Towards Linked Data for Ecosystems of Digital Twins

Samuele Burattini
Department of Computer Science and Engineering, Alma Mater Studiorum University of Bologna
Cesena, Italy
samuele.burattini@unibo.it

Antoine Zimmermann
Mines Saint-Étienne, Univ Clermont Auvergne, INP Clermont Auvergne, CNRS, UMR 6158 LIMOS
Saint-Étienne, France
antoine.zimmermann@emse.fr

Marco Picone
marco.picone@unimore.it
Department of Sciences and Methods for Engineering, University of Modena and Reggio Emilia
Reggio nell'Emilia, Emilia Romagna, Italy

Alessandro Ricci
Department of Computer Science and Engineering, Alma Mater Studiorum University of Bologna
Cesena, Italy
a.ricci@unibo.it

ABSTRACT

Due to either the inherent complexity of the domain or the evolving nature of systems, we can envision solutions that digitalize assets in a complex domain using an *ecosystem* of distributed Digital Twins instead of a single monolithic one. To effectively tackle interoperability in such ecosystems, this paper advocates for the introduction of a representation based on Semantic Web technologies enabling the discovery of both Digital Twin structure – i.e. the static information about the asset model and offered services – and state – i.e. the data and metrics collected at runtime – to support the management of ecosystems and the creation of application mashups. A review of the state of the art suggests that currently investigated ways to describe a Digital Twin are not sufficient to achieve this objective. A proposal of key requirements for a Digital Twin representation is outlined leading to the proposal of a core ontology and a Linked Data approach for state management.

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems**; • **Information systems** → **Semantic web description languages**; *Mashups*; **Ontologies**.

KEYWORDS

Digital Twins, Semantic Web, Interoperability, Ontologies

ACM Reference Format:
Samuele Burattini, Antoine Zimmermann, Marco Picone, and Alessandro Ricci. 2024. Towards Linked Data for Ecosystems of Digital Twins. In *ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems (MODELS Companion '24), September 22–27, 2024, Linz, Austria.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3652620.3688245

MODELS Companion '24, September 22–27, 2024, Linz, Austria

ACM ISBN 979-8-4007-0622-6/24/09
https://doi.org/10.1145/3652620.3688245

1 INTRODUCTION

Digital Twins (DTs) have risen from a concept tied to the product lifecycle management [13] to be a generic paradigm for the digitalization of physical assets (PAs) in many application domains [27]. Still, DT solutions mostly either focus on the mirroring of individual assets, failing to capture the relationships that exist among them, or adopt a monolithic approach in the digitalization of large-scale entities (e.g. a full city [8]) leading to issues concerning the scale and quality of the model and its evolvability.

Using multiple DTs to create a modular digital representation might be the key to effectively model large-scale scenarios – e.g., a nationwide Digital Twin [6] – and interoperability is considered the key challenge to achieving full maturity of the DT approach [20]. So-called *ecosystems* of Digital Twins [26] would then be composed of several entities – possibly connected by relationships – each mirroring crucial assets of an application domain and together achieving the digital representation of a complex system.

On the one hand, this would allow stakeholders to incrementally build the Digital Twin layer following their needs. The evolution of the DTs would be easier as they would be essentially independent from both a modelling and technological standpoint. The quality of the overall model would also improve, as it would be possible to keep track of the relationship between PAs and to model them at different relative scales for different application purposes, instead of relying on a single monolithic view. Moreover, this would open the way towards a vision of mashup applications that can use DT features *as-a-service* [26], even potentially leveraging DTs made available by different stakeholders and organizations belonging to different application domains.

On the other hand, managing the complex network of DTs opens up several challenges on how to integrate them effectively: with issues concerning data and service discoverability, deployment management, data and knowledge integration etc.

In this paper, following the trends from academia, industry and standardization bodies [9], we argue that a common representation model for DTs could address some of these issues. We investigate an approach based on Linked Data [17] principles and Semantic Web technologies to foster the adoption of a layer allowing Digital Twins to be represented as Web resources through a rich semantic model that could address interoperability concerns for the discoverability of DTs, their data and their exposed services. We ground our contribution with respect to well-known conceptual models for Digital Twins and present the main features and technologies tied to Linked Data in Section 2. We then survey the state-of-the-art Digital Twin representations in Section 3 to identify gaps that we relate to the core problem of interoperability in DT ecosystems. Finally, we propose two directions to handle the representation of Digital Twins – a core ontology and DT state management based on the Linked Data Platform – in Section 4 and highlight the main open challenges towards the achievement of the overall goal of interoperable ecosystems of DTs in Section 5 that we plan to explore in future works.

