Available online at www.sciencedirect.com

ScienceDirect

Procedia Structural Integrity 85 (2026) 1121–1128

Structural Integrity Procedia

www.elsevier.com/locate/procedia

SMAR 2026 - 8th International Conference on Smart Monitoring, Assessment and Rehabilitation of Civil Structures

Ontology-based Integration of Spatial Sensor Information in Digital Twins

Fabian Rolke[a]*, Henrik Becks[b], Anne Göbels[a], Oliver Schulz[a], Jakob Beetz[a]

[a] *Design Computation, RWTH Aachen University, 52072 Aachen, Germany*
[b] *Chair and Institute of Structural Concrete, RWTH Aachen University, 52072 Aachen, Germany*

Abstract

Structural Health Monitoring (SHM) systems generate measurement data whose interpretation depends on the precise spatial contextualization of sensors relative to the monitored structure, allowing engineers to relate measurements to specific structural components, load paths, or environmental influences. In practice, sensor positions are often documented across distributed heterogeneous resources such as technical drawings, photographs, and textual descriptions, limiting accessibility, interoperability, and machine-readability. As these formats are limited to 2D space, relating them and their implicit spatial sensor and observation information to the physical infrastructure asset typically requires cross-referencing multiple resources.

This paper presents a method for transforming spatially implicit sensor documentation into a Resource Description Framework (RDF) dataset, utilizing the Space Types Ontology (SPOT) and the Relative Location Ontology (RELOC) for expressing spatial relations, and the Sensor, Observation, Sample, and Actuator Ontology (SOSA) for sensor semantics. The resulting graph enables qualitative and quantitative queries on topological and directional relations between sensors, their observations, and the resources documenting them. The approach contributes to the broader goal of predictive maintenance through Digital Twins for condition assessment and prognosis models, supported by spatially accessible heterogeneous data.

© 2026 The Authors. Published by ELSEVIER B.V.
This is an open access article under the CC BY-NC-ND license (https://creativecommons.org/licenses/by-nc-nd/4.0)
Peer-review under responsibility of SMAR 2026 Organizers

Keywords: Spatial Linking; Structural Health Monitoring; Digital Twins; Sensor Localization

* Corresponding author. Tel.: +49-241-80-952-35.
E-mail address: rolke@dc.rwth-aachen.de

2452-3216 © 2026 The Authors. Published by ELSEVIER B.V.
This is an open access article under the CC BY-NC-ND license (https://creativecommons.org/licenses/by-nc-nd/4.0)
Peer-review under responsibility of SMAR 2026 Organizers
10.1016/j.prostr.2026.08.097

1. Introduction

Current efforts to extend the service life of aging infrastructure, such as bridges, including the development of Digital Twins, increasingly rely on Structural Health Monitoring (SHM) systems (Kang et al., 2025). SHM systems employ distributed sensor networks to continuously observe structural behavior and environmental influences affecting infrastructure assets.

Most sensors measure phenomena along specific directions or within defined spatial extents, such as longitudinal displacement, vertical acceleration, or localized temperature fields. Precise spatial contextualization of a sensor and its observation space is therefore essential for engineers to relate measurements to specific structural components, load paths, and environmental influences, supporting anomaly identification and maintenance planning (Glisic et al., 2014). However, information about the position, orientation, and extent of a sensor and its observation is often distributed across heterogeneous and semi-structured documentation resources, including drawings, photographs, and textual descriptions. This fragmentation creates interoperability challenges, increases the risk of inconsistent information, and limits machine-readable integration into digital asset management workflows.

In addition, textual descriptions frequently encode spatial relations in qualitative terms (e.g., “The temperature sensor is contained in the bottom slab of the girder.”). This motivates the need for explicit qualitative representations of spatial relations between sensors, observations, monitored components, and the resources documenting them.

A common approach to improve interoperability is the integration of sensor information into the Industry Foundation Classes (IFC) schema. Such approaches provide type-level interoperability by enabling stakeholders to interpret information according to a shared schema (Törmä, 2013). However, type-level interoperability does not explicitly capture that information contained in different resources refers to the same real-world entity.

Such relationships are expressed through instance-level interoperability (Törmä, 2014), for which the Resource Description Framework (RDF) (W3C, 2014) offers a well-established mechanism based on Uniform Resource Identifiers (URIs) and semantic relationships between entities. While RDF provides the technical foundation for instance-level integration, vocabularies capable of qualitatively representing three-dimensional resources and their spatial relationships remain limited.

