www.itcon.org - Journal of Information Technology in Construction - ISSN 1874-4753

A KNOWLEDGE GRAPH-BASED DIGITAL TWIN SYSTEM FOR BRIDGE NETWORK CONDITION MONITORING

SUBMITTED: June 2025
PUBLISHED: August 2026
EDITOR: Robert Amor
DOI: 10.36680/j.itcon.2026.036

Carlos Ramonell Cazador, PhD candidate
Universitat Politècnica de Catalunya
https://orcid.org/0000-0003-0390-3603
carlos.ramonell@upc.edu

Pieter Pauwels, Associate Professor
Eindhoven University of Technology
https://orcid.org/0000-0001-8020-4609
p.pauwels@tue.nl

Rolando Chacón, Associate Professor
Universitat Politècnica de Catalunya
https://orcid.org/0000-0002-7259-5635
rolando.chacon@upc.edu

SUMMARY: *While a lot of research has examined semantics and digital twinning (DT), their use in bridge network monitoring has received less attention. In particular, the lifecycle management of information on bridge infrastructure is a critical challenge, and this paper aims to address it through a linked data framework that meets current needs and aligns with common practices in bridge maintenance. The paper proposes a technological framework based on knowledge graphs, open linked data ontologies, and industry standards such as IFC, INSPIRE, and ICDD to achieve multi-scale data integration, interoperability and standard data delivery. This framework serves as the baseline architecture for a knowledge graph-based DT system. It uses a microservices approach, assembling a suite of open-source Extract, Transform and Load (ETL) components, a graph database, and a file database. The application on a bridge network in the Metropolitan Area of Barcelona demonstrates that the proposed framework automatically abstracts Bridge Condition Indices (BCI) at the network level by aggregating state updates from individual bridge components. Where official inspection scoring procedures were unavailable, these indices utilise heuristic approximations to illustrate the prototype workflow. Furthermore, the system enables the resolution of context-rich queries that combine GIS, BIM, and inspection data. This study demonstrates how Digital Twins can be conceptualised and developed to operate with standardised, flexible, modular, interoperable, and multi-scale data, setting the groundwork for future Bridge Management Systems (BMS).*

KEYWORDS: *Digital twin, BMS, knowledge graphs, ontologies, ICDD.*

REFERENCE: *Ramonell Cazador, C., Pauwels, P., & Chacón, R. (2026). A knowledge graph-based digital twin system for bridge network condition monitoring. Journal of Information Technology in Construction (ITcon), 31, 835-865. https://doi.org/10.36680/j.itcon.2026.036*

1. INTRODUCTION

The socio-economic development of any region significantly depends on the scope, efficiency, and resilience of its infrastructure systems. Transportation networks, in particular, are paramount. Within these networks, bridges serve as crucial links, reducing travel times and enhancing accessibility across diverse geographical landscapes. However, bridges are especially prone to deterioration. Traffic loads and environmental phenomena continuously stress them. In the event of damage or collapse, the impact is significant, and recovery times can jeopardise regional economies. Therefore, it is necessary to maintain consistent monitoring of bridge conditions to uphold their structural integrity efficiently and, in turn, the serviceability of the whole network. Nowadays, infrastructure managers are facing significant challenges in maintaining bridge networks. Many bridges are ageing simultaneously, resulting in a collective reduction of their structural capacity. An unforeseen increase in traffic over the past few decades has subjected numerous bridges to loads occasionally exceeding those for which they were originally designed. Additionally, environmental uncertainty increases the perceived risk of failure and the lack of resilience in infrastructure systems, prompting the need for more extensive monitoring of bridge network conditions.

This necessitates faster, more detailed information gathering across large territories, which must be processed and interpreted to assist in decision-making, all while operating with limited economic resources. Bridge Management Systems (BMS) are used for this purpose. However, these systems have limitations in their conformance with current information modelling and management standards, as well as in their capacity to flexibly incorporate emerging information sources within a rapidly evolving AECO (Architecture, Engineering, Construction, and Operations) digital ecosystem. This has led to the exploration of advanced digital solutions (Gkoumas et al., 2019), among which the concept of semantic digital replicas of built assets, i.e., semantic Digital Twins. This exploration facilitates the development of new information systems that can overcome information management limitations, as also pointed out by Boje et al. (2020).

This paper introduces the concept of a semantic Digital Twin to the bridge maintenance domain. It presents a generalizable approach to developing Digital Twin systems to guide the development of future BMS systems. This paper serves as a reference for integrating domain knowledge with standard Geographic Information Systems (GIS) and Building Information Modelling (BIM) data, as well as digitised versions of standardised inspection protocols. This integration is achieved through a knowledge graph (Hogan et al., 2021) based on openly available linked data ontologies, which incorporates information exchange mechanisms compliant with the Industry Foundation Classes (IFC) and Information Containers for linked Document Delivery (ICDD) standards, enabling stakeholders to deliver information to DT systems seamlessly. This paper thus assumes knowledge graphs in the Resource Description Framework (RDF) data modelling paradigm. Results of this research demonstrate how existing technologies and standards can support more comprehensive, flexible, and interoperable BMS.

Accordingly, the contribution of this paper is threefold, spanning the conceptual, technological, and operational dimensions of implementing semantic digital twins for bridge maintenance. Conceptually, the study bridges the gap between the information lifecycle of built assets, the requirements for bridge maintenance and the theoretical capabilities of Digital Twins, providing a robust framework for their application in the built environment. Technologically, it demonstrates that a knowledge graph-based approach is instrumental for the semantic orchestration of heterogeneous data, integrating BIM, GIS, and inspection records while ensuring data provenance and long-term interoperability using open standards. Additionally, the research offers a reproducible methodology and an open-source toolkit, including custom-developed ontologies and ETL tools. By making these software components and data models publicly available, this work provides a functional reference that researchers and practitioners can use to replicate the developments described throughout the study.

The remainder of the paper is organised as follows: Section 2 discusses the implementation of digital twins in the built environment, outlines the fundamental concepts of the bridge maintenance domain, and introduces knowledge graphs as the key technological enabler for DT systems. Section 3 presents the design and architecture of the DT system. Section 4 presents the implementation of the DT system in a case study, based on a simulation of a series of scenarios representing various bridge information management scenarios, including the delivery of an as-built Bridge Information Model (BrIM) after construction, the integration of inspection data for a single bridge, and the updating and querying of inspections and bridge state information across a bridge network. Finally, Section 5 provides concluding remarks.

2. BACKGROUND

The following sections outline the fundamental concepts and state-of-the-art technologies that underpin this research. Section 2.1 discusses the evolution of Digital Twins in the AECO industry and identifies critical gaps in current applications. Section 2.2 introduces Knowledge Graphs and ontologies as the primary mechanisms for achieving semantic interoperability and provides a detailed overview of existing digital frameworks for modelling complex bridge structures.

2.1 Digital twins in the built environment

This section discusses the current state of Digital Twins within the AECO industry. It defines the lifecycle expectations for these digital replicas and identifies the systemic barriers, both technical and cultural, that currently limit their widespread implementation.

2.1.1 Expectations and requirements

Digital Twins in the built environment are envisioned as digital replicas of built assets that harmonise information production and management throughout their entire life cycle (Boje et al., 2020; Sacks et al., 2020). They should be conceived during the design phase, enriched and verified against their physical counterparts during construction, and utilised and updated during the operation and maintenance (O&M) stage. This process is illustrated in Fig. 1.

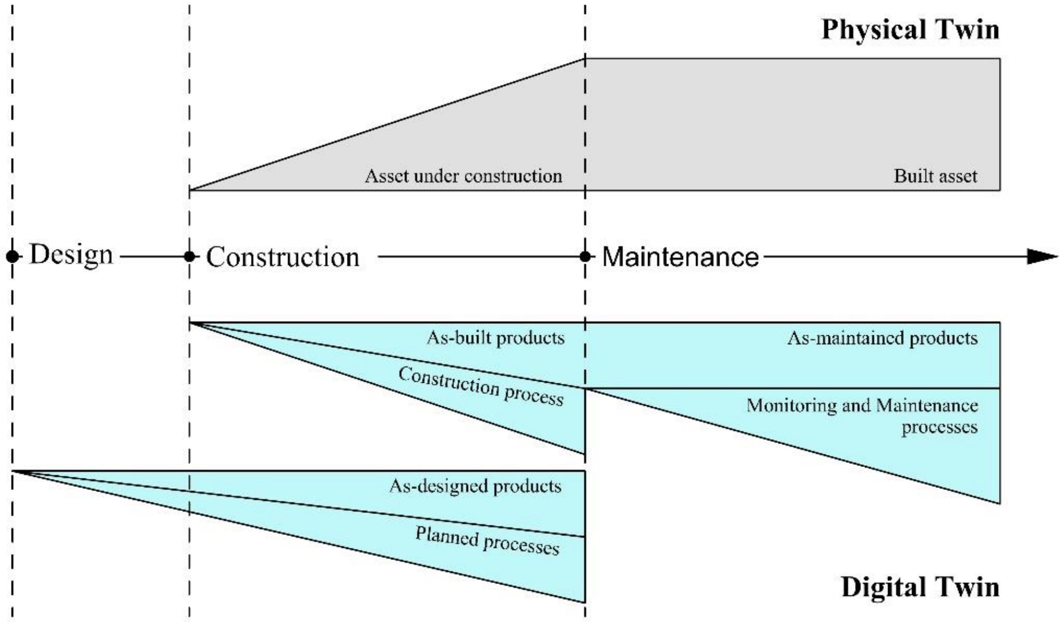

Figure 1: Digital Twin across the asset life cycle. Adapted from Sacks et al. (2020).

In this process, the DT begins with an as-designed digital model conforming to BIM standards. Before construction begins, this model would ideally be enriched with data related to planned construction processes. There is considerable research interest in establishing DT-based systems for construction sites. These twins aim to streamline the control of construction progress. The literature refers to these systems as Digital Twin Construction (DTC) (Boje et al., 2020; Sacks et al., 2020). Successful application of DTC involves establishing systematic information pipelines that collect data from actual construction processes and as-built components and interpreting this data to compare it with planned processes and as-designed components (Pal et al., 2023; Chacón et al., 2024). This approach aims to enhance construction managers' awareness of construction processes, enabling better-informed decisions that improve overall efficiency and safety on sites. Additionally, DTC ensures that the most reliable and accurate as-built representation of the asset is delivered to asset managers.

Once construction is finished, information is handed over to asset managers. Adequately organised as-built information is expected to provide a long-lasting repository of data. This information construct ideally integrates various forms of monitoring data, maintenance-related information, and relevant contextual geographical information to provide asset managers with a reliable basis for decision-making during the O&M phase (Chacón et al., 2023). This requires Digital Twin systems that can integrate multiple sources of information (Ramonell et al., 2023). Digital twin systems need to provide stakeholders with controlled access to this information and enable updates as new data is obtained, allowing for continuous tracking of asset state and behaviour at a required frequency. Trends observed from accumulated monitoring and maintenance data form the basis for developing predictions and decision support for asset managers. Thus, a digital twin during the O&M phase is expected to serve as a mirror of the asset's current state, a reliable archive of interconnected digital information describing its history, and a window through which to forecast its near-future potential states. Additionally, built assets are part of a more complex and interconnected environment within cities or infrastructure networks. Asset managers would eventually need to aggregate and integrate information from multiple digital twins to establish management strategies that span different scales, from individual building components to entire buildings and, ultimately, cities or asset networks (Lu et al., 2020; Zhang et al., 2022).

The DT data and system components should be defined based on their scope, considering the type of asset being represented, its intended users, and the workflows it will support (Lu and Brilakis, 2019; Pregnolato et al., 2022). Furthermore, the implementation context significantly influences their development; this requires consistency with regional maintenance and management practices that align with international, national, and regional standards and guidelines.

2.1.2 Challenges

The primary challenges in implementing DTs stem from the fragmented organisational structure and the cultural disconnect from digital technologies in the AECO industry. There is a generalised lack of fundamental digital skills related to programming languages, data science, and software system development (Sacks et al., 2020; Tuhaise et al., 2023; Broo and Schooling, 2023; Posada et al., 2024), and the combination of both digital skills and domain knowledge is essential to developing digital twin solutions. The consequences of this situation are shown in DT-related research. A significant number of studies remain theoretical, reiterating similar definitions, challenges, and topics without significant advancements in recent years (Halmetoja, 2022; Broo and Schooling, 2023; Albalkhy et al., 2024). Other studies provide technical implementations that rely on connecting various existing off-the-shelf tools through plug-ins or Application Programming Interfaces (APIs), resulting in approaches that are difficult to generalise (Tuhaise et al., 2023; Albalkhy et al., 2024).

These cultural conditions impede the resolution of the technical challenges that need to be addressed for the development of DT systems:

- **Transition to Web-based Environments:** The web is the ideal platform for real-time information exchange among industry stakeholders. Modern web-based applications rely on dynamic, lightweight information transactions with cloud-based systems, facilitated via Web APIs. These APIs interface with system functionality and allow querying and filtering of large amounts of information in databases. Information is transferred using encoding formats such as JavaScript Object Notation (JSON) (see https://www.json.org/), Extensible Markup Language (XML) (see https://www.w3.org/XML/) or Resource Description Framework (RDF) (see https://www.w3.org/RDF/). In contrast, most of the data managed by industry stakeholders is stored in siloed files, and much of it is not even machine-readable (Pauwels et al., 2017b). Other data may contain machine-readable information, such as BIM models, but the format is not openly accessible or suitable for web environments (Pauwels et al., 2022). More data-driven, granular information management and exchange through the web is required.
- **Lack of Shared Data Models (Interoperability):** Data interoperability is a requirement in the AECO industry to enable collaboration among stakeholders who use different software tools. A digital twin system is yet another software solution that requires combining diverse data sources. While open BIM data modelling standards such as IFC enable interoperability among BIM software solutions, they cover only a portion of the data needed to support the management of a built asset. The information in an IFC model is a key source for representing a built asset, but it must be supplemented with many additional layers of information (Boje et al., 2020). To meet interoperability requirements in the industry, these data layers should be structured according to open data models shared across industry software. This approach would result in an ecosystem of models that can be flexibly assembled, enabling them to adapt to the needs of their final users.

- **Data Integration:** Assuming all data layers within a digital twin to be machine-readable, interoperable, and adapted to web environments, they must also be semantically interrelated. Such semantic structuring enables applications to access, interpret, and integrate data efficiently, leveraging these interrelations to generate meaningful insights and support advanced functionalities for specific use cases. For this purpose, ontologies and knowledge graphs are widely adopted across many domains (Chaudhri et al., 2022) and can also be used to integrate domain-specific data to represent the built environment (Pauwels et al., 2021).

2.2 Knowledge graphs

To address the data silos and interoperability issues identified in the previous section, this research leverages Knowledge Graphs. The following subsections detail the structural advantages of graph-based data, the role of formal ontologies in standardising domain knowledge, and the current applications of these technologies in representing complex infrastructure assets.

2.2.1 Why graphs?

A graph is a structure represented by a collection of nodes (or vertices) that are connected by relationships (or edges). Nodes typically represent entities or concepts that are used in specific domains. Relationships define how these concepts interrelate. These relations are normally defined with a direction (directed graphs). For instance, using a node representing the concept *Bridge* and another representing *City*, we can establish a relation such as *located in, which* is directed from *Bridge* to *City*. This establishes a subject-predicate-object structure, called a triple, that conveys meaning in a way similar to human language. The predicate establishes a relation from the subject to the object. This enables the natural formation of rich conceptual networks that describe real-world elements and their contexts. Figure 2 illustrates a graph example partially representing a bridge. Consequently, graphs are simple yet powerful general-purpose tools to model complex systems semantically.

