisionary, if deliberately open-ended, conclusion. The investigation of digital twins does require us to revise our ontology in one precise sense: it shows that existing metaontological frameworks each encounter serious difficulties when confronted with entities that are abstractly constituted, dynamically updated, and causally integrated into physical systems. This is not a merely negative result. Identifying the structural inadequacy of available frameworks is itself an ontological contribution, one that clarifies what a satisfactory account of digital twins would need to achieve. Digital twins do not privilege any single framework; but they do set new, demanding criteria that any adequate framework must meet. In this sense, they serve not merely as a “point of reflection” that reopens old debates, but as a productive constraint on where the next phase of metaontological work must go.

5 Conclusion: digital twins as a contemporary ontological case

The philosophical investigation of digital twins shows that, in addition to being objects of technological interest, they can also serve as useful cases for ontological analysis. Considered in this way, they help bring into focus how different metaontological frameworks handle entities that are abstractly constituted, dynamically updated, and embedded in cyber-physical systems. Digital twins bring together the classic challenges posed by fictional objects like Sherlock Holmes and abstract objects like the number 7, while introducing a new, critical dimension: *direct causal efficacy* within the mind-independent, physical world. This combination of attributes makes them a valuable contemporary case, especially because it concentrates in one setting several pressures that are often discussed separately in ontological debate.

This article highlights three interconnected ways in which digital twins function as a useful test case for metaontology. First, their role in contemporary engineering practice places pressure on Quinean approaches by extending indispensability considerations to abstract yet causally embedded entities. Second, their autonomous causal role exposes limitations in both fictionalist and Meinongian strategies that rely on inert or non-existent objects. Third, dependency-based frameworks such as grounding approaches illustrate one way of mapping the hybrid structure of digital twins, showing how their causal efficacy can be understood as derived from the physical processes that realize them.

These observations are not intended to settle metaontological disputes, nor to declare any single framework victorious. A theory's ability to handle digital twins is one data point among many in assessing its overall plausibility. The value of the exercise lies elsewhere: in demonstrating how emerging technological practices can reinvigorate philosophical inquiry. For frameworks concerned with ontology, digital twins provide a contemporary case in which questions of abstraction, dependence, and causal integration become especially vivid. In this sense, they complement more familiar examples such as Holmes and numbers by bringing contemporary technological practice into contact with foundational ontological questions. Their value lies not in overturning existing debates, but in helping to reframe them under conditions shaped by increasingly hybrid forms of representation and intervention.

Acknowledgements The author would like to thank M. Emin Pilgir and the two anonymous reviewers for their careful reading and constructive feedback, which led to substantial improvements in the manuscript. The Ontology/Metaontology reading group at Abbara is gratefully acknowledged for providing the intellectual context in which the central idea of this paper took shape.

Author contributions The author is solely responsible for the conception of the study, the analysis and interpretation of the arguments, and the writing of the manuscript.

Funding Open access funding provided by the Scientific and Technological Research Council of Türkiye (TÜBİTAK).

Data availability Not applicable.

Declarations

Competing interests The authors have no competing interests to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bennett, J. (2002). What events are. *The blackwell guide to metaphysics* (pp. 43–65). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9780470998984.ch3>
- Berto, F. (2008). Modal Meinongianism for fictional objects. *Metaphysica*, 9(2), 205–218. <https://doi.org/10.1007/s12133-008-0033-z>
- Berto, F., & Plebani, M. (2015). *Ontology and Metaontology: A contemporary guide* (M. Plebani, Ed.). Bloomsbury Academic. <https://philarchive.org/rec/BEROAM-2>
- Boolos, G. (1984). To be is to be a value of a variable (or to be some values of some variables). *The Journal of Philosophy*, 81(8), 430–449. <https://doi.org/10.2307/2026308>
- Braun, M. (2021). Represent me: Please! Towards an ethics of digital twins in medicine. *Journal of Medical Ethics*, 47(6), 394–400. <https://doi.org/10.1136/medethics-2020-106134>
- Braun, M. (2022). Ethics of digital twins: Four challenges. *Journal of Medical Ethics*, 48(9), 579–580. <https://doi.org/10.1136/medethics-2021-107675>
- Casati, R., & Varzi, A. (2015). Events. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Winter 2015). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2015/entries/events/>
- Chalmers, D. (2017). The virtual and the real. *Disputatio*, 9(46), 309–352. <https://doi.org/10.1515/disp-2017-0009>
- Chalmers, D. (2022). *Reality+: Virtual worlds and the problems of philosophy*. W. W. Norton.
- Crane, T. (2012). What is the problem of non-existence? *Philosophia*, 40(3), 417–434. <https://doi.org/10.1007/s11406-011-9354-1>
- Deniz, E. M. (2018). Yersiz Yurtsuz Nesneler Mantıkta Yer Tutar Mı? [VIII. Mantık Çalıştayı Kitabı, Aralık 2018, s. 121–128]. *VIII. Mantık Çalıştayı*. https://www.academia.edu/38147565/Yersiz_Yurtsuz_Nesneler_Mant%C4%B1kta_Yer_Tutar_M%C4%B1_VIII_Mant%C4%B1k_%C3%87al%C4%B1%C5%9Ftay%C4%B1_Kitab%C4%B1_Aral%C4%B1k_2018_s_121_128
- Eklund, M. (2006). Metaontology. *Philosophy Compass*, 1(3), 317–334. <https://doi.org/10.1111/j.1747-9991.2006.00026.x>
- Field, H. H. (1980). *Science without numbers: A defence of nominalism*. Princeton University Press.
- Field, H. H. (1989). *Realism, mathematics & modality*. Blackwell.
- Fraassen, B. C. van. (1980). *The scientific image*. Oxford University Press. <https://doi.org/10.1093/0198244274.001.0001>
- Frege, F. L. G. (1884). *Die Grundlagen der Arithmetik*. Koebner.
- Friedell, D. (2016). Abstract creationism and authorial intention. *The Journal of Aesthetics and Art Criticism*, 74(2), 129–137. <https://doi.org/10.1111/jaac.12271>
- Friedell, D. (2020). Abstracta are causal. *Philosophia*, 48(1), 133–142. <https://doi.org/10.1007/s11406-019-00073-9>
- Genot, E. J. (2021). The Holmesian logician: Sherlock Holmes' science of deduction and analysis and the logic of discovery. *Synthese*, 198(11), 10169–10186. <https://doi.org/10.1007/s11229-020-02709-w>