2 BACKGROUND

In this section, we will discuss the conceptual models of DTs that we take as reference to identify requirements for their representation. We further introduce the principles and technologies tied to the (Semantic) Web and Linked Data [17] highlighting their role in enabling interoperability in heterogeneous distributed systems.

2.1 Digital Twin Conceptual Models

Despite being a well-recognized concept, Digital Twins are still in a very dynamic evolution phase, with new proposals for conceptual models emerging from different domains and areas of application, ranging from computer science to manufacturing, medicine, etc.

Among those proposals, the five-dimensional model presented in [25] is becoming one of the most relevant in the area as it captures the essential features of a DT in a domain-agnostic fashion. It is composed by a *Physical Dimension* concerning the actual PA being mirrored, a *Connection Dimension* concerning the data flows from the real world to the DT (and vice versa) and towards external services, a *Data Dimension* that stores and process the data generated by the DT, a *Model Dimension* which is a virtual model of the properties and behaviour of the PA and a *Service Dimension* which enables external applications to access the data of the DT and exposes services such as prediction and simulation of the DT behaviour. This model already gives an intuition on the inherent complexity of capturing the full nature of a DT.

The need for a common meta-model for Digital Twins is an essential feature in the vision for DT ecosystems. In the proposal for a Web of Digital Twins [26] this is identified with the possibility of conceptualizing DTs as a set of *properties* the DT measures, *actions* the DT exposes to change the state of the corresponding asset, *events* that can notify observers of changes and *relationships* with other DTs. This very simple model is generally sufficient to represent the features of any kind of PA, enabling its representation in the digital world. Very similar models are adopted in the Internet of Things domain such as the W3C[1] Web of Things (WoT) with its *Thing Description* [22] and also in the data model [23] designed by Microsoft for its Azure Digital Twins cloud-based platform. This proves, empirically, its effectiveness in providing a general model to represent assets across different domains.

Despite such general models exist, for a DT to be interoperable in complex ecosystems an integrated representation of all its dimensions is needed.

[1] World Wide Web Consortium

2.2 Interoperability in the Web and Linked Data

The vision of the Semantic Web [4] consists of introducing new technologies, standards, and best practices to the Web to ease the exploitation of information at large online. In particular, Semantic Web technologies should enable machines to automatically process information from arbitrary sources to perform tasks on our behalf. A first step towards this vision is to build a Web of data that complements the Web of documents that people consume. While multimedia documents serve humans well, accompanying structured data simplifies information consumption by software systems. To improve the exploitation of data worldwide, standard data formats should be used across the Web. Explicit schemas facilitate processing further. Finally, ontologies help pinpoint the exact meaning of the data structures. These roughly outline the building blocks of the Semantic Web. As information on the Web often describes physical reality this enables machines to use the Semantic Web to understand reality and take action in it. In many ways, this goal aligns with the very core idea of DTs, making their overlap interesting to study.

We take the Web as inspiration because it is grounded on principles of interoperability between heterogeneous components, that are summarized in the REpresentational State Transfer (REST) constraints derived from the Web architecture [10]. REST suggests that interoperability is achieved through a uniform interface, which on the Web is built leveraging URIs to commonly identify resources and the manipulation of resources through the exchange of representations—using the semantics encoded in the HTTP protocol methods and negotiated serialization formats.

The Semantic Web enriches this idea by encoding semantics in the representation through the general model of the Resource Description Framework (RDF) [7]. This allows representations of resources to be expressed in machine-readable *triples* that state facts using a common vocabulary encoded within ontologies, themselves encoded in the Web Ontology Language (OWL) [24]. A set of triples defining relationships between resources and data constitutes a Knowledge Graph [18]. Linked Data [17], further completes the idea by applying a hypermedia-based design to RDF, using dereferenceable URIs as names for anything, and pointing those URIs at useful resources that describe the concept they are identifying. This essentially allows clients to follow them and discover additional information, by browsing it. In this vision, the whole Semantic Web becomes a distributed knowledge graph.