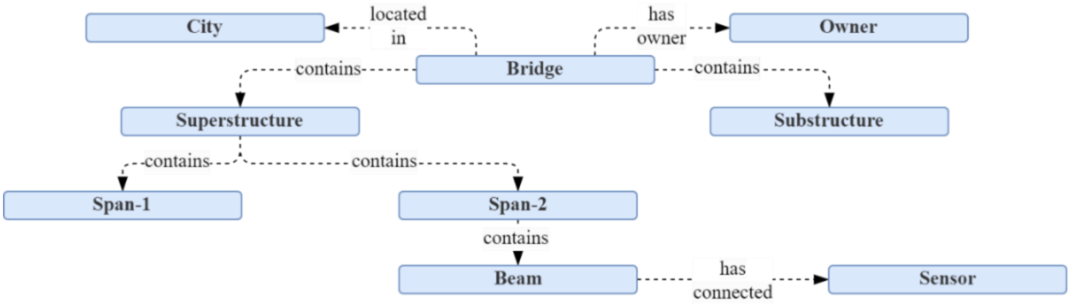

Figure 2: Example of a graph representing a bridge and its context.

Graphs can be encoded in a machine-readable form using the Resource Description Framework (RDF) (McBride, 2004), a standard developed by the W3C standardisation group for linked data on the Semantic Web (Berners-Lee et al., 2001). In RDF graphs, each element of a triple is uniquely identified by a Uniform Resource Identifier (URI), enabling data linking across the graph and the web. RDF data can be serialised using multiple formats, such as RDF/XML (https://www.w3.org/TR/rdf-syntax-grammar/), N-Triples (https://www.w3.org/TR/rdf12-n-triples/), Turtle (https://www.w3.org/TR/turtle/), or JSON-LD (https://json-ld.org/), and can be stored in specialised database systems referred to as semantic stores. Data within these stores is managed using SPARQL (https://www.w3.org/TR/sparql11-query/), a W3C-endorsed standard language and protocol for querying and manipulating RDF data.

2.2.2 Ontologies and knowledge graphs

Ontologies are defined in a machine-readable format, representing a shared conceptualisation of real-world entities or systems, to provide computers with the capability to process data that conforms to domain-specific rules, vocabulary, and context (Studer et al., 1998). Ontologies are developed using the Web Ontology Language (OWL), recommended by W3C and based on the RDF framework, making them directly compatible with RDF graphs (Hitzler et al., 2009).

Ontologies can be developed independently by different industry stakeholders, shared on the web, and connected to other existing ontologies to progressively build an interconnected network of concepts and their interrelations.

This modular approach fosters data interoperability and knowledge sharing among users. For this purpose, ontologies need to be developed and published in accordance with best practices to meet the Findable, Accessible, Interoperable and Reusable (FAIR) principles (Garijo and Poveda-Villalón, 2020). High-level tools such as Protégé can be used for ontology development (Jain and Singh, 2013). More low-level options, such as the RDFLib library for Python, allow the creation and management of ontologies and instance graphs through programming (see https://rdflib.readthedocs.io/). There are other useful tools available for ontology visualisation (Lohmann et al., 2015), and automated validation and documentation (Poveda-Villalón et al., 2014; Garijo, 2017).

Through ontologies, information contained in independent, yet complementary knowledge graphs can be structured and extended. By mapping graph data to entities and relationships in one or more ontologies, graphs can be constructed and queried from distinct yet overlapping domain-specific perspectives. Additionally, inference engines can utilise axioms and rules defined within ontologies to infer data not explicitly present in the graph, thereby expanding the semantics by which the graph can be consulted (Bienvenu et al., 2019). Figure 3 illustrates the concept of a knowledge graph built according to multiple ontologies.

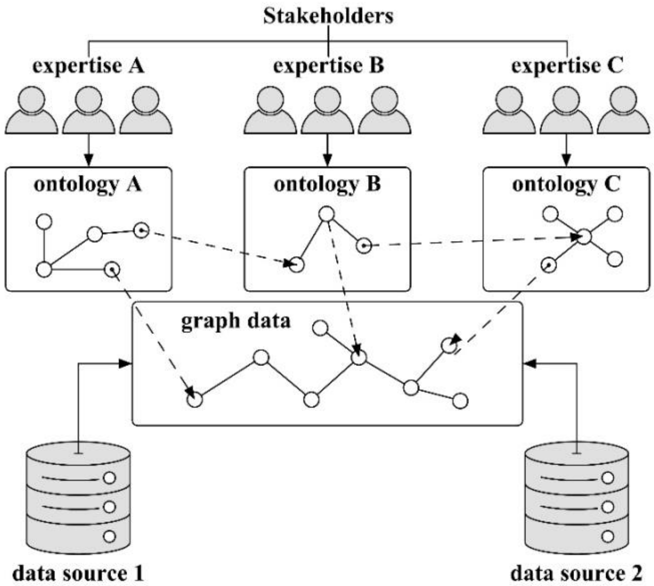

Figure 3: Knowledge graph concept.

2.2.3 Applications of knowledge graph-based systems in the built environment

The idea of using knowledge graphs and semantic web technologies to enable DTs in the built environment has been explored recently. Akroyd et al. (2021) presented this concept using a multi-domain knowledge graph dynamically operated by a fleet of distributed computational agents. These agents are focused on tasks that require querying, creating or updating specific subsets of graph data. Additionally, their distributed nature allows for the hosting of diverse information in distinct locations while providing links to their original repositories. By using specialised ontologies spanning multiple domains, data is interoperable across agents, enabling a continuous, seamless flow of information. Ontologies, data, and agents can therefore be added modularly as needed and composed to automate complex tasks.

Research efforts taking a knowledge graph-based approach have been conducted under the umbrella of The World Avatar Project (see https://theworldavatar.io) (Lim et al., 2022). The project promotes the use of a knowledge graph to create a DT of Singapore. The graph is dynamically accessed by agents that modify, update, process and visualise the information that it contains. Its applicability is demonstrated in large-scale building management (Quek et al., 2024), flood assessment (Hofmeister et al., 2024), and geospatial city management applications (Chadzynski et al., 2023). Another example is the CReDo project (Climate Resilience Demonstrator: see https://digitaltwinhub.co.uk/climate-resilience-demonstrator-credo), carried out as part of the digital twin program in the UK, where the knowledge graph combined the description of assets from energy, water, and telecommunication networks with data from flood simulation and models that describe the effect of the flood on the assets (Hayes et al., 2022; Akroyd et al., 2022).

This technological concept for a DT has the potential to meet the requirements and address the challenges described in Section 2.1. However, related research projects focus on knowledge graph-based DT implementations to support the operations of large-scale assets, such as entire cities or territories. Narrowing the scale to individual built assets, these types of systems are very limited in the academic literature (Ramonell et al., 2023). However, ongoing efforts to use semantic web technologies in the field of BIM have already established a set of backbone ontologies that enable the use of this approach at the building level as well (Pauwels et al., 2021). The developments in this paper are inspired by the knowledge graph-based architecture of The World Avatar project. The proposed DT system adopts the concept and focuses specifically on bridges, providing an ontological framework to model the bridge maintenance domain and providing tools that facilitate the integration of domain-specific information modelling and management standards.

2.2.4 Ontologies for representing built assets

One of the most complete ontologies enabling the digital representation of built assets is the IFC schema. Its OWL version, IfcOWL, was initially proposed as a way of using IFC as RDF graphs (Pauwels and Terkaj, 2016). The updated version of IfcOWL (see https://www.w3id.org/ifc/IFC4X3_ADD2) also contains classes for representing bridges included in the IFC4X3 ADD2 version of the EXPRESS schema. However, IfcOWL is a large and complex ontology that exhibits inefficiencies when using and querying data (Pauwels et al., 2017a; Terkaj and Pauwels, 2017). To achieve a simpler, extendable and modular approach, the W3C Linked Building Data (LBD) Community Group promotes the shift towards using Semantic Web technologies in the AECO industry, curating and aligning an ecosystem of lightweight, connectable ontologies to be used in various applications related to the life-cycle management of built assets (visit https://www.w3.org/groups/cg/lbd).

Central to this ecosystem is the Building Topology Ontology (BOT) (Rasmussen et al, 2021), which is inspired by the IFC core and designed to be extended with other ontologies from multiple domains. For instance, IFC-derived taxonomies such as the Built Element Ontology (BEO) and Distribution Element Ontology (MEP) are available to extend BOT with IFC products via OWL adaptations rooted in the *IfcBuiltElement* and *IfcDistributionElement* classes, respectively. Although these ontologies are focused on buildings, they can be used to represent other types of built assets. See https://www.w3id.org/bot for BOT, https://www.w3id.org/beo for BEO, and https://www.w3id.org/mep for MEP.

Similar studies have been presented, focusing on modelling bridges. Hamdan and Scherer (2020) presented a bridge ontology framework. Mirroring the role of BOT for buildings, the Bridge Topology Ontology (BROT) serves as the central framework (see https://w3id.org/brot). It provides the basic semantics for representing bridge structures, including components, zones, and their topological relationships. The ontology is complemented with the BRIDGE ontology for a general description of bridge construction types (see https://w3id.org/bridge), the BRCOMP ontology for a further classification of bridge components defined in BROT (see https://w3id.org/brcomp), the BRSTR ontology for defining the structural behaviour semantics of the bridge (see https://w3id.org/brstr), and the BMAT ontology for defining bridge-specific materials (see https://w3id.org/bmat).

The Damage Topology Ontology (DOT) (Hamdan et al, 2019) provides the base semantics for representing asset inspections and damage assessments. This ontology can be extended to detailed damage classification, damage assessment, mechanical degradation and other application scenarios (see https://w3id.org/dot). One example is the Concrete Damage Ontology (CDO), which adds semantics to model the types of damage that occur in concrete-made built components (available at https://github.com/Alhakam/dot/blob/master/Extension/cdo.ttl).

The entities defined in these ontologies can be semantically linked to geometric information through the Ontology for Managing Geometry (OMG) (see https://w3id.org/omg). This ontology is extended with the File Ontology for Geometry Formats (FOG) to include known geometry formats such as OBJ, E57, IFC, DWG, etc. (see https://w3id.org/fog), and with the Geometry Metadata Ontology (GOM) to add geometry-related metadata (see https://w3id.org/gom). Additionally, the Ontology for Property Management (OPM) (see https://w3id.org/opm) facilitates following the evolution of specific properties as they are updated over time (Rasmussen et al., 2018).

The combination of these ontologies forms the basis for interlinking information on building topology and metadata with descriptions of construction products, their state and identified damages, as well as their geometric descriptions, within a knowledge graph. That constitutes the main idea behind the notion of Linked Building Data (LBD), as presented by the W3C LBD Community Group (see https://www.w3.org/community/lbd/). This approach thus also represents the essential core model for the proposed Digital Twin system for bridge maintenance presented in this study.

2.3 Bridge maintenance

This section details the fundamental strategies for managing and preserving bridge infrastructure. It outlines the transition from reactive to proactive maintenance models and establishes that structural performance indicators are the primary metrics for evaluating asset health over time.

2.3.1 Maintenance approaches

Maintenance strategies for bridge networks are fundamentally based on monitoring structural performance indicators, which standardise the analysis of deterioration and facilitate data-informed decision-making. As reviewed by Sanchez-Silva et al. (2016), this structural degradation is driven by a combination of gradual factors such as material ageing, environmental conditions, and traffic loads, as well as abrupt shocks caused by events such as flooding or earthquakes. To counteract these effects and extend the structure's operational life, maintenance is categorised by the timing of intervention: corrective maintenance is implemented reactively when performance drops to a minimum acceptable threshold, whereas preventive maintenance adopts a proactive approach to ensure indicators remain above a specified higher service level (Figure 4).

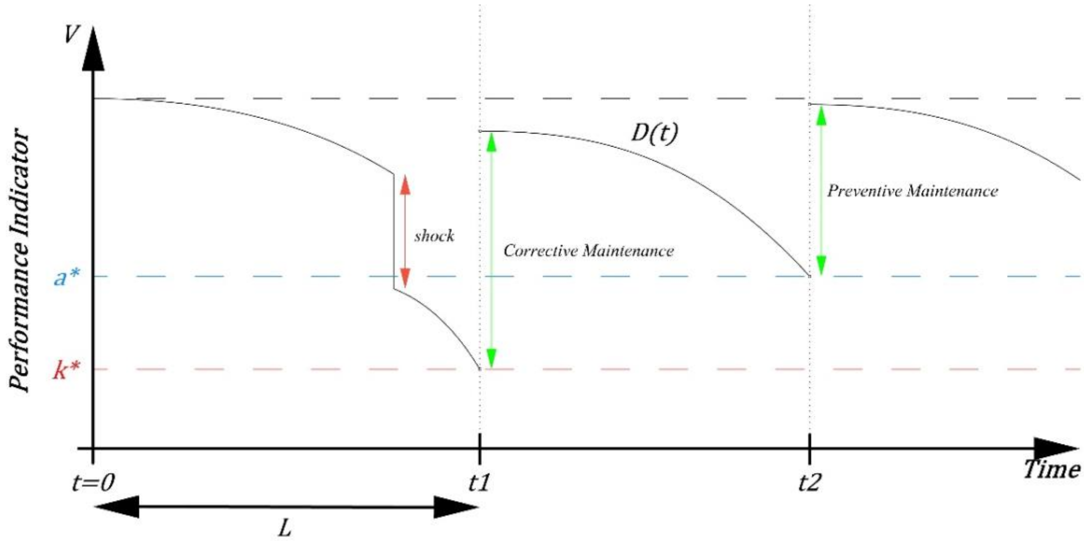

Figure 4: Parametrisation of bridge deterioration based on PIs and associated maintenance strategies, where V is the structural performance, D(t) is the deterioration function, k and a* are the corrective and preventive maintenance thresholds, respectively.*

In current practice, preventive maintenance strategies based on continuous structural monitoring are preferred for their long-term cost-effectiveness and ability to identify risks early (Sanchez-Silva et al., 2016; Kim and Yeo, 2016; Daro et al., 2022). These strategies typically rely on visual inspections to record damage, which are increasingly complemented by sophisticated assessments using sensors, imaging, and point clouds. Furthermore, techniques such as load testing and non-destructive evaluation are widely employed to establish accurate baselines for structural performance (Lee et al., 2014; Abudabous and Feroz, 2020; Chacón et al., 2023b).

A digital twin of a bridge, designed to support maintenance workflows, aligns closely with the principles of preventive maintenance. Therefore, these twins should be able to integrate at least the information collected during visual inspections. Integrating the insights generated from these activities into a Bridge Digital Twin could yield significant benefits:

- Historical Records: Established inspection programmes provide an extensive historical record of bridge conditions, enabling Digital Twins to leverage existing longitudinal datasets.

- Workflow Integration: Supporting these procedures within a Digital Twin framework ensures seamless alignment with current industrial maintenance workflows.
- Contextualisation: Contextualising inspection data alongside relevant datasets, such as BrIM and GIS models, enhances the understanding of bridge condition and improves the planning of future maintenance activities.