- Goodman, J. (2003). Where is Sherlock Holmes? *Southern Journal of Philosophy*, 41(2), 183–197. <https://doi.org/10.1111/j.2041-6962.2003.tb00948.x>
- Iglesias, S., Earp, B. D., Voinea, C., Mann, S. P., Zahiu, A., Jecker, N. S., & Savulescu, J. (2025). Digital dopelgängers and lifespan extension: What matters? *The American Journal of Bioethics*, 25(2), 95–110. <https://doi.org/10.1080/15265161.2024.2416133>
- Irmak, N. (2012). Software is an abstract artefact. *Grazer Philosophische Studien*, 86(1), 55–72. https://doi.org/10.1163/9789401209182_005
- Irmak, N. (2018). An ontology of words. *Erkenntnis* (1975–), 84(5), 1139–1158. <https://doi.org/10.1007/s10670-018-0001-0>
- Irmak, N. (2023). Type realism reconsidered. *The Journal of Aesthetics and Art Criticism*, 81(4), 497–507. <https://doi.org/10.1093/jaac/kpad043>
- Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020). Characterising the digital twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29, 36–52. <https://doi.org/10.1016/j.cirpj.2020.02.002>
- Kaplan, D. (1990). Words. *Proceedings of the Aristotelian Society Supplementary Volumes*, 64, 93–119. <https://www.jstor.org/stable/4106880>
- Kim, J. (1998). *Mind in a physical world: An essay on the mind-body problem and mental causation*. MIT Press.
- Linnebo, Ø. (2012). Metaontological minimalism. *Philosophy Compass*, 7(2), 139–151. <https://doi.org/10.1111/j.1747-9991.2011.00471.x>
- Meinong, A. (1904/1960). The theory of objects. In R. M. Chisholm (Ed.), *Realism and the Background of Phenomenology* (pp. 76–117). Free Press.
- Meinong, A. (1988). *Über Gegenstandstheorie. Selbstdarstellung* (A. Meinong Ed.). Philosophische Bibliothek.
- Mittelstadt, B. (2021). Near-term ethical challenges of digital twins. *Journal of Medical Ethics*, 47(6), 405–406. <https://doi.org/10.1136/medethics-2021-107449>
- Nyholm, S. (2021). Should a medical digital twin be viewed as an extension of the patient's body? *Journal of Medical Ethics*, 47(6), 401–402. <https://doi.org/10.1136/medethics-2021-107448>
- Onder, O., & Ilkilic, I. (2026). Can health technologies be care optimizers? A normative evaluation of digital health technologies in light of postphenomenological reflections. *Medicine Health Care and Philosophy*. <https://doi.org/10.1007/s11019-026-10337-3>
- Parsons, T. (1980). *Nonexistent objects*. Yale University Press.
- Priest, G. (2005). *Towards Non-Being: The Logic and Metaphysics of Intentionality*. Oxford University Press.
- Priest, G. (2016). *Towards non-being* (2nd ed.). Oxford University Press.
- Putnam, H. (2014). *Philosophy of logic (Routledge Revivals)*. Routledge. (Original work published 1971)
- Quine, W. V. O. (1948). On what there is. *Review of Metaphysics*, 2(5), 21–38.
- Quine, W. V. O. (1964). *Word and object*. MIT Press.
- Rosen, G. (1990). Modal fictionalism. *Mind*, 99(395), 327–354. <https://doi.org/10.1093/mind/xcix.395.327>
- Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale, & A. Hoffmann (Eds.), *Modality: Metaphysics, logic, and epistemology* (pp. 109–135). Oxford University Press.
- Routley, R. (1980). *Exploring Meinong's Jungle and Beyond: An investigation of noneism and the theory of items*. Australian National University.
- Russell, B. (1919). *Introduction to Mathematical Philosophy*. Dover Publications.
- Sainsbury, R. M., & Tye, M. (2012). *Seven puzzles of thought: And how to solve them: An originalist theory of concepts*. Oxford University Press.
- Schaffer, J. (2009). *On What Grounds What* (R. Wasserman, D. Manley, & D. Chalmers, Eds; pp. 347–383). Oxford University Press.
- Thomasson, A. L. (1999). *Fiction and metaphysics*. Cambridge University Press.
- Walton, K. L. (1990). *Mimesis as make-believe: On the foundations of the representational arts*. Harvard University Press.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University Press.
- Yablo, S. (2001). Go figure: A path through fictionalism. *Midwest Studies in Philosophy*, 25(1), 72–102. <https://doi.org/10.1111/1475-4975.00040>
- Yablo, S. (2005). The myth of the seven. In M. E. Kalderon (Ed.), *Fictionalism in metaphysics* (pp. 88–115). Oxford University Press.