To further streamline how to manage Linked Data, the W3C produced a recommendation in the form of the Linked Data Platform (LDP) [29] that provides patterns to handle Linked Data using RESTful interactions and defines the concept of a *container* of RDF resources. LDP Containers can be used to group resources and guide clients in their manipulation with standard behaviour as when an LDP Container is discovered, it is then possible to get the list of contained resources, request their representations and (optionally) modify them or create new ones.

3 DIGITAL TWIN REPRESENTATIONS

In this section, we first define what are the desirable features a DT representation should offer drawing from the conceptual models of DTs and with regard to the problem of interoperability in DT ecosystems that we here summarize as:

> Given an ecosystem of DTs we want to allow consumers to easily discover what DTs are present in it, what assets they mirror and how they are mirroring them, the services that the DTs expose and the real-time and past data the DTs are gathering, regardless of the underlying implementing technology.

We then analyze some notable works in the area to highlight how the idea of introducing a level of representation of DTs is well-recognized in the literature and industry, but at the same time, no existing proposal makes it possible to capture both the structure and the state of a DT in a coherent framework.

3.1 Requirements for Digital Twins Representations

When adopting a vision for ecosystems of Digital Twins where several entities of the target domain are digitalized, each through individual DTs, the need to have suitable representations for such DTs emerges quite naturally to achieve several objectives. This is reflected by DT standardization efforts [9] that are advocating for descriptions and representations to improve interoperability among different DTs. In this section we will analyze key features that make such a level of representation desirable and will use them as requirements for an ideal representation model for a Digital Twin, to later compare with existing related works.

Two main perspectives are adopted, on one hand, the representation is useful for DT *consumers* i.e., users and application developers that may wish to access the DT data, act on the physical world through the DT or leverage the services offered by the DT. On the other hand, the representation can be useful for DT *operators* i.e. DT developers and maintainers that need to manage a complex ecosystem of running instances of DTs, update their models, verify their state of operation and so on. For both kinds of users, the representations offer a uniform interface to access information about the DT, regardless of the underlying implementing technology.

Identification. The need to identify uniquely a DT is essential when considering open ecosystems. Moreover, the DT should clearly identify also the unique PA it is representing, as more DTs of the same asset may exist with potentially very different models.

Physical Asset Modelling. As different DTs for the same asset can be built, the representation of a unique DT should also include information about the models it is offering for the asset. An example of this could be in terms of which properties of the asset are measured and what temporal scale the DT is considered to update their values. Different models of the same room could very well in fact monitor the current temperature, the average temperature of the last hour or the energy consumption.

Physical Asset State. As DT digitalize PAs, one of the core objectives in their representation is making the real-time and past state of the mirrored entity available to consumers. This allows querying for assets that are in a given state at any given time (e.g. shutting off all devices that are consuming energy over a given threshold).

Service Description. DTs offer a bi-directional interaction with PAs which allows both to monitor and control the physical reality. Often, their models can also be used to predict or simulate the behaviour of the asset, and in some cases, visualization of the PA state is offered to allow human operators to understand the behaviour of the asset in real time. In open ecosystems of DTs, all of these features can be exposed as *services* for consumers. The description of these services is then essential to discovering the capabilities of a DT and using it accordingly, enabling building mashup applications on top of newly discovered DTs, leveraging their services.

Deployment Context. Technological details concerning how the DT is developed, alongside its version, authors and deployment requirements (e.g. dependency on MQTT brokers, GPU usage) give additional information to system administrators needing to update the DT models, or deploy instances of DTs. Consumers may also choose among different replicas of the same DT the ones that are deployed *closer* to either the source of data or the application depending on their needs.

Metrics and State of Operation. Metrics concerning its coupling with regard to the PA (e.g. [3]) allow for assessing the ability of the DT to effectively represent the PA in real time. Moreover, the lifecycle of a DT can be fairly complex, and a DT could sometimes be disconnected from its physical counterpart. The state of operation of the DT is essential for operators who may need to investigate bugs and consumers who may need to understand that the DT is offline and cannot be used to control the PA.