2.3.2 Bridge inspections

The most basic assessment form is a visual inspection. Visual inspections offer a versatile and cost-effective method for monitoring the visible components of a bridge. They can be carried out at different levels of detail, from individual components to entire bridges, in some cases without requiring highly specialised personnel, and in most cases utilising only simple equipment (e.g., cameras, flashlights, and an inspection sheet). A traditional concern is the costs associated with specialised equipment or operators when restricted access to specific bridge components is required (e.g., high altitudes, water bodies). In this regard, the use of UAVs equipped with high-resolution cameras is increasingly being implemented as a solution to the accessibility problem and to reduce time and equipment costs (Seo et al., 2018; Mandirola et al., 2022). Presently, visual inspections are widely adopted as the default method for monitoring bridge deterioration in many countries.

The results of these inspections are typically reported to bridge managers as a Bridge Condition Index (BCI). This parameter abstracts the structural condition of the bridge at a given point in time (See Figure 5 as a specification of Figure 4). The approach to calculating it varies from country to country. Iacovino et al. (2022) presented a thorough review of the methods that are used in different countries to calculate it, with references to the corresponding guidelines.

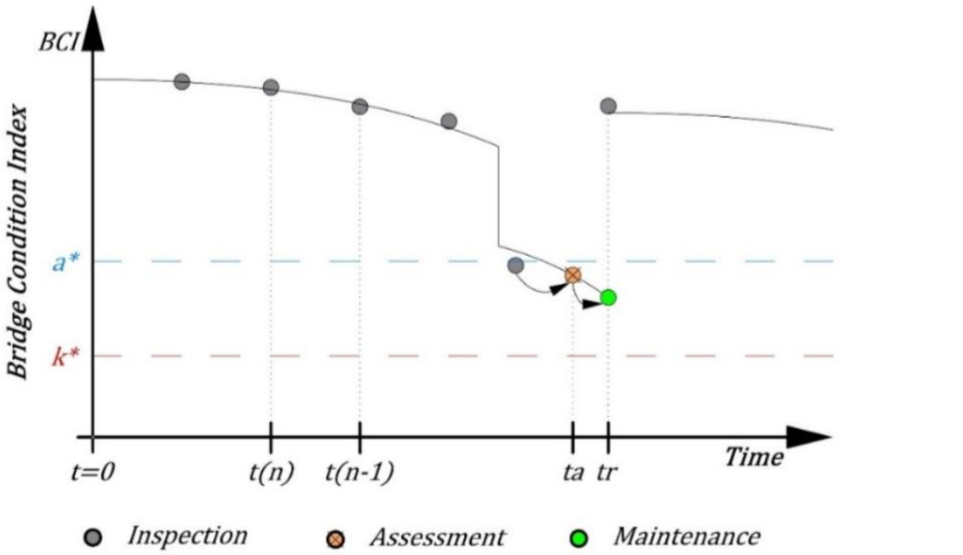

Figure 5: BCI-based bridge maintenance.

All methods are based on a qualitative approach. Inspectors annotate a set of parameters related to the damages observed on bridge components according to the requirements specified in the regional guidelines. Based on these annotated damages, the Condition Indices (CI) of individual bridge components or defects are derived, depending on different rules. To calculate the BCI, some countries use a weighted averaging method that assesses the overall structural condition by aggregating the CIs weighted according to their importance or the impact of the bridge component or defect on the bridge's structural integrity. Others take a worst-condition component approach, in which the BCI corresponds to the CI of the component in the worst condition. Others assign a rating associated with a qualitative value, e.g., from 1 (good) to 5 (bad), based on the inspector's judgment (Liu et al., 2023). In any case, BCIs serve as the quantitative basis for many models that predict the remaining life of the bridge structure (Srikanth and Arockiasamy, 2020). In this way, bridge owners can draw the deterioration curves of multiple bridges in a network, providing a basis for deciding whether more in-depth assessments are needed or to trigger maintenance actions.

3. DIGITAL TWIN SYSTEM DESIGN

This section details the architecture and technical framework of the proposed Bridge Digital Twin. It describes the integration of heterogeneous data sources, ranging from geospatial networks to specific bridge components, within a unified Knowledge Graph. The following subsections outline the system's conceptual logic, the specific ontologies used to define its semantics, and the functional components that manage data ingestion and delivery.

3.1 System concept

The system aims to enhance the observability of bridge networks, providing bridge administrators with all necessary information to track and understand the evolution of their bridges' condition. Unlike traditional Bridge Management Systems that operate as disjointed data silos, this system leverages a Knowledge Graph (KG) to integrate BIM, GIS, and inspection data standards, enabling dynamic analysis at both the component and network scales. The KG constitutes the semantic "twin" of a bridge network, interconnecting the "twins" of the bridges that form it by:

1. Linking data contained in different silos, forming a unified data network.
2. Classifying data according to multiple, overlapping ontologies, semantically enriching it and enabling interoperability.
3. Enabling context-rich queries that traverse from representations of individual bridge components to the scale of the whole bridge network.
4. Providing flexibility to the system, enabling it to evolve as new data is added or new ontologies are needed to describe the data.

The KG integrates bridge representations with the representation of the transport network they support (roads) and their broader geospatial context (e.g., railway lines and waterways). This structure is progressively enriched through the timely ingestion of inspection data, including structural damage records, imagery, and reports (see Figure 6), which enables the automatic aggregation of component-level condition updates into network-level indices (BCIs). Consequently, this study characterises the digital twin as a dynamic mirror of the network's current state, establishing a robust foundation for future applications for prescriptive maintenance prioritisation and condition forecasting based on historical BCI trends.

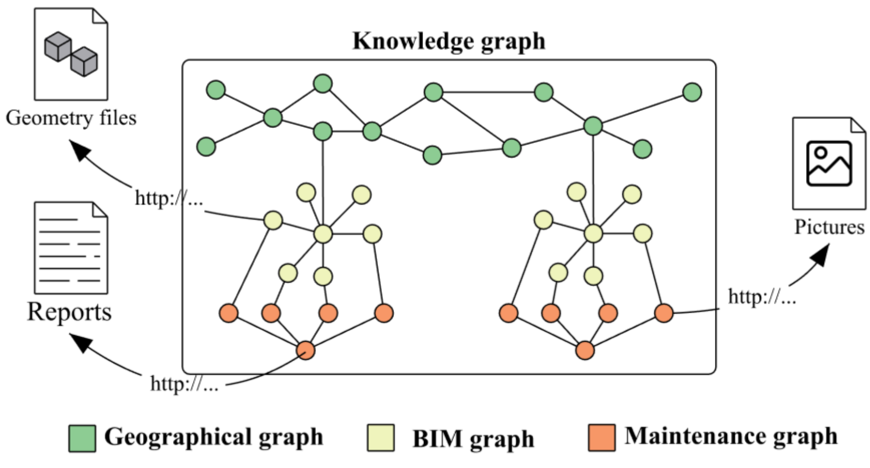

Figure 5: Knowledge Graph-based data integration concept for the proposed digital twin system.

The ontological framework used for the KG comprises the ontologies described in Section 2.2.4 and a set of newly developed ontologies for geospatial and inspection data integration. This framework is presented in detail in Section 3.2. Data integration and delivery are enabled through the orchestration of various functional components. These components are publicly accessible and are described in Section 3.3.

3.2 Ontologies

The ontologies used in the system facilitate structuring geospatial information, BIM, and inspection and damage-detection data as a graph. For each of these layers, multiple interconnected ontologies are used, as described in the subsequent sections.

3.2.1 Geospatial ontologies

Geospatial information is abundant and can be obtained from many sources. In this case, the INSPIRE Directive (Infrastructure for Spatial Information in the European Community) (Bartha and Kocsis, 2011) is considered the primary data source for representing the geospatial elements used throughout this study. Similar to IFC in BIM, INSPIRE provides a set of open GIS data models that cover 34 spatial data themes. These themes include, among others, administrative boundaries, hydrography, cadastral boundaries, risk zones, land use, buildings, and transportation networks, which are particularly relevant for asset management.

To represent geospatial information in accordance with the INSPIRE data models, a new set of OWL ontologies has been created to partially implement the classes and relations available in the INSPIRE Consolidated UML Models for hydrology and infrastructure networks (see https://inspire-mif.github.io/uml-models/approved/html/index.htm). The Network ontology (see www.w3id.org/inspireOnts/Network) defines a *Network* as a geographical feature composed of Network Elements, such as Nodes, Links, Areas, and various types of Connections to other *Networks*. Figure 7 shows this ontology pictorially. From the Network Ontology, the RoadNetworks (visit http://www.w3id.org/inspireOnts/RailTransportNetwork#), RailNetworks (visit http://www.w3id.org/inspireOnts/HydroNetwork#), and HydroNetworks (visit http://www.w3id.org/inspireOnts/RoadTransportNetwork#) ontologies are derived. These ontologies specify classes to define railway lines, roads, and rivers, respectively.

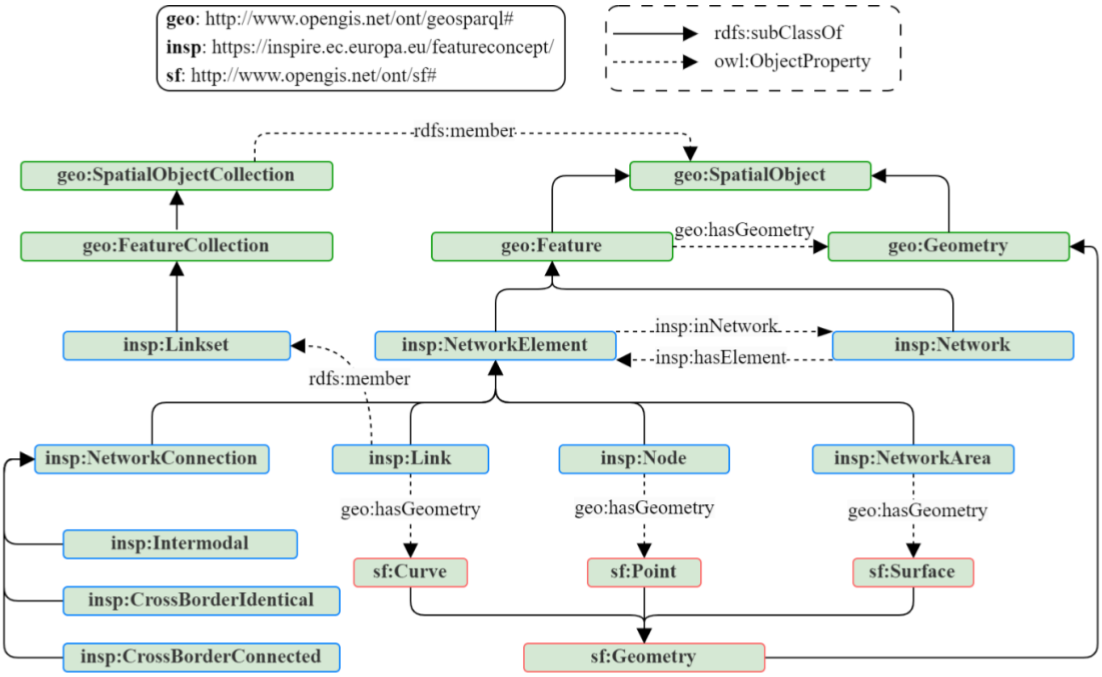

Figure 6: Network ontology.

These ontologies are a specification of existing standard ontologies for representing geospatial data: the GeoSPARQL Ontology (see https://opengeospatial.github.io/ogc-geosparql/geosparql11/geo.html) and the Simple Features Vocabulary (see https://opengeospatial.github.io/ogc-geosparql/geosparql11/sf_geometries.html). GeoSPARQL is a standard from the Open Geospatial Consortium (OGC) designed to represent and query geospatial linked data (Battle and Kolas, 2011). It provides the ontology and an online service for performing basic operations with spatial geometry. The Simple Features Vocabulary is a taxonomy that extends the GeoSPARQL ontology to represent two-dimensional geographic features, such as points, lines, and polygons.

Currently, these ontologies are used to semantically distinguish among road, hydrological, and railway links, serving the system's demonstration purposes.

3.2.2 BrIM ontologies

The ontologies described in this section represent bridges and their components as RDF graphs. All the ontologies already exist as part of an ontology suite curated by the LBD group (visit https://www.w3.org/community/lbd/). The Building Topology Ontology (https://w3id.org/bot) and Built Element Ontology (https://w3id.org/beo) serve as the basis for representing bridge spatial structures and their built components, respectively. The Built Element Ontology has been updated in the frame of this study to incorporate the new classes and class properties described in the IFC4X3 Addendum 2 schema. The ontology is automatically generated from the EXPRESS schema definition using the open-source IFCEXPRESS2BEO converter (see https://github.com/cramonell/beo). Figure 8 shows a subset of the classes in this ontology that are useful for modelling bridge components.

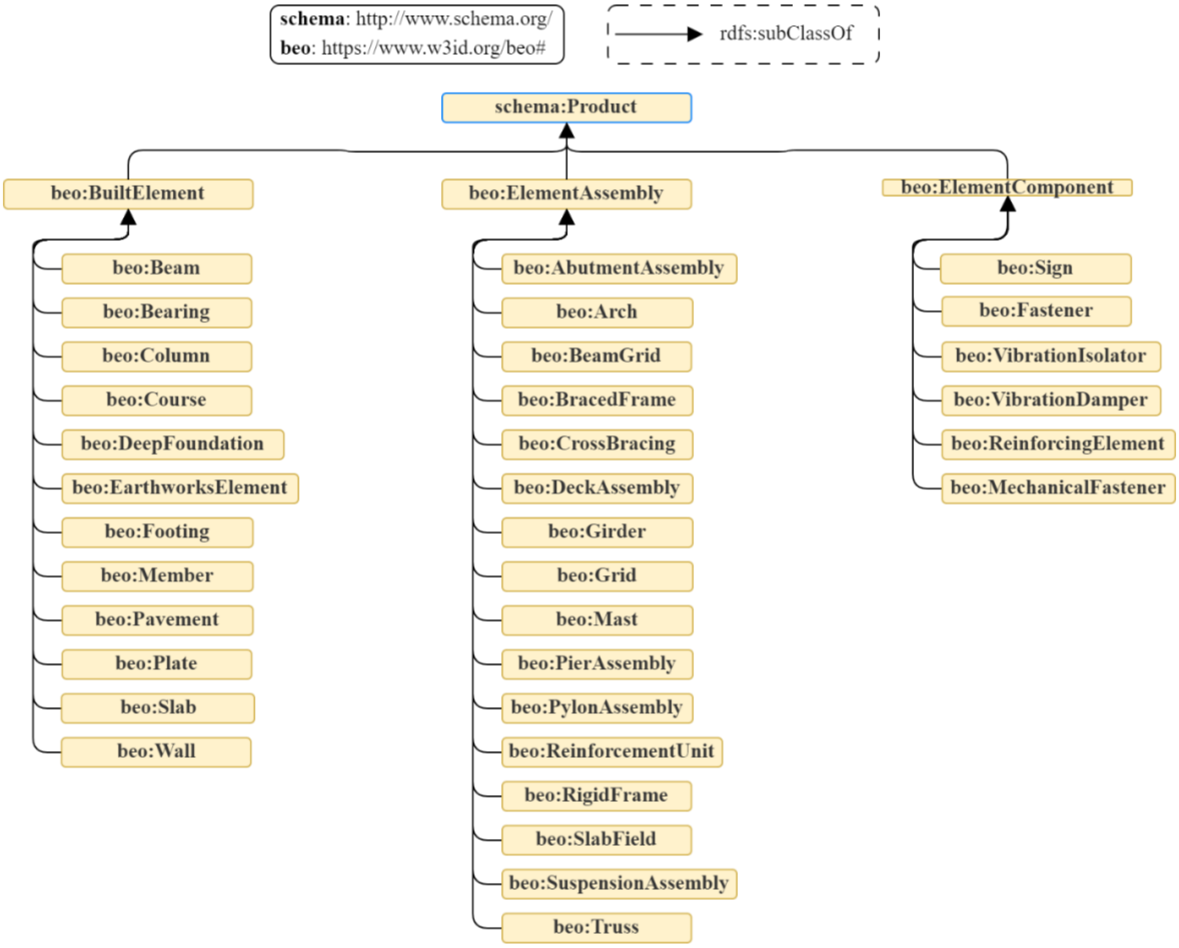

Figure 7: A subset of classes in the Built Element Ontology (BEO) that enable the representation of bridge components.