The Relative Location Ontology (RELOC) (Göbels & Beetz, 2024) addresses this gap by representing directional and topological relationships qualitatively, e.g., reloc:containedInBottom. The Space Types Ontology (SPOT) (Göbels et al., 2024, 2025) complements this by providing a vocabulary for describing different space types, e.g., spot:EntitySpace, spot:DocumentSpace, enabling the distinction between a real-world entity and its representations across documentation resources.

This paper examines the combined use of RELOC, SPOT, and the Sensor, Observation, Sample, and Actuator Ontology (SOSA) (Compton et al., 2012; Janowicz et al., 2019) to semantically and spatially interlink heterogeneous sensor documentation resources, enabling machine-readable spatial contextualization of the physical sensor instances and their observations. The proposed approach is guided by the following research questions:

RQ-1: *How can heterogeneous resources, documenting sensors, be explicitly mapped to the infrastructural asset using spatial, machine-readable vocabularies?*

RQ-2: *How can sensors and their observations be formally distinguished and expressed using shared spatial vocabularies?*

To answer these questions, we analyze heterogeneous sensor documentation resources and spatially align them to a digital representation of the built asset through manual annotation. Primarily utilizing the SPOT, RELOC, and SOSA ontologies, the annotations are transformed into an RDF graph. The graph's ability to explicitly express semantic and spatial relationships among sensors, their observations, and features of interest is evaluated using SPARQL (W3C, 2013) queries. The approach is tested on the sensor installation documentation of the Nibelungen Bridge in Worms, Germany. The monitoring setup is described in more detail in Eisermann et al. (2024).

2. Related Work

Approaches to sensor spatialization in the Architecture, Engineering, Construction, and Operation (AECO) sector broadly fall into Building Information Modeling (BIM)-based and ontology-based strategies, each with distinct capabilities and limitations regarding interoperability and data heterogeneity.

BIM-based approaches confine sensor spatialization to specific authoring environments and schemas. Valinejadshoubi et al. (2017) demonstrate this within a single proprietary tool, achieving neither type-level compatibility across models nor instance-level linking. Additionally, remodeling and positioning sensors within BIM software or a federated BIM model (Preidel et al., 2018) requires expert domain knowledge and explicit geometric and semantic knowledge of the asset. As project scale increases, dedicated BIM management roles are required (Borrmann et al., 2018), resulting in a high-effort, cost-intensive process.

While ontology-based approaches offer tool-agnostic, vocabulary-driven integration, achieving sufficient spatial expressivity without resorting to non-standard SPARQL extensions remains a persistent challenge.

Eneyew et al. (2022) propose a multi-layer digital twin architecture, BIM Internet-of-Things Data-Integration (BIM-IoTDI), addressing semantic interoperability between BIM and IoT data sources through shared vocabularies such as SOSA and BOT. SOSA and SSN are primarily used to describe semantic relationships among sensors, observable properties, observations, and features of interest, as well as technical sensor properties such as accuracy, latency, and battery lifetime. Ontologies such as BOT (Rasmussen et al., 2020), BROT/BRCOMP (Hamdan & Scherer, 2020), and GOM/FOG (Bonduel et al., 2019) describe building topology and geometry but provide limited support for expressing spatial relationships between entities.

BimSPARQL (Zhang et al., 2018) partially addresses this by extending SPARQL with custom functions that operate on WKT-serialized geometry derived from IfcOWL (Beetz et al., 2009; Pauwels et al., 2017), enabling a range of diverse quantitative queries. However, such language extensions reduce portability across standard SPARQL environments, limiting their widespread use. Furthermore, because BimSPARQL relies on quantitative encodings, its applicability to qualitative representations of spatial information is restricted.

No existing approach sufficiently combines sensor semantics and their heterogeneous representations with instance-level interoperability to express where a sensor is located relative to the structure, how its observation relates to monitored components, and how these relationships are depicted across heterogeneous resources. This gap motivates the approach presented in this paper.

3. Methodology

The proposed method transforms spatially implicit sensor documentation into an RDF dataset, making spatial information fragments distributed across heterogeneous documentation resources queryable (Fig. 1).

After analyzing the sensor documentation and identifying potential information needs (Section 3.1), the individual pages are loaded into a custom annotation application. Each page is annotated manually (Section 3.2). The annotations are then transformed into an RDF graph (Section 3.3). SPOT-, RELOC-, and SOSA-aligned SPARQL queries are used to retrieve spatial relations among sensors, their observations, monitored components, and the resources in which they appear (Section 4). By cross-referencing multiple 2D sensor appearances in separate reference spaces with complementary orientations, the sensor's 3D position and orientation can be inferred (Section 4).