Data Sources and Storage. Describing how data is obtained – i.e., by means of which sources – allows keeping track of provenance and eases the management of the systems supporting the DT. Additionally, a DT may point to dedicated storage which allows consumers to perform data-intensive operations.

3.2 State of the Art Analysis

Having identified the requirements for Digital Twin representations, we believe a complete model should capture all of them in a coherent representation. We then compare relevant works in the state of the art for DT representations with regard to the support of those requirements. We selected them based on their relevance to the topic of creating descriptions for ecosystems of Digital Twins picking from well-known ontologies and data models used in the context of Digital Twins from both industry and academia and selecting among recent works with the keywords "Digital Twins" and "Ontology". Table 1 reports the main analysed works. We do not claim this collection to be exhaustive, but we believe it is a sufficient selection to highlight the open gaps in the area. Additionally, our interest was specifically in works aiming to achieve a description that makes the DT easier to operate with for external consumers. We refer interested readers to a recent survey [19] for a more general analysis of the use of ontologies in Digital Twins.

The Digital Twin Definition Language (DTDL) [23] is an RDF-based model for the Azure Digital Twins[2] platform by Microsoft. Its model allows to define models for DTs in terms of their properties, commands and relationships with other twins and to instantiate

[2]https://azure.microsoft.com/en-us/products/digital-twins

Table 1: Surveyed works in the state of the art on DT representations, with the respective features they support: (✓) means the feature is directly supported, (~) means the feature can be represented within the framework but not explicitly addressed or not completely supported, (×) means the feature is not supported.

Representation	Identification	PA Modelling	PA State	Services	Deployment	Metrics	Data
Azure DTDL [23]	~	✓	✓	×	×	~	×
WoT Thing Description [22]	~	✓	×	✓	×	~	×
ETSI SAREF [11]	×	✓	✓	~	×	×	×
Barros et al. [1]	✓	×	×	×	×	×	×
Barth et al.[2]	✓	×	×	~	~	×	✓
Singh et al. [28]	✓	✓	✓	×	×	×	✓
Steinmetz et al. [30]	✓	✓	×	✓	×	×	✓
Gonzalez et al.[12]	✓	~	~	✓	×	×	✓

them within the platform, storing the last updated values of the properties in a so-called *twin graph*. Even though each instance has an identifier, the represented PA is not identified explicitly. Azure DT keeps track of the timestamps of the updates of each properties automatically, but does not represent any other metric explicitly, nor it allows to indicate data provenance.

The W3C WoT Thing Description (TD) [22] is a model to describe connected devices under the abstraction of *things* offering interaction affordances to read properties, execute actions or register to events. A TD identifies ambiguously both a device or a software entity implementing the *thing*—that as per the WoT architecture [21] could also be a Digital Twin. Interaction affordances in a TD are equipped with forms to detail how to invoke them, using a specific communication protocol, they thus allow new consumers to understand how to act on the thing and can represent services to interact with the DT.

The Smart Application REference (SAREF) ontology [11] by the European Telecommunications Standards Institute (ETSI), is an ontology for IoT applications, that promotes interoperability across several domains. It suits the purpose of representing DTs and a dedicated group is investigating how to adapt it better for the purpose [9]. SAREF includes concepts such as services, properties and observations to collect the value of specific properties making it a good candidate to represent the model and state of a PA. It is not clear though how it would be possible to identify the software counterpart of the DT, nor how to decorate it with information about the deployment context, metrics or data sources.

Other proposals from the academia [1, 2, 28, 30] have different objectives with regard to the level of description of a DT. While [1] attempts to link the concept of Digital Twins with foundational ontologies [16] achieving a very high-level description, [2] attempts a systematization with respect to the data sources that compose a DT and its value on the target context. Differently [28] and [30] focus on the IoT domain, with the former focusing on data, whereas the latter focusing on the services and APIs that a DT offers, extending the IoTLite[3] ontology to explicitly identify DTs.