Additionally, the Bridge Ontology (https://w3id.org/bridge) and the Bridge Topology Ontology (https://w3id.org/brot) are used to add semantics for bridges and bridge types. To include geometric semantics in the graph and link graph elements to external files containing their geometric descriptions, the Ontology for Managing Geometry (https://w3id.org/omg), the Geometry Metadata Ontology (https://w3id.org/gom), and the File Ontology for Geometry Formats (https://w3id.org/fog) are combined.

3.2.3 Ontologies for modelling bridge inspections

Inspection procedures are governed by guidelines that vary significantly across jurisdictions. Consequently, defining a single, universally generalizable ontology is unfeasible. Instead, it is the responsibility of each relevant authority to define ontologies that align with the specific information requirements of their regulatory context. To represent inspection results in this case study, the Ontology for Visual Inspection of Bridges in the Spanish Road Network has been developed (see https://cramonell.github.io/bridge/inspection/iopce/actual/index-en.html). This ontology adheres to the information requirements of the Spanish national guidelines, which mandate that every bridge component be inspected to identify and characterise damage based on *severity, extent, and evolution* (see Table 1). Based on these parameters, a Damage Index is derived for each component and then aggregated to calculate the global Condition Index of the structure (see Table 2).

Table 1: Damage description parameters according to the Spanish guidelines for bridge inspection.

Index	Severity	Extent	Evolution
0	Low probability of causing loss of structural function	-	Slow speed and non-aggressive environment
1	Medium probability of causing loss of structural function	Lower than 25%	Low speed and aggressive environment, or high speed and non-aggressive environment
2	Damage with a high probability of causing loss of structural function	Between 25% and 50%	Fast speed and aggressive environment
3	-	Between 50% and 75%	-
4	-	Higher than 75%	-

As the specific algorithms for these calculations are proprietary and not publicly available, these parameters were approximated using heuristic functions informed by engineering judgment. This approach ensures the resulting indices are consistent with the guideline's logic without requiring access to the restricted official formulations, which does not impact the validation of the system's semantic architecture.

Table 2: Description of the Damage Index value ranges according to the Spanish guidelines for bridge inspection.

Index	Damage Index	Condition Index
0-20	Damage with no significant consequences ’a priori’	Structure without evident pathologies or with damages that have no relevant consequences for the durability or service conditions of the structure.
21-40	Damage that may have a pathological progression or reduce the serviceability or durability of the component if not repaired promptly.	Structure with damages that may have a pathological evolution affecting the durability or service conditions of the structure. It is advisable to monitor its temporal evolution to determine it objectively.
41-60	Damage indicating a pathology that results in a reduction of serviceability or durability of the component.	Structure with damages that indicate a pathology leading to a reduction in the service conditions of the structure. It will be necessary to track the evolution of the pathology in subsequent inspections. Medium-term action may be required to enhance the structure’s durability.
61-80	Damage that can translate into a modification of the structural or functional behaviour.	Structure with damages or pathologies that can result in a modification of the load-bearing behaviour or a significant reduction in service levels. Requires short-to-medium-term action. Depending on the nature of the damage, a special inspection may be necessary.
81-100	Deterioration that compromises the component's safety.	Structure with damages or pathologies compromising the safety of the component/structure. Requires a special inspection and urgent intervention. In some cases, usage restrictions may be necessary.

The ontology is built upon the semantics provided by the Damage Topology Ontology (see https://w3id.org/dot) and the Ontology for Property Management (https://w3id.org/opm). DOT is utilised to model identified damage and its topological relationship with affected bridge components or zones (representative spatial volumes that contain multiple bridge components or other zones). Meanwhile, OPM enables the modelling of properties that evolve over time; this temporal tracking is critical for recording the progression of structural damage and fluctuating Condition Indices. These indicators are updated within the graph following the methodologies established by Rasmussen et al. (2018). The proposed ontology is illustrated in Figure 9.

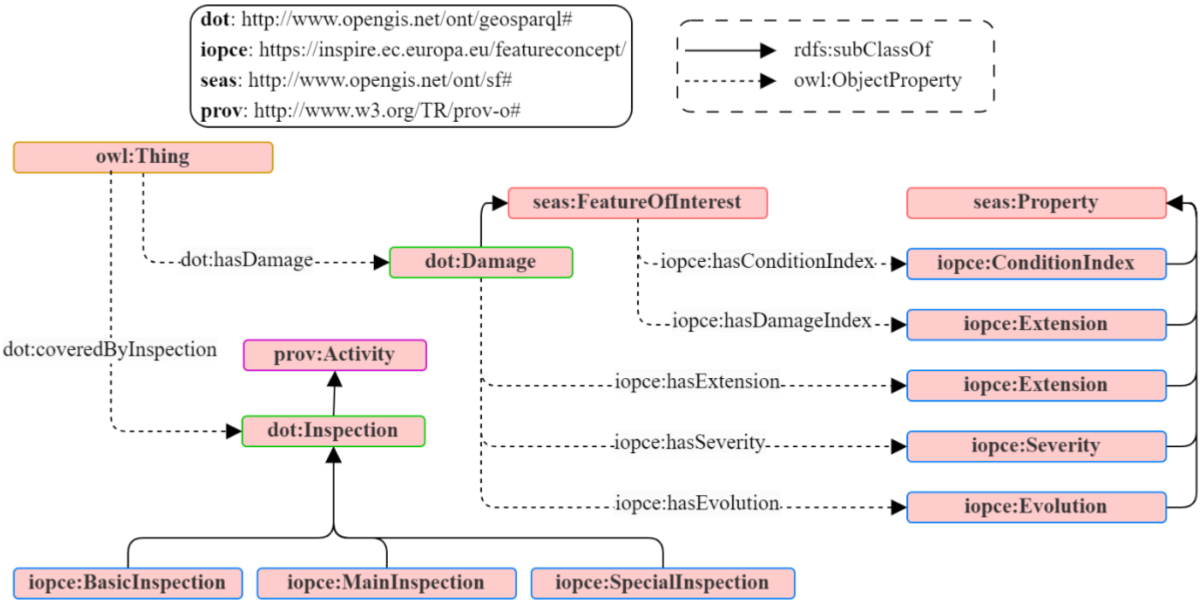

Figure 8: Ontology for bridge inspection conforming to the Spanish guidelines.

3.3 System components

The system architecture is based on a microservices framework comprising Extract, Transform, and Load (ETL) tools and a hybrid database structure. The ETL components facilitate the conversion of heterogeneous external data sources (including geospatial, BIM, and inspection-related data) into knowledge graphs. Ontotext GraphDB serves as the primary triple store for the knowledge graph, while a dedicated file database stores and links auxiliary documents (e.g., imagery and PDF reports), according to regular LBD recommendations and best practices. These components are briefly described in Table 3.

Table 3: System components.

Name	Type	Source	Description
Graph DB	Database	GraphDB (DockerHub)	A SPARQL-compliant triple store for storing the system's RDF data.
Files DB	Database	Authors (Local)	A local file system for storing binary data (geometry files, images, reports).
IFC2KG	ETL	Authors	Manages the conversion of IFC semantic and geometric data into RDF. See Section 3.3.1.
INSPIRE2KG	ETL	Authors	Converts INSPIRE-compliant GIS shapefiles into RDF. See Section 3.3.2.
ICDD	ETL	Authors	ETL that manages the ingestion of ICDD containers within the system databases. See Section 3.3.3.

The interaction between external applications and the system is facilitated by a Web API developed using Python and FastAPI (see https://fastapi.tiangolo.com). The API provides several endpoints to query the knowledge graph and trigger ETL processes, as summarised in Table 4.

Table 4: System API endpoints.

API Endpoint	Description
/sparql	An endpoint to interact with the knowledge graph through SPARQL queries.
/files	An endpoint that enables the retrieval of specific files from the file database.
/ifc	An endpoint that activates the functionalities of the IFC2KG ETL.
/inspire	An endpoint that activates the functionalities of the INSPIRE2KG ETL.
/icdd	An endpoint that activates the functionalities of the ICDD ETL.

In this study, three different types of connected applications are envisioned for the demonstration of the proposed system:

1. Inventory Application: Allows users to deliver the digital representation of a new bridge to the system's bridge inventory.
2. Inspector Application: Enables inspectors to interact with previous inspection data and deliver new inspection results.
3. Bridge Manager Application: Visualises and queries information about the network state, showing the BCI as the main performance indicator.

The development of these applications, as well as their corresponding interfaces, falls outside the scope of this study. However, their interaction with the system is described in detail in Section 4, which covers the system validation in a case study. Figure 10 schematizes the system components, the connected applications, and the information flow.

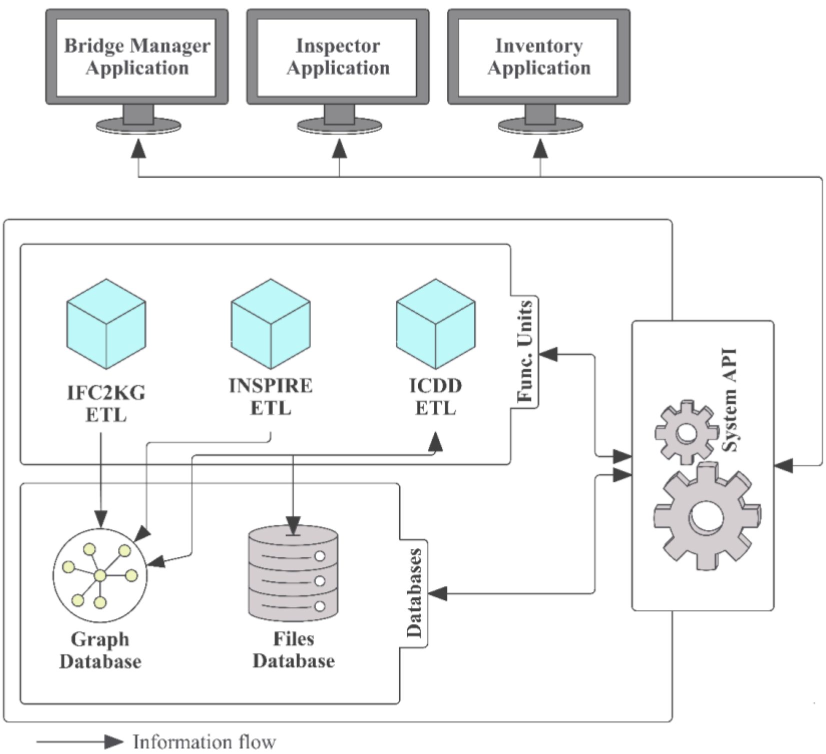

Figure 9: System architecture and connected applications.

3.3.1 IFC2KG ETL

The IFC2KG ETL has been developed as part of this study. It enables the import of IFC files containing bridge representations as RDF graphs. The ETL is based on a tool with the same name, available in the following repository: https://github.com/cramonell/KG4BE. The tool is based on IfcOpenshell and RDFLib. To perform the conversion, it uses a JSON file to map classes, attributes, and relations in the IFC schema with entities defined in multiple ontologies. The idea is to use this mapping file to configure data conversion and, therefore, adapt it to the needs of different users. The conversion tool produces an RDF file. The mapping file can be found in this repository: https://github.com/cramonell/KG4BE/blob/main/IFC-converter/conversion-map.json. The JSON file and the test files are available in the following repository: https://github.com/cramonell/KG4BE/blob/main/IFC-converter/.

3.3.2 INSPIRE2KG ETL

The INSPIRE2KG ETL is designed for importing shapefiles, defined according to the INSPIRE data models, as RDF graphs. It includes three conversion tools for hydrology, rail, and road networks, based on a customised version of code that provides a library for converting shapefiles to RDF files in accordance with GeoSPARQL. The code is available in the following repository: https://github.com/nvitucci/shape2geosparql/tree/master/shape2geosparql. Currently, the converters are configured to process data obtained from the INSPIRE geoportal provided by official Spanish agencies (visit https://centrodedescargas.cnig.es/CentroDescargas). This data aligns with entities in INSPIRE data models and also includes additional properties to describe geospatial features according to National counterparts. Consequently, the converters focus solely on entity conversion, not properties, and are not generalizable to other datasets. The code is available for inspection in the following repository: https://github.com/cramonell/KG4BE/tree/main/INSPIRE-converter/. It is important to note that this tool is not expected to be used dynamically. The geospatial information serves as the semantic basis for the graph that interconnects the rest of the data related to bridges and bridge inspections, and is loaded only once during system setup. Various test files to inspect the conversion results are available in the following repository: https://github.com/cramonell/KG4BE/tree/main/INSPIRE-converter/tests.

3.3.3 ICDD ETL

The ICDD standard enables the exchange of documents linked to semantic information contained in RDF graphs. However, the structure of an ICDD container is not directly compatible with the proposed architecture, which uses a single graph linked to documents. Thus, the ICDD ETL is designed to ingest data from ICDD containers into the system, enabling inspectors to unify the delivery of semantic descriptions of inspections and damage, linked to supporting documents such as photographs and technical reports. The ETL utilises a custom Python library which facilitates the creation and modification of ICDD containers. It was developed during this research and is publicly available in the following repository: https://github.com/cramonell/ICDD.

The ETL processes an ICDD container as input, unbundles its contents to the file database, and generates a new RDF file for the graph database. This RDF file contains updated URI links to the new persistent storage locations. The construction of the RDF file follows a three-step logic:

1. All documents within the Payload Folder that contain RDF serialisations are imported into the new RDF graph. Then, linkset files in the Payload Triples Folder are processed, and *Linkset:Links* are also added to the graph.

2. *For Linkset:LinkElements* pointing to *container:Documents* with a *container: filetype* property that is not an RDF format (.ttl, .nt, .rdf, etc.), the ETL adds their *linkset:Identifier* and the referenced *container:Document* in the index file. The document URLs are redirected to their new location in the system file database.

3. If the *container:filetype* property is an RDF serialisation format, the *linkset:URIBasedIdentifier* is used to instantiate it in the graph with such identifier as a *linkset:LinkElement*, establishing a direct link between a node in the graph and the document.

Once this process is finished, the resulting RDF file is imported into the graph database. Figure 11 shows the results of the data transformation between the ICDD container and the system databases.

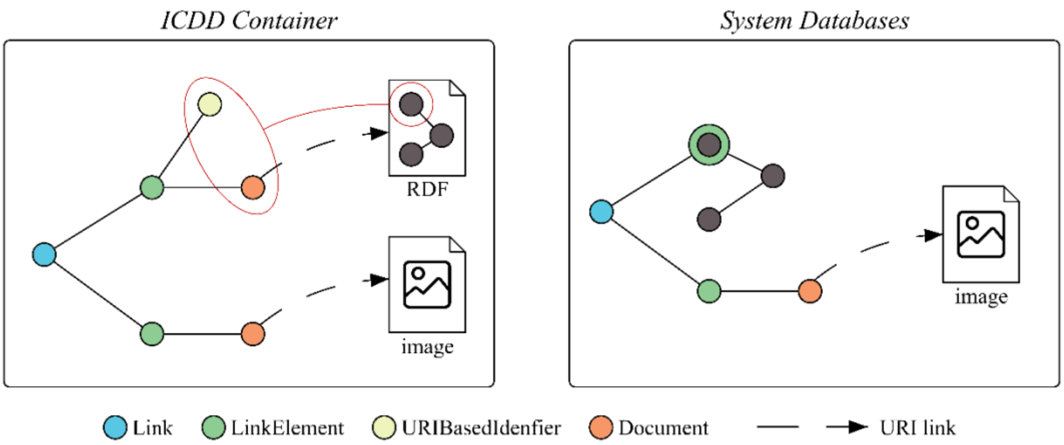

Figure 10: Difference between the structure of link data inside an ICDD container and the system databases.

4. SYSTEM IMPLEMENTATION: CASE STUDY

The system is demonstrated through a case study of a bridge network in the Metropolitan Area of Barcelona, Spain. The objective is to demonstrate how the knowledge graph accommodates heterogeneous bridge representations, interlinks them with their geographical context, and dynamically updates the network state. The case study highlights the system's ability to span multiple scales, from individual bridge components to the entire infrastructure network, supported by a continuous influx of inspection data.

At the network scale, a segment of the A-2 highway containing 11 bridges is considered, interlinked with adjacent road, railway, and hydrological networks. At the bridge and component scale, the PR-04-B015 bridge is used as a primary example. This 846-metre-long structure is a continuous, horizontally curved composite bridge spanning the Llobregat River. Figure 12 illustrates the bridge's structural details, while Figure 13 provides its geographical context within the A-2 highway network.

Figure 11: PR-04-B015 bridge, passing over the Llobregat river.

It is assumed that the system is operated by regional bridge administrators and pre-populated with geospatial data representing the A-2 highway and its intersecting networks. These data were ingested via the INSPIRE2KG ETL (Section 3.3.2) using public shapefiles from official geoportals. In this scenario, the inventory initially contains all bridges within the highway segment except for PR-04-B015. The case study then demonstrates the enrichment of the system's data through three sequential steps:

Step 1: Asset Information Delivery. From the perspective of bridge managers using the Inventory Application, the as-built representation of the PR-04-B015 bridge is delivered to the system. This step illustrates the integration of the bridge model with existing geospatial data via the knowledge graph (Section 4.1).

Step 2: Inspection Delivery. From the perspective of bridge inspectors using the Inspector Application, inspection outcomes for PR-04-B015 are submitted. This step details the interaction between the application and the system using SPARQL queries and ICDD containers, showing how inspection data are linked to the structural model (Section 4.2).

Step 3: Network Management. From the perspective of bridge managers using the Bridge Manager Application, network-wide BCIs are recalculated, and the state history is queried. This step showcases the use of SPARQL to assess network conditions and prioritise maintenance actions (Section 4.3).

It is important to note that the presented inspection results and calculated indices are intended strictly for demonstrative purposes. They serve to illustrate the system's functionality and do not constitute an exhaustive or structurally certified assessment of the bridge and network's actual condition.

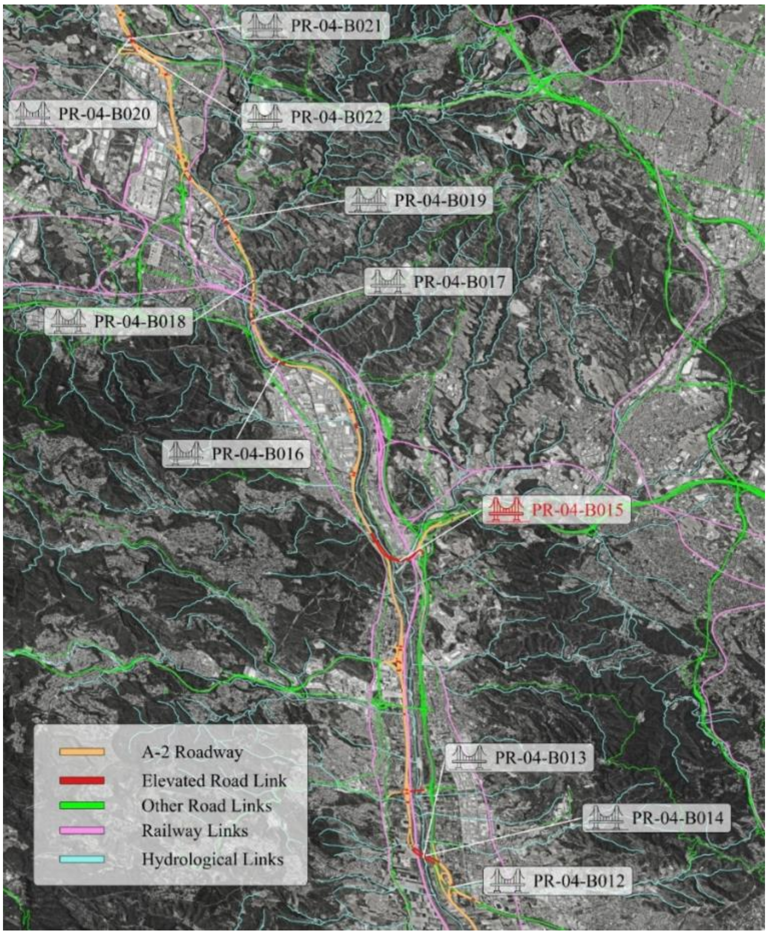

Figure 12: Representation of the geospatial information and the bridges used to test the system. This includes the A-2 roadway segment, additional roads, the railway network and the hydrological network of the region. The demonstrator bridge is the PR-04-B015 (marked in red).

4.1 Step 1: Asset information delivery

The ingestion of a new bridge into the digital twin inventory occurs after the construction phase, when as-built data are transferred, or after the retrospective digitisation of an existing asset. It is assumed that the delivery format is an IFC file. The PR-04-B015 model used here was developed within the H2020 ASHVIN project (visit https://www.ashvin.eu/). It is shown in Figure 14 and available in the following repository: https://github.com/cramonell/KG4BE/blob/main/IFC-converter/tests/IFC/llobregat-bridge.ifc.

Bridge managers upload the IFC file via the Inventory Application. The system API directs this to the IFC2KG ETL, which bifurcates the data: semantic triples are stored in the Graph DB, while geometric representations are stored in the File DB. Figure 15 shows a snippet of the converted IFC file, describing one of the bridge's internal frames. The full RDF representation of the bridge is available at https://github.com/cramonell/KG4BE/blob/main/IFC-converter/tests/RDF/llobregat-bridge.ttl.

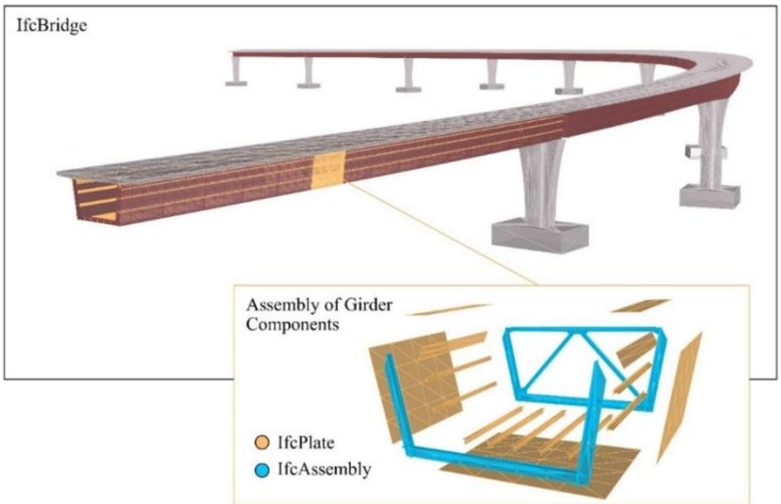

Figure 13: Geometric representation of the IFC model of the PR-04-B015 bridge and detail of the assembly of multiple components to model the bridge girder.

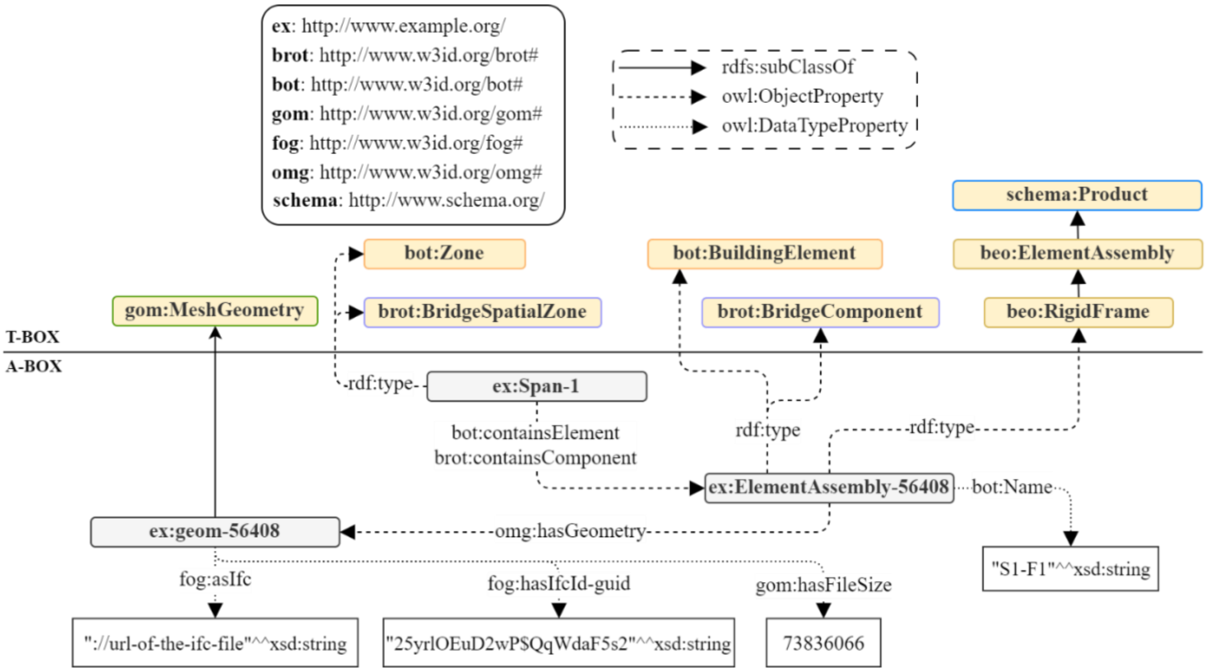

Figure 14: Steel frame of the PR-04-B015 bridge represented as a knowledge graph using the BOT, BROT, BEO and GOM ontologies.

Initially, the bridge and its geographical context coexist in the database as disconnected subgraphs. These can also be stored in multiple databases and multiple servers, in line with the philosophy of linked data and the semantic web (scalability). The location of the bridge, together with its geographical context, is represented in Figure 16. The geospatial graph uses *inspire:HydroLink*, *inspireF:RailwayLink*, and *inspireF:RoadLink* (subclasses of *geo:Feature*) with geometries encoded as *sf:LineString* in GeoJSON (see https://geojson.org/). Road and hydrological links are members of *inspireF:LinkSet*, which groups links into broader entities (in this case, the A2 highway and the Llobregat river). Figure 17 illustrates an example of a geospatial graph.

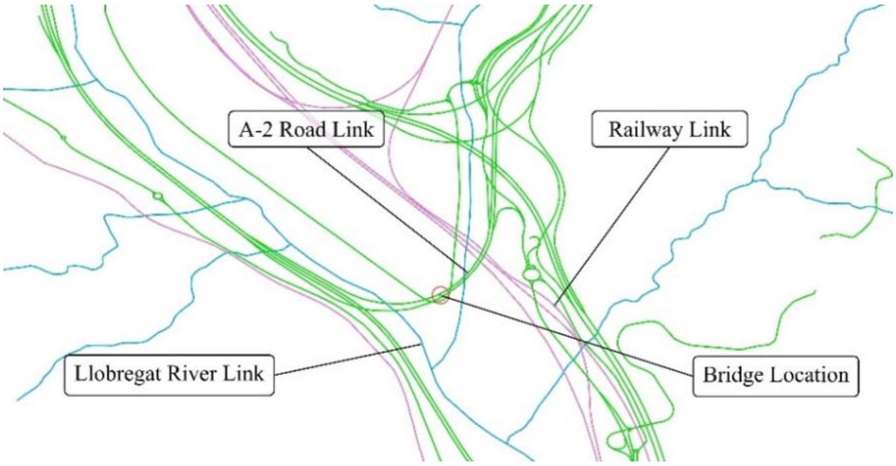

Figure 15: Geometrical representation of the bridge's geographical context.

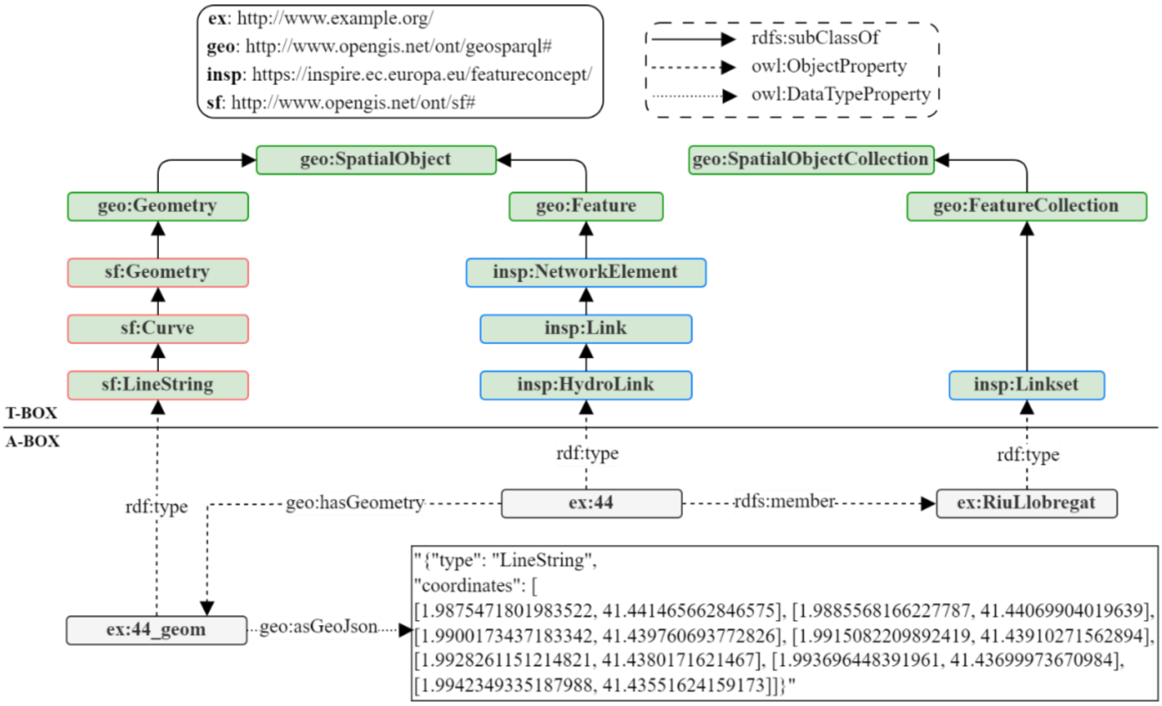

Figure 16: Knowledge graph representing a Hydrological link, member of the Llobregat river.