A notable work is the recent model proposed by Gonzalez et al. [12] that extended the WoT TD ontology to represent Digital Twins, explicitly modelling the five-dimensions of DTs [25]. With this extension, they can represent the relationship between a DT and its physical counterpart, despite not providing detailed instructions on describing its model or state. Services are represented through TD interaction affordances, and data sources and repositories linked to the DT are represented explicitly through the DCAT[4] ontology, allowing consumers to access also non-semantic data.

4 PROPOSING SEMANTIC DIGITAL TWIN REPRESENTATIONS

Having identified requirements for interoperability through DT representations that capture the DT structure and state and recognized that existing models generally fall short in representing DTs in their entirety, we here propose an approach based on Linked Data to represent DTs for interoperable ecosystems.

The approach is based on two main elements: the first is a draft of a core ontology to explicitly identify concepts that are required to represent a DT, the second is an approach to handle browsable DT history using a Linked Data approach.

4.1 Towards a Core Ontology for Digital Twins

From the requirements outlined in the previous section, we here sketch a model for a core ontology that can support the definition of DT representations (Figure 1). Developing an ontology is a long iterative process and as the goal is to support interoperability among the most possible types of DTs we just show here a tentative model that we plan to continue refining in future works.

First, we distinguish clearly the DT from the PA it is mirroring. Each has a state (and a history, see section 4.2). The state of the PA includes data about the asset, whereas the state of the DT includes its state of operation, metrics and deployment context. The DT also has a model, that is described in terms of the features that compose it (properties, actions and events). Finally, the DT will expose Services, that can be related to other components of the DT – e.g. the service for subscribing to an event or invoking an action that could be described with a WoT Thing Description [22] – or not – e.g. a simulation service that is generic for the whole DT.

We believe this simple model to be a good start towards the definition of a core ontology, although it will need further refinement and testing with different kinds of DTs. Also, we present this as a core model, since we expect the development of extensions

[3]https://www.w3.org/submissions//iot-lite/

[4]Data Catalog Vocabulary: https://www.w3.org/TR/vocab-dcat-3/

for specific kinds of DTs. Still, having a core root that can already address most requirements can drive interoperability to build DT ecosystems.

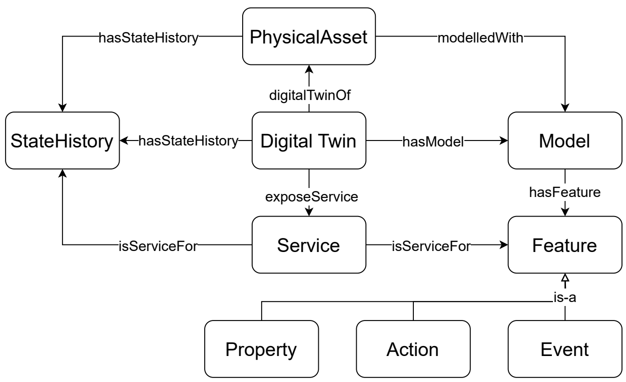

Figure 1: A schema of the entities for a draft core ontology for DT representations. Relationships are not necessarily 1:1 (i.e., a DT can have multiple Models).

4.2 A Linked Data Approach for State History

Digital Twins by definition evolve over time to reflect the changes of their physical counterparts. As such, their representation must evolve as well. As can be seen from Table 1, most existing approaches present a level of description that does not include the state of the PA. Nevertheless, this is valuable as it allows clients to directly interrogate the ecosystem of DTs with queries on the current state of a system of DTs. This is supported, for instance, in the Azure Digital Twins platform that offers a representation of a DT in terms of the latest values of its properties, but not its history which is delegated to a separate data storage and as such needs to be accessed with a different tool and query language.

A different approach is instead adopted in other kinds of descriptions such as the one that can be built with the SAREF ontology [11] that includes an `Observation` as a first-class concept to represent the operation that leads to collecting a measurement of some property value at a given time. Managing the evolving state of the asset would then mean adding new observations and property values to the same description. This would make the representation more complete, but also heavier: better for querying but less effective when simply browsing the ecosystem of DTs.

In general, handling the evolution of knowledge graphs over time is a complex subject, but necessary if one wants to keep the representation of a DT up to date. At any given time the representation of the DT should give a consistent and updated representation of its state, while at the same time offering the possibility to browse back in time. For this reason, we envision supporting such a mechanism using the rules of the Linked Data Platform [29] to host and allow access to the data of a DT.