Currently, users need to manually link the bridge node with the surrounding geospatial element nodes. For this, a graphical user interface (GUI) is needed to trigger a SPARQL query that establishes *brot:locatedAbove* or *brot:locatedBelow* relationships. Also, this use of regular GUIs is in line with common guidelines in the use of linked data and semantic web technologies. These relations represent the relative spatial position of bridge components and can also apply to the bridge and its surrounding geospatial elements. The bridge is located below the roads that it supports and above all of the geospatial elements it spans.

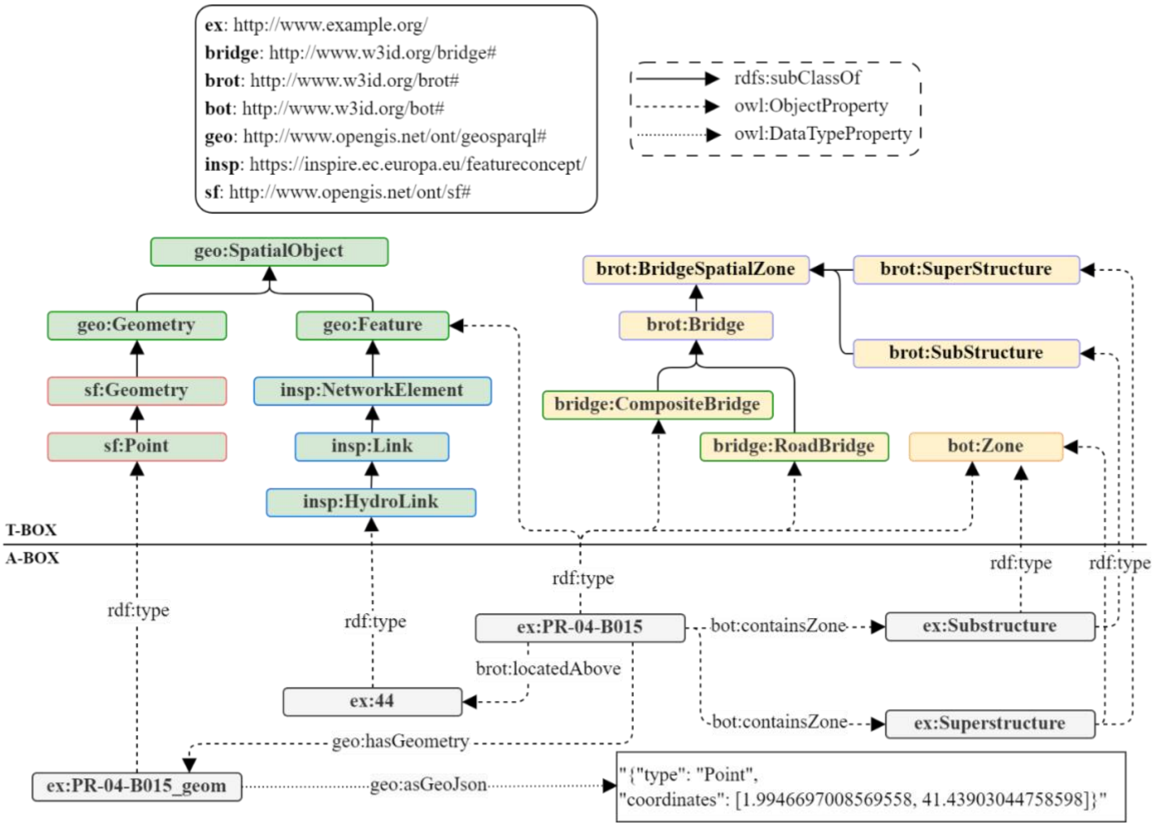

Figure 17: Knowledge graph integrating the bridge with the geographical representation of the road it supports.

4.2 Step 2: Inspection delivery

This stage demonstrates the delivery of results from a visual inspection conducted on the PR-04-B015 bridge. The Inspector Application enables personnel to submit all data collected during site visits in accordance with Spanish national guidelines. This includes the creation of a new inspection instance, the association and updating of damage records linked to specific bridge components, and the delivery of supporting documentation. The semantic integration of this information is realised through the user's interaction with the application's GUI. Figure 19 illustrates this process.

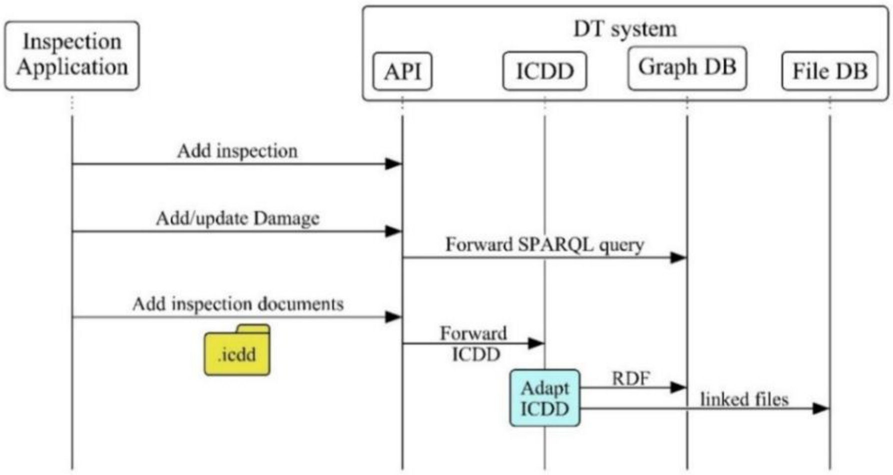

Figure 18: Interaction between the inspection application and the DT system to add new inspection data.

The SPARQL queries are stored in different files in an online repository. The links to the files are available in Table 5. The Add-inspection query records essential metadata, such as the unique inspection code, start time, and description, while establishing relationships with the inspected structural components. To record new damage and its associated indicators, the Add-damage query is employed. Each damage indicator is modelled as a *seas:Property* with an *opm:PropertyState* containing timestamped values. The *opm:CurrentPropertyState* class is used to annotate the most recent indicator value semantically. When damage identified in a previous inspection is reassessed, the indicators are updated using the Update-damage query. In this process, the previous *opm:CurrentPropertyState* is reclassified as a historical *opm:PropertyState*, and the new data is instantiated as the new *opm:CurrentPropertyState*. This approach, based on the methodology proposed by Rasmussen et al. (2018), ensures that the graph maintains a full longitudinal record of damage evolution, as shown in Figure 20.

Table 5: Summary of the SPARQL queries used by the Inspector Application to interact with the digital twin system.

Query ID	File link
Add-inspection	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/INS-add-inspection.sparql
Add-damage	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/INS-add-inspection-damage.sparql
Update-damage	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/INS-update-inspection-damage.sparql

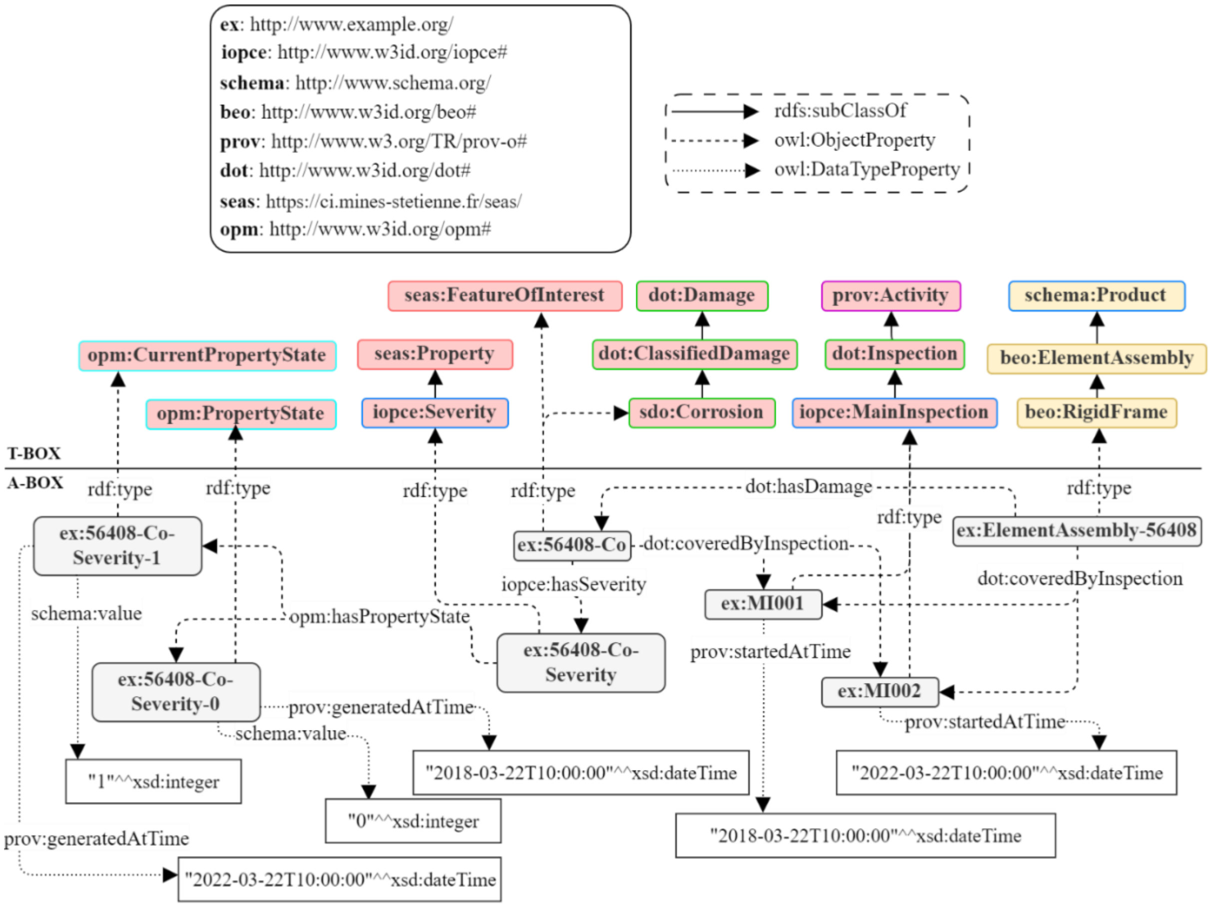

Figure 19: Knowledge graph representing corrosion severity evolution after two different inspections.

Additionally, the inspector may upload one or multiple files associated with the inspection data. These files are initially stored locally on the workstation where the Inspector Application is operating. To facilitate the transfer, these assets are bundled into an ICDD container and transmitted to the system API. Upon arrival, the container is deconstructed, and its contents are integrated into the system's databases. Figure 21 illustrates the internal ICDD structure employed to transfer images associated with specific structural damage.

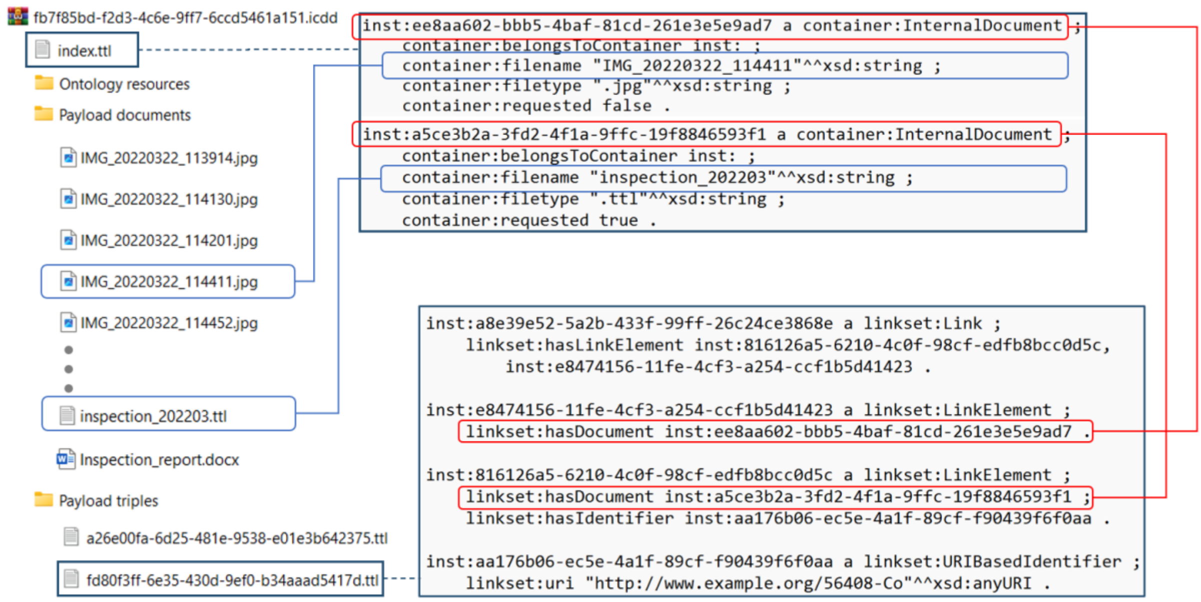

Figure 20: Exemplary ICDD container for delivering inspection results.

Figure 22 illustrates the integration of an image delivered via the ICDD container into the knowledge graph, where it is semantically linked to the specific damage instance that it documents.