With every state update, the DT will produce a *snapshot graph* that describes its state at a given time. The DT will then be granted authorized access to POST it as a resource in an LDP Container. If intended as a *named graph* [5] each snapshot of the DT state can further be described in terms of the time frame it represents

Listing 1: An example of how the history of the states of the DT can be represented using LDP containers

```
@base <http://localhost:8080/>.
@prefix ldp: <http://www.w3.org/ns/ldp#>.
@prefix dto: <http://example.org/dtOntology#>.
@prefix time: <http://www.w3.org/2006/time#>.
@prefix dcterms: <http://purl.org/dc/terms/>.

<myDT> a dto:DigitalTwin;
   dto:hasStateHistory <myDT/stateContainer>.

##Container Description##

<myDT/stateContainer> a dto:StateHistory;
   a ldp:BasicContainer;
   dcterms:title "The state of the DT";
   ldp:contains <s1>, <s2>, <s3>.

<s1> time:hasTime _:t1.
<s2> time:hasTime _:t2.
<s3> time:hasTime _:t3.

_:t1 time:before _:t2.
_:t2 time:before _:t3.
```

and with an ordering relationship with the other ones, creating a chain of immutable representations that consumers can browse (Listing 1). The advantage of using the LDP standard would offer a uniform interface for browsing the DT state history.

It must be noted that the semantics of named graphs that we adopted in this paper are still an object of debate in the Semantic Web community[5] We consider this an interesting use case that motivates the need to allow the attribution of names to a graph to further describe it as a whole.

Of course, this practice is not meant to replace more efficient data storage solutions (e.g. time series databases) that might still be more performing when requiring data in bulk for analytics purposes. This would allow though to have uniform, Linked Data-compliant access to all the features of a DT, enabling browsing through time with pure RESTful interactions.

5 OPEN CHALLENGES AND FUTURE WORKS

In this paper, we have outlined the requirements for representations of DTs while showing how existing works fail to capture all of them. We then suggested how representations based on Linked Data principles can address such requirements: We advocate for a core ontology for DTs and propose a Linked Data approach to manage DT history under a uniform interface that can be browsed by consumers.

The vision of interoperable ecosystems of DTs comes with several challenges, both at the technical and societal levels. Here we highlight the main ones and comment on future works.

At the societal level, stakeholders may not be willing to open the data and services of DTs to other consumers, limiting interoperability and preferring the creation of closed systems, especially when DTs are built by companies that sell both the hardware and software counterparts. At the same time, institutions have already shown interest in the realization of DTs for the management of the public infrastructure and open ecosystems of DTs might improve transparency in public settings. Of course, for this framework to

[5]see https://www.w3.org/TR/2014/NOTE-rdf11-datasets-20140225/ for an in-depth discussion

be effective, it must be possible to allow fine-grained access control to data and guarantee security and privacy. We believe the Solid[6] project to be an interesting direction to achieve this goal in decentralized Linked Data ecosystems.

On the technical level, the main challenges concern the distributed nature of ecosystems of DTs which makes it difficult to create links between DTs, especially when such links are dynamic relationships. Managing the huge amount of data DTs produce, also challenges the possibility of storing them and interrogating them effectively with knowledge graphs. We believe a hybrid approach to be a plausible solution with DTs' representations pointing to dedicated data storages that use more suitable formats for bulk access. Finally, generating the representations for DTs could be a daunting task for developers, and limit their applicability in practice. We believe development tools should ease this task. A basic core ontology could provide shared structure across different DT development frameworks, leaving only the more domain-oriented aspects to consider for DT developers.

Addressing interoperability is a multi-faceted problem. On one hand, using semantic technologies does not grant interoperability on its own [15] and at the same time, we acknowledge that a core ontology may create a false-agreement problem [14]. Nevertheless, this direction can be a step forward: in future works, we plan to continue developing such core ontology testing the representation of DTs in practical use cases, paired with the Linked Data Platform container approach of *snapshot graphs* for DT history to evaluate its effectiveness. Finally, an essential step will be the creation of alignments with other existing ontologies, to better support representations for DTs across different domains.

ACKNOWLEDGMENTS

This work is partially supported by the project "Digital Twins Ecosystems for the clinical, strategic and process governance in Healthcare" funded by the European Union - NextGenerationEU and AUSL della Romagna (PNRR M4C2, Investment 3.3 DM 352/2022) - CUP J33C22001400009

REFERENCES

[1] Claudio Barros, Rebecca Salles, Eduardo Ogasawara, Giancarlo Guizzardi, and Fabio Porto. 2021. Requirements for an ontology of digital twins. In *CEUR workshop proceedings*, Vol. 2941. Rheinisch Westfälische Technische Hochschule.

[2] Linard Barth, Matthias Ehrat, Rainer Fuchs, and Jens Haarmann. 2020. Systematization of Digital Twins: Ontology and Conceptual Framework. In *Proceedings of the 3rd International Conference on Information Science and Systems (ICISS '20)*. Association for Computing Machinery, New York, NY, USA, 13–23. https://doi.org/10.1145/3388176.3388209

[3] Paolo Bellavista, Nicola Bicocchi, Mattia Fogli, Carlo Giannelli, Marco Mamei, and Marco Picone. 2024. ODTE: A Metric for Digital Twin Entanglement. *IEEE Open Journal of the Communications Society* 5 (2024), 2377–2390. https://doi.org/10.1109/OJCOMS.2024.3385659

[4] T. Berners-Lee, J. Hendler, and O. Lassila. 2001. The Semantic Web. *Scientific American* 284, 5 (May 2001), 34–43. https://doi.org/10.1038/scientificamerican0501-34

[5] Jeremy J. Carroll, Christian Bizer, Patrick J. Hayes, and Patrick Stickler. 2005. Named graphs. *Journal of Web Semantics* 3, 4 (2005), 247–267.

[6] Gemini Council and Kirsten Lamb. 2022. *Gemini Papers: Summary.* Technical Report. Centre for Digital Built Britain.

[7] Richard Cyganiak, David Wood, and Markus Lanthaler. 2014. *RDF 1.1 Concepts and Abstract Syntax, W3C Recommendation 25 February 2014.* W3C Recommendation. World Wide Web Consortium. http://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/

[8] Li Deren, Yu Wenbo, and Shao Zhenfeng. 2021. Smart city based on digital twins. *Computational Urban Science* 1 (2021), 1–11.

[9] ETSI Specialist Task Forces (STF) 628. 2024. *SmartM2M; Digital Twins Communication Requirements.* TS TS-103-845-V1.1.1. European Telecommunications Standards Institute (ETSI).

[10] Roy T. Fielding and Richard N. Taylor. 2002. Principled Design of the Modern Web Architecture. *ACM Trans. Internet Technol.* 2, 2 (May 2002), 115–150. https://doi.org/10.1145/514183.514185

[11] Raúl García-Castro, Maxime Lefrançois, María Poveda-Villalón, and Laura Daniele. 2023. The ETSI SAREF ontology for smart applications: a long path of development and evolution. *Energy Smart Appliances: Applications, Methodologies, and Challenges* (2023), 183–215.

[12] Salvador González-Gerpe, Andrea Cimmino, Socorro Bernardos, María Poveda-Villalón, and Raúl García-Castro. 2024. WoTDT: an Extension of the WoT Thing Description Ontology for Digital Twins in the Construction Domain. In *LDAC'24: 12th Linked Data in Architecture and Construction Workshop.*

[13] Michael Grieves and John Vickers. 2017. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. *Transdisciplinary perspectives on complex systems: New findings and approaches* (2017), 85–113.

[14] Nicola Guarino. 1998. *Formal ontology in information systems: Proceedings of the first international conference (FOIS'98), June 6-8, Trento, Italy.* Vol. 46. IOS press.