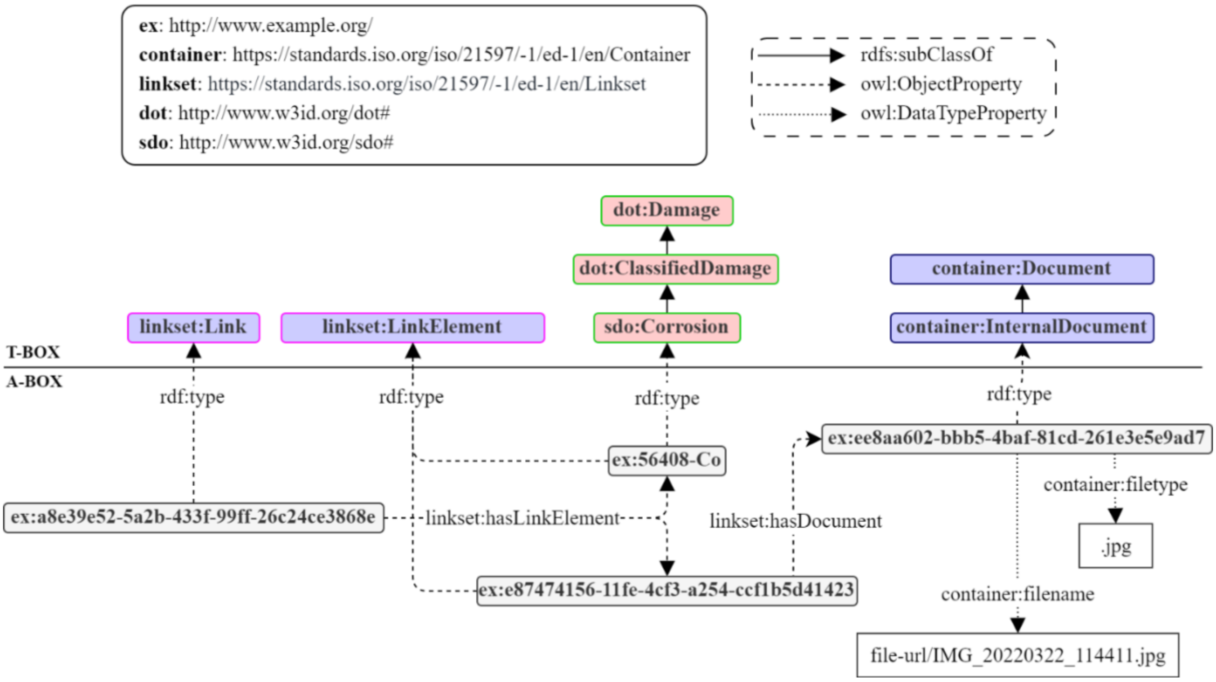

Figure 21: Exemplary integration of files (pictures) in the knowledge graph (damage representation).

4.3 Step 3: Network management

The delivery of new inspection data and its integration with the system's bridge representations occur concurrently for multiple assets across the network. As the system accumulates this information, it facilitates the inference of current bridge conditions, the reconstruction of historical deterioration patterns, and comparative benchmarking across the bridge portfolio. A key step for this is abstracting to a BCI the information delivered during inspections, which refers only to specific bridge components. This scenario demonstrates how the system can automatically abstract the BCIs of the network and give bridge managers a holistic view of the state of the network through the Bridge Manager Application. Figure 23 illustrates the BCI calculation pipeline and the interaction between the Bridge Manager Application and the Digital Twin system. Table 6 summarises the SPARQL queries that are available in an online repository.

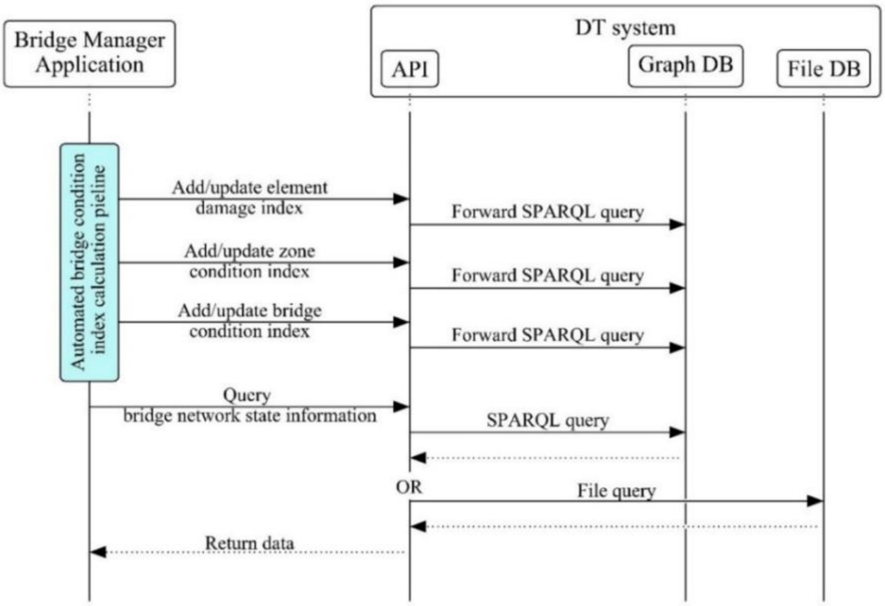

Figure 22: Interaction between the Bridge Manager Application and the DT system.

Table 6: Summary of the SPARQL queries used by the Bridge Manager Application to interact with the digital twin system.

Query ID	File link
Add-component-CI	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-add-element-CI.sparql
Update-component-CI	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-update-element-CI.sparql
Add-zone-CI	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-add-zone-CI.sparql
Update-zone-CI	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-update-zone-CI.sparql
Add-BCI	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-add-BCI.sparql
Update-BCI	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-update-BCI.sparql
BCI-History	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-BCI-history.sparql
BCI-A2	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-BCI-A2.sparql
BCI-River	https://github.com/cramonell/KG4BE/blob/main/paper-KG4BMS/SPARQL/BM-BCI-Llobregat.sparql

The BCI calculation pipeline executes a sequence of SPARQL queries that evaluate the structural state across three distinct semantic levels: the bridge component, the bridge zone, and the entire bridge assembly.

1. Component Condition Index calculation: The damage index for each bridge component is derived by aggregating the current state values of damage indicators identified during the most recent inspection. The Add-component-CI query instantiates new indices, while Update-component-CI modifies existing ones.

2. Zone Condition Index calculation: Each bridge can be semantically divided into multiple zones, understood as areas of interest for bridge maintainers to have a medium-scale abstraction of where the damage on a bridge might be located. For instance, each span and pier of a bridge could be considered an independent bridge zone. The condition index of a zone is calculated by summing all the component condition indexes calculated in the previous step. The Add-zone-CI query adds a new zone condition index, while Update-zone-CI modifies existing ones.
3. Bridge Condition Index (BCI) calculation: The final BCI is determined by taking the highest condition index among the bridge zones. The Add-BCI query adds a new bridge condition index, while Update-BCI modifies existing ones.

The persistent storage of BCI history within the graph database allows for longitudinal performance tracking. The BCI-History query demonstrates the retrieval of these temporal data, which are visualised in Figure 24 to reveal deterioration trends over successive inspection cycles.

Figure 23: Bridge condition index evolution over time.

Once the bridge-specific data are merged with the broader portfolio, the Bridge Manager Application can perform complex queries that leverage geospatial semantics. The system can filter asset health based on infrastructural or environmental contexts. For instance, the BCI-A2 query in Table 6 shows a method to get the current condition indexes of the bridges located on the A2 highway. As another example, the BCI-River query shows a method to get the current condition indexes of the bridges crossing the Llobregat River. Figure 25 displays the results of the latter.

The analytical capabilities enabled by the system's knowledge graph extend significantly beyond the static observation of Bridge Condition Indices (BCIs). By providing a multi-scale perspective, the framework effectively bridges the gap between granular, component-level structural data and high-level network analytics. The semantic interlinking of heterogeneous data sources allows for queries that traverse varying levels of abstraction, offering a holistic understanding of infrastructure health. This capability empowers bridge administrators to perform bespoke analyses tailored to specific operational requirements. Consequently, the transition toward such a semantically enriched Digital Twin ensures that infrastructure management becomes a fundamentally data-informed process, enhancing the long-term resilience and maintenance of built assets.

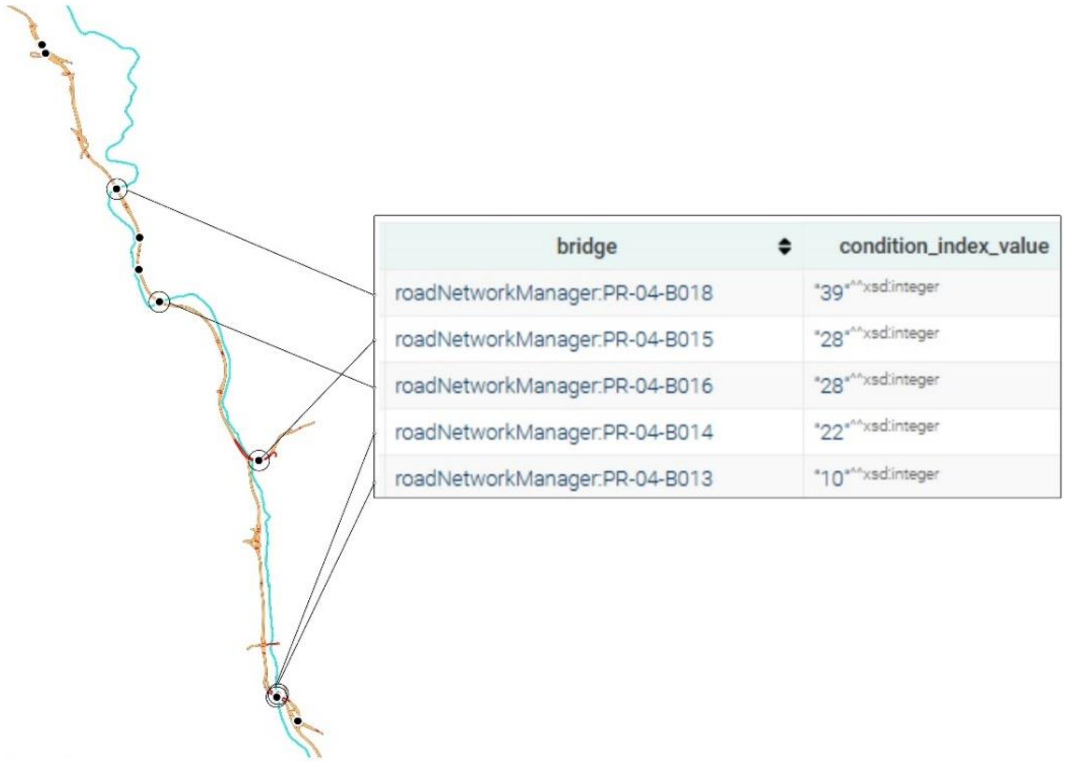

Figure 24: Ordered comparison of the condition indices of the bridges passing over the Llobregat River.

5. CONCLUSIONS

This study aims to provide the groundwork for the next generation of Bridge Management Systems (BMS). It establishes a comprehensive Digital Twin (DT) framework for bridge network monitoring that addresses critical information lifecycle management challenges inherent to built assets. By proposing the use of Knowledge Graphs and open Linked Data ontologies, the study achieves a flexible, multi-scale Digital Twin architecture that operates in alignment with industry standards such as IFC, INSPIRE, and ICDD. Finally, it demonstrates the application of the proposed framework within a realistic case study. In this way, we demonstrate our theoretical approach and provide useful tools that facilitate broader adoption.

Ontologies facilitate a modular architecture where stakeholders can independently maintain semantic models. This enables the composition of a customised framework by aligning selected standard ontologies with domain-specific extensions, allowing international standards to coexist with local regulatory definitions and proprietary stakeholder data flexibly. This has been exemplified in the case study, where existing Linked Building Data (LBD) ontologies were complemented with the INSPIRE and Ontology for Visual Inspection of Bridges in the Spanish Road Network to complete the information needs for bridge network monitoring specific to Spanish roads.

Knowledge Graphs operationalise this semantic framework by interconnecting data and ontologies within a unified environment, enabling context-rich querying across disjointed domains. Furthermore, leveraging Linked Data capabilities facilitates web-based data integration, effectively dismantling traditional information silos. This capability was demonstrated in the case study, where the system successfully integrated heterogeneous data sources, including IFC models, INSPIRE geospatial data, and visual inspection records, into a coherent graph structure. Central to this implementation are the developed open-source ETL tools, which enable the conversion of IFC and ICDD containers into a graph and can be reused by academia and industry.

Additionally, the relations in the graph enable multi-scale navigation, providing users with new ways of understanding and visualising data. Furthermore, the automated abstraction of the network-level BCIs from new inputs of damage in individual components of the bridge exemplifies how the in-built logics in the KG can be leveraged to create automations that enrich the information across the graph. Automations like this are very easy to configure through SPARQL-based pipelines and can be flexibly expanded as the graph increases its semantics and data.

However, while this study demonstrates the functional advantages of Knowledge Graph-based systems, the current implementation is limited to a demonstrative scope. Transitioning to industrial application requires a structured evaluation to substantiate these benefits quantitatively. This includes performance benchmarking (API response times, scalability under increasing asset volumes, concurrent user load testing) and usability evaluation through representative asset-management scenarios. Furthermore, limitations related to large-scale deployment, legacy system integration, and real-world institutional adoption are barriers that need to be explicitly addressed.

From an academic perspective, future work to broaden the scope of Digital Twins for bridge network monitoring requires extending the suite of open ontologies to incorporate additional monitoring activities, such as Load Tests, Non-Destructive Testing (NDT), and Minor Destructive Testing (MDT) techniques, alongside data from distributed sensor networks. Expanding the semantics and the amount of interconnected data within the knowledge graph will enable new possibilities for analysing infrastructure across vast territories, supporting increasingly complex decision-making processes.

The successful adoption of these systems hinges on empirical evaluation through practical use cases and collaboration among system developers, administrators, academia, and industry stakeholders. A strong collaborative approach will enable the industry to leverage existing data, ensure interoperability, and drive a seamless transition toward more efficient and resilient information systems that support long-term bridge maintenance practices. By fostering an open and adaptable digital ecosystem, these advancements will empower stakeholders to engage in proactive infrastructure management, contributing to the sustainability and reliability of bridge networks.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial support from the Agency for Management of University and Research Grants (AGAUR) through the FI-2021 grant, which provided the primary funding for the doctoral research in which this article is framed. Additionally, the authors acknowledge the financial support of the European Union's Erasmus+ program.

REFERENCES

AlBalkhy W., Karmaoui D., Ducoulombier L., Lafhaj Z., and Linner T. (2024). Digital twins in the built environment: Definition, applications, and challenges. Automation in Construction, Vol. 162, 105368. DOI: https://doi.org/10.1016/j.autcon.2024.105368.

Akroyd J., Mosbach S., Bhave A. and Kraft M. (2021). Universal digital twin: a dynamic knowledge graph. Data-Centric Engineering, Vol. 2 No. 14. DOI: https://doi.org/10.1017/dce.2021.10.

Akroyd, J. (2022). CReDo Methodology Papers: Implementation. Digital Twin Hub. Available at: https://digitaltwinhub.co.uk/climate-resilience-demonstrator-credo/reports-resources/

Bartha G. and Kocsis S. (2011). Standardization of geographic data: The European inspire directive. European Journal of Geography, Vol. 2, No. 2, 79-89. Retrieved from https://eurogeojournal.eu/index.php/egj/article/view/36

Battle R. and Kolas D. (2011). Geosparql: enabling a geospatial semantic web. Semantic Web Journal, Vol. 3, No. 4, 355-370. Retrieved from https://www.semantic-web-journal.net/sites/default/files/swj176_1.pdf

Berners-Lee T., Hendler J. and Lassila O. (2001). The Semantic Web. Scientific American, Vol. 284, No. 5, 28-37. DOI: https://doi.org/10.1038/scientificamerican0501-34.

Bienvenu, M., Leclère, M., Mugnier, M.-L., & Rousset, M.-C. (2020). Reasoning with Ontologies. In P. Marquis, O. Papini, & H. Prade (Eds.), A Guided Tour of Artificial Intelligence Research: Volume I: Knowledge

Representation, Reasoning and Learning (pp. 185–215). Springer, Cham. DOI: https://doi.org/10.1007/978-3-030-06164-7_6

Boje C., Guerriero A., Kubicki S., and Rezgui Y. (2020). Towards a semantic construction digital twin: Directions for future research. Automation in Construction, Vol. 114, 103179. DOI: https://doi.org/10.1016/j.autcon.2020.103179.