[15] Giancarlo Guizzardi. 2020. Ontology, Ontologies and the "I" of FAIR. *Data Intell.* 2, 1-2 (2020), 181–191. https://doi.org/10.1162/DINT_A_00040

[16] Giancarlo Guizzardi, Alessander Botti Benevides, Claudenir M. Fonseca, Daniele Porello, João Paulo A. Almeida, and Tiago Prince Sales. 2022. UFO: Unified Foundational Ontology. *Appl. Ontology* 17, 1 (2022), 167–210. https://doi.org/10.3233/AO-210256

[17] Tom Heath and Christian Bizer. 2011. *Linked Data: Evolving the Web into a Global Data Space.* Morgan & Claypool. http://linkeddatabook.com/editions/1.0/

[18] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d'Amato, Gerard de Melo, Claudio Gutiérrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan F. Sequeda, Steffen Staab, and Antoine Zimmermann. 2021. *Knowledge Graphs.* Number 22 in Synthesis Lectures on Data, Semantics, and Knowledge. Springer. https://doi.org/10.2200/S01125ED1V01Y202109DSK022

[19] Erkan Karabulut, Salvatore F. Pileggi, Paul Groth, and Victoria Degeler. 2024. Ontologies in digital twins: A systematic literature review. *Future Generation Computer Systems* 153 (April 2024), 442–456. https://doi.org/10.1016/j.future.2023.12.013

[20] Robert Klar, Niklas Arvidsson, and Vangelis Angelakis. 2024. Digital Twins' Maturity: The Need for Interoperability. *IEEE Syst. J.* 18, 1 (2024), 713–724. https://doi.org/10.1109/JSYST.2023.3340422

[21] Ryuichi Matsukura, Michael McCool, Kunihiko Toumura, and Michael Lagally. 2023. *Web of Things (WoT) Architecture 1.1.* W3C Recommendation. W3C. https://www.w3.org/TR/2023/REC-wot-architecture11-20231205/.

[22] Michael McCool, Ege Korkan, and Sebastian Käbisch. 2023. *Web of Things (WoT) Thing Description 1.1.* W3C Recommendation. W3C. https://www.w3.org/TR/2023/REC-wot-thing-description11-20231205/

[23] Microsoft. 2024. Digital Twins Definition Language (DTDL) v3. https://azure.github.io/opendigitaltwins-dtdl/DTDL/v3/DTDL.v3.html.

[24] Boris Motik, Peter F. Patel-Schneider, and Bijan Parsia. 2012. *OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax (Second Edition), W3C Recommendation 11 December 2012.* W3C Recommendation. World Wide Web Consortium. http://www.w3.org/TR/2012/REC-owl2-syntax-20121211/

[25] Qinglin Qi, Fei Tao, Tianliang Hu, Nabil Anwer, Ang Liu, Yongli Wei, Lihui Wang, and A.Y.C. Nee. 2021. Enabling technologies and tools for digital twin. *Journal of Manufacturing Systems* 58 (2021), 3–21. https://doi.org/10.1016/j.jmsy.2019.10.001 Digital Twin towards Smart Manufacturing and Industry 4.0.

[26] Alessandro Ricci, Angelo Croatti, Stefano Mariani, Sara Montagna, and Marco Picone. 2022. Web of Digital Twins. *ACM Transaction on Internet Technology* 22, 4 (2022), 101:1–101:30. https://doi.org/10.1145/3507909

[27] Concetta Semeraro, Mario Lezoche, Hervé Panetto, and Michele Dassisti. 2021. Digital twin paradigm: A systematic literature review. *Comput. Ind.* 130 (2021), 103469. https://doi.org/10.1016/J.COMPIND.2021.103469

[28] Sumit Singh, Essam Shehab, Nigel Higgins, Kevin Fowler, Dylan Reynolds, John A Erkoyuncu, and Peter Gadd. 2021. Data management for developing digital twin ontology model. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture* 235, 14 (Dec. 2021), 2323–2337. https://doi.org/10.1177/0954405420978117 Publisher: IMECHE.

[29] Steve Speicher, John Arwe, and Ashok Malhotra. 2015. *Linked Data Platform 1.0.* W3C Recommendation. W3C. https://www.w3.org/TR/2015/REC-ldp-20150226/.

[30] Charles Steinmetz, Achim Rettberg, Fabiola Goncalves C. Ribeiro, Greyce Schroeder, and Carlos E. Pereira. 2018. Internet of Things Ontology for Digital Twin in Cyber Physical Systems. In *2018 VIII Brazilian Symposium on Computing Systems Engineering (SBESC).* IEEE, Salvador, Brazil, 154–159. https://doi.org/10.1109/SBESC.2018.00030

[6] https://solidproject.org/