Broo D. G. and Schooling J. (2023). Digital twins in infrastructure: definitions, current practices, challenges, and strategies. International Journal of Construction Management, Vol. 23, No. 7, 1254–1263. DOI: https://doi.org/10.1080/15623599.2021.1966980.

Chacón R., Casas J. R., Ramonell C., Posada H., Stipanovic I., and Škarić S. (2023). Requirements and challenges for infusion of SHM systems within Digital Twin platforms. Structure and Infrastructure Engineering, Vol. 21, No. 4, 1–17. DOI: https://doi.org/10.1080/15732479.2023.2225486.

Chacón R., Ramonell C., Posada H., Sierra P., Tomar R., Martínez de la Rosa C., Rodriguez A., Koulalis I. and Wagmeister, S. (2023b). Digital twinning during load tests of railway bridges - case study: the high-speed railway network, Extremadura, Spain. Structure and Infrastructure Engineering, 20(7–8), 1105–1119. https://doi.org/10.1080/15732479.2023.2264840

Chacón R., Posada H., Ramonell C., Jungmann M., Hartmann T., Khan R., and Tomar R. (2024). Digital twinning of building construction processes. Case study: A reinforced concrete cast-in structure. Journal of Building Engineering, Vol. 84, 108522. DOI: https://doi.org/10.1016/j.jobe.2024.108522.

Chadzynski A., Li S., Grišiute A., Chua J., Hofmeister M., Yan J., Tai H. Y., Lloyd E., Tsai Y. K., Agarwal M. (2023). Semantic 3D city interfaces—Intelligent interactions on dynamic geospatial knowledge graphs. Data-Centric Engineering, Vol. 4 No. 20, DOI: https://doi.org/10.1017/dce.2023.14.

Chaudhri V. K., Baru C., Chittar N., Dong X. L., Genesereth M., Hendler J., Kalyanpur A., Lenat D. B., Sequeda J., Vrandečić D. and Wang K. (2022). Knowledge graphs: Introduction, history, and perspectives. AI Magazine, Vol. 43, No. 1, 17–29. DOI: https://doi.org/10.1002/aaai.12033.

Daró P., Alovisi I., Mancini G., Negri S., Bigaj-van Vliet A. and van Meerveld H. (2022). Lessons learned from proactive maintenance practices for concrete bridges. Proceedings of the fib Congress, Oslo, Norway, 12–16. Retrieved from https://im-safe-project.eu/wp-content/uploads/2021/01/WP3-Lessons-learned-from-proactive-maintenance-practices-for-concrete-bridges.pdf

Garijo D. (2017). WIDOCO: a wizard for documenting ontologies. In The Semantic Web. ISWC 2017: 16th International Semantic Web Conference, Vienna, Austria, October 21-25, 2017, Proceedings, Part II (Springer), Vol. 16, 94-102. DOI: https://doi.org/10.1007/978-3-319-68204-4_9

Garijo D. and Poveda-Villalón M. (2020). Best practices for implementing fair vocabularies and ontologies on the web. Applications and practices in ontology design, extraction, and reasoning (IOS Press), 39-54. DOI: https://doi.org/10.3233/SSW200034.

Gkoumas K., Marques Dos Dantos F., Van Balen M., Tsakalidis A., Ortega Hortelano A., Grosso M., Haq G. and Pekár F. (2019). Research and innovation in bridge maintenance, inspection and monitoring. Publications office of the Euopean Union. Retrieved from https://www.researchgate.net/profile/Konstantinos-Gkoumas/publication/330831101_Research_and_innovation_in_bridge_maintenance_inspection_and_monitoring_A_European_perspective_based_on_the_Transport_Research_and_Innovation_Monitoring_and_Information_System_TRIMIS/links/5c55a6e4299bf12be3f66f7e/Research-and-innovation-in-bridge-maintenance-inspection-and-monitoring-A-European-perspective-based-on-the-Transport-Research-and-Innovation-Monitoring-and-Information-System-TRIMIS.pdf

Halmetoja E. (2022). The role of digital twins and their application for the built environment. Industry 4.0 for the Built Environment (Springer International Publishing), 415–442. DOI: https://doi.org/10.1007/978-3-030-82430-3_18.

Hamdan A. H., Bonduel M. and Scherer R. J. (2019). An ontological model for the representation of damage to constructions. CEUR Workshop Proceedings, Vol. 2389, 64-77. Retrieved from https://lirias.kuleuven.be/retrieve/0c190b27-242b-4e2b-a6f2-1c3872251cc4

Hamdan A. H. and Scherer R. J. (2020). Integration of BIM-related bridge information in an ontological knowledgebase. Linked Data in Architecture and Construction (LDAC), 77-90. Retrieved from https://ceur-ws.org/Vol-2636/06paper.pdf

Hayes S., Dent C., Mawdsley B., Judson R., and Collingwood T. (2022). CReDo: Overview Report. Technical report, Centre for Digital Built Britain. Retrieved from https://www.repository.cam.ac.uk/items/a5824ab5-606a-4597-9ba1-2c7595a94b13

Hitzler P., Krötzsch M., Parsia B., Patel-Schneider P. F., and Rudolph S. (2009). OWL 2 Web Ontology Language Primer. W3C recommendation, Vol. 27, No. 1, 123. Retrieved from https://www.w3.org/TR/2009/PR-owl2-primer-20090922/all.pdf

Hofmeister M., Bai J., Brownbridge G., Mosbach S., Lee K. F., Farazi F., Hillman M., Agarwal M., Ganguly S., Akroyd J. (2024). Semantic agent framework for automated flood assessment using dynamic knowledge graphs. Data-Centric Engineering, Vol. 5 No. 14. DOI: https://doi.org/10.1017/dce.2024.11.

Hogan A., Blomqvist E., Cochez M., d'Amato C., De Melo G., Gutierrez C., Kirrane S., Labra Gayo J. E., Navigli R., Neumaier S. and others (2021). Knowledge graphs. ACM Computing Surveys (Csur). DOI: https://doi.org/10.1145/3447772.

Iacovino C., Turksezer Z.I., Giordano P.F. and Limongelli M.P. (2022). A Survey of Bridge Condition Rating Systems. Proceedings of the 1st Conference of the European Association on Quality Control of Bridges and Structures. EUROSTRUCT 2021. Lecture Notes in Civil Engineering (Springer), Vol 200. DOI: https://doi.org/10.1007/978-3-030-91877-4_2

Jain V. and Singh M. (2013). Ontology development and query retrieval using protégé tool. International Journal of Intelligent Systems and Applications, Vol. 9, No. 9, 67-75. DOI: https://doi.org/10.5815/ijisa.2013.09.08.

Kim Y. A. J. and Yeo H. (2016). A comparative study of time-based maintenance and condition-based maintenance for optimal choice of maintenance policy. Structure and Infrastructure Engineering, Vol. 12, No. 12, pp. 1525–1536. DOI: https://doi.org/10.1080/15732479.2016.1149871.

Mandirola M., Casarotti C., Peloso S., Lanese I., Brunesi E. and Senaldi I. (2022). Use of UAS for damage inspection and assessment of bridge infrastructures. International Journal of Disaster Risk Reduction, Vol. 72, 102824. DOI: https://doi.org/10.1016/j.ijdrr.2022.102824.

McBride B. (2004). The resource description framework (RDF) and its vocabulary description language RDFS. Handbook on ontologies (Springer), 51-65.

Lim M. Q., Wang X., Inderwildi O. and Kraft M. (2022). The world avatar—a world model for facilitating interoperability. Intelligent Decarbonisation: Can Artificial Intelligence and Cyber-Physical Systems Help Achieve Climate Mitigation Targets? (Springer), 39-53. DOI: https://doi.org/10.1007/978-3-030-86215-2_4.

Lohmann S., Link V., Marbach E. and Negru S. (2015). WebVOWL: Web-based visualization of ontologies. Knowledge Engineering and Knowledge Management: EKAW 2014 Satellite Events, VISUAL, EKM1, and ARCOE-Logic, Linköping, Sweden, November 24-28, 2014. Revised Selected Papers (Springer), Vol. 19, 154-158. DOI: https://doi.org/10.1007/978-3-319-17966-7_21.

Liu P, Shi Y., Xiong R. and Tang P. (2023). Quantifying the reliability of defects located by bridge inspectors through human observation behavioral analysis. Developments in the Built Environment, Vol. 14, 100167. DOI: https://doi.org/10.1016/j.dibe.2023.100167.

Lu Q., Parlikad A. K., Woodall P., Don Ranasinghe G., Xie X., Liang Z., Konstantinou E., Heaton J., and Schooling J. (2020). Developing a digital twin at building and city levels: Case study of West Cambridge campus. Journal of Management in Engineering, Vol. 36, No. 3, 05020004. DOI: https://doi.org/10.1061/(ASCE)ME.1943-5479.0000763.

Lu R. and Brilakis I. (2019). Digital twinning of existing reinforced concrete bridges from labelled point clusters. Automation in Construction, 105, 102837. DOI: https://doi.org/10.1016/j.autcon.2019.102837

Pal A., Lin J. J., Hsieh S.-H., and Golparvar-Fard M. (2023). Automated vision-based construction progress monitoring in the built environment through digital twin. Developments in the Built Environment, Vol. 16, No. 100247. DOI: https://doi.org/10.1016/j.dibe.2023.100247.

Pauwels P. and Terkaj W. (2016). EXPRESS to OWL for construction industry: Towards a recommendable and usable ifcOWL ontology. Automation in Construction, Vol. 63, 100-133. DOI: https://doi.org/10.1016/j.autcon.2015.12.003.

Pauwels P., Krijnen T., Terkaj W. and Beetz J. (2017). Enhancing the ifcOWL ontology with an alternative representation for geometric data. Automation in Construction, Vol. 80, 77-94. DOI: https://doi.org/10.1016/j.autcon.2017.03.001.

Pauwels P., Zhang S. and Lee Y. C. (2017). Semantic web technologies in AEC industry: A literature review. Automation in Construction, Vol. 73, 145-165. DOI: https://doi.org/10.1016/j.autcon.2016.10.003.

Pauwels P., Costin A. and Rasmussen M. H. (2021). Knowledge graphs and linked data for the built environment. Industry 4.0 for the Built Environment: Methodologies, Technologies and Skills (Springer), 157–183. DOI: https://doi.org/10.1007/978-3-030-82430-3_7.

Pauwels P., Shelden D., Brouwer J., Sparks D., Nirvik S. and McGinley T. P. (2022). Open data standards and BIM on the cloud. Buildings and Semantics (CRC Press), 101–136. DOI: https://doi.org/10.1201/9781003204381-6.

Posada H., Chacón R. and Ramonell C. (2024). Job roles for digital twinning building construction processes. Introducing the digital twin manager position. International Journal of Construction Management, 1-17. DOI: https://doi.org/10.1080/15623599.2024.2417632

Poveda-Villalón M., Gómez-Pérez A. and Suárez-Figueroa M. C. (2014). Oops! (ontology pitfall scanner!): An on-line tool for ontology evaluation. International Journal on Semantic Web and Information Systems, Vol. 10, No. 2, 7-34. DOI: https://doi.org/10.4018/ijswis.2014040102.

Pregnolato M., Guner S., Voyagaki E., De Risi R., Carhart N., Gavriel G., Tully P., Tryfonas T., Macdonald J. and Taylor C. (2022). Towards Civil Engineering 4.0: Concept, workflow and application of Digital Twins for existing infrastructure. Automation in Construction, Vol. 141, 104421. DOI: https://doi.org/10.1016/j.autcon.2022.104421.

Quek H. Y., Hofmeister M., Rihm S. D., Yan J., Lai J., Brownbridge G., Hillman M., Mosbach S., Ang W., Tsai Y. K. (2024). Dynamic knowledge graph applications for augmented built environments through "The World Avatar". Journal of Building Engineering, 109507. DOI: https://doi.org/10.1016/j.jobe.2024.109507.

Ramonell C., Chacón R. and Posada, H. (2023). Knowledge graph-based data integration system for digital twins of built assets. Automation in Construction, 156, 105109. DOI: https://doi.org/10.1016/j.autcon.2023.105109

Ramonell C. (2025) Graph-driven digital twins as assistants to bridge maintenance. Tesi doctoral, UPC, Departament d'Enginyeria Civil i Ambiental, 2025. DOI: https://doi.org/10.5821/dissertation-2117-432851

Rasmussen M. H., Lefrançois M., Bonduel M., Hviid C. A. and Karlshøj J. (2018). OPM: An ontology for describing properties that evolve over time. Proceedings of the 6th Linked Data in Architecture and Construction Workshop

Rasmussen M. H., Lefrançois M., Schneider G. F. and Pauwels P. (2020). BOT: The building topology ontology of the W3C linked building data group. Semantic Web, Vol. 12, No. 1, 143-161. DOI: https://doi.org/10.3233/SW-200385

Sacks R., Brilakis I., Pikas E., Xie H. S., and Girolami M. (2020). Construction with digital twin information systems. Data-Centric Engineering, Vol. 1, No. 14. DOI: https://doi.org/10.1017/dce.2020.16.

Sanchez-Silva M., Frangopol D. M., Padgett J. and Soliman M. (2016). Maintenance and Operation of Infrastructure Systems: Review. Journal of Structural Engineering, Vol. 142, No. 9, F4016004. DOI: https://doi.org/10.1061/(ASCE)ST.1943-541X.0001543.

Seo J., Duque L. & Wacker J. (2018). Drone-enabled bridge inspection methodology and application. Automation in Construction, Vol. 94, 112–126. DOI: https://doi.org/10.1016/j.autcon.2018.06.006.

Srikanth I. and Arockiasamy M. (2020). Deterioration models for prediction of remaining useful life of timber and concrete bridges: A review. Journal of Traffic and Transportation Engineering, Vol. 7, No. 2, 152–173. DOI: https://doi.org/10.1016/j.jtte.2019.09.005.

Studer R., Benjamins V. R. and Fensel D. (1998). Knowledge engineering: Principles and methods. Data & Knowledge Engineering, Vol. 25, No. 1-2, 161-197. DOI: https://doi.org/10.1016/S0169-023X(97)00056-6.

Terkaj W. and Pauwels P. (2017). A method to generate a modular ifcOWL ontology. 8th International Workshop on Formal Ontologies meet Industry: Episode 3: The Tyrolean Autumn of Ontology (CEUR-WS). Retrieved from https://ceur-ws.org/Vol-2050/FOMI_paper_3.pdf

Tuhaise V. V., Tah J. H. M., and Abanda F. H. (2023). Technologies for digital twin applications in construction. Automation in Construction, Vol. 152, 104931. DOI: https://doi.org/10.1016/j.autcon.2023.104931.

Vergara C., Bahsoon R., Theodoropoulos G., Yanez W. and Tziritas N. (2023). Federated digital twin. IEEE/ACM 27th International Symposium on Distributed Simulation and Real Time Applications (DS-RT), pp. 115-116. DOI: https://doi.org/10.1109/DS-RT58998.2023.00027

Zhang H., Qi Q., and Tao F. (2022). A multi-scale modelling method for digital twin shopfloor. Journal of Manufacturing Systems, Vol. 62, 417–428. DOI: https://doi.org/10.1016/j.jmsy.2021.12.011.