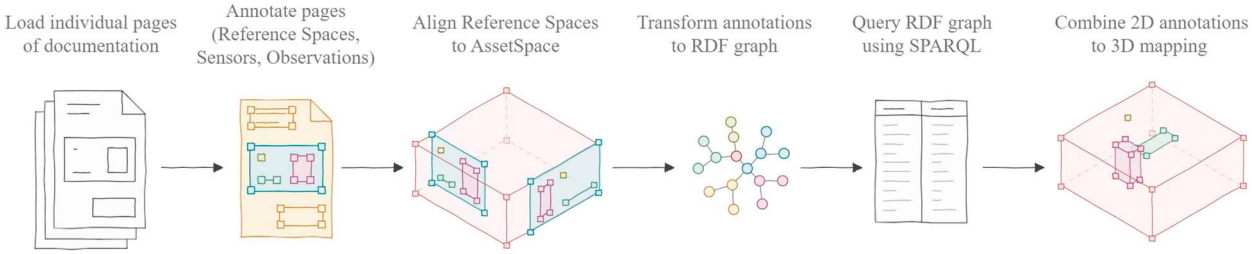


Fig. 1. Proposed method for transforming fragmented sensor documentation resources into a queryable RDF graph, enabling explicit spatial contextualization of sensors and their observations.

3.1. Analysis of Sensor Documentation

Sensor documentation typically contains information on the types, properties, and locations of installed sensors, as well as the sensor network architecture. This information is often stored in non-machine-readable or semi-structured formats. Common data types include technical drawings, photographs, and text.

Technical drawings provide information about a sensor's approximate position within a larger frame of reference but rarely include details about its orientation or extent. Photographs offer a limited frame of reference but are a reliable source of information about the precise position, orientation, and extent of sensors. Text cannot be directly mapped to the asset, but it provides explicit, sometimes metric, information about locational or directional relations among other resources, e.g., “Temperature Sensor […] 5.5 m east of the pier edge”.

3.2. Annotation of Sensor Documentation

To streamline the annotation process, an application was developed (Fig. 2). Reference spaces (technical drawings, photographs, text) are framed, classified, and spatially aligned with the built asset using axis and boundary-point mappings. Text within reference spaces is automatically extracted via Optical Character Recognition (OCR) and stored as a comment for that space. The position, orientation, and extent of sensors, along with their observed properties, are annotated relative to the reference space in which they appear. Each sensor appearance is linked to its corresponding sensor instance. After checking for severe spatial compliance errors, e.g., multiple appearances of the same sensor with different alignments to the asset, the annotations are exported as an RDF graph.

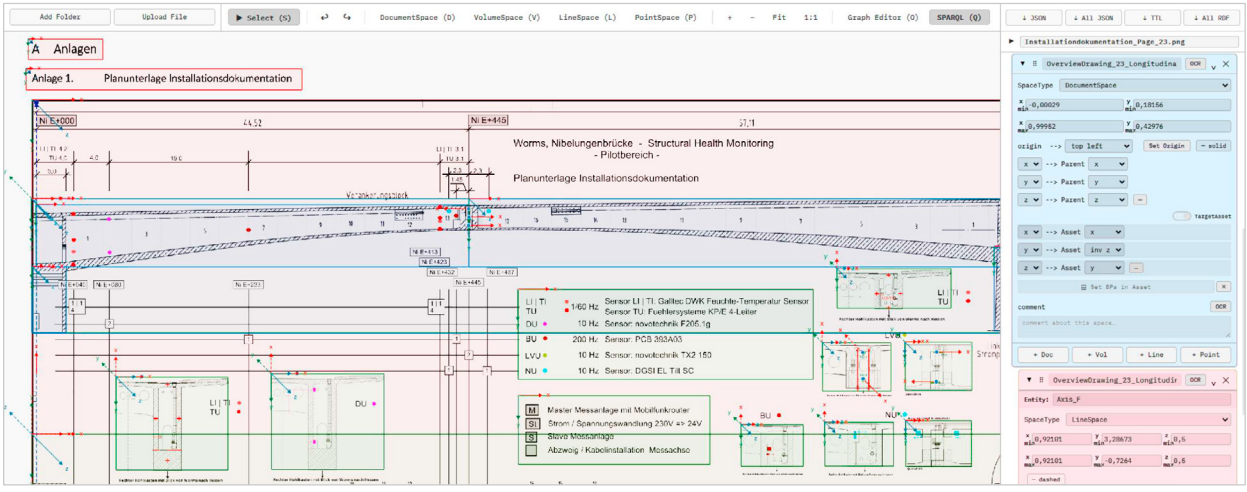


Fig. 2. Interface of the application during annotation of an overview drawing. (Annotations made by authors - Data-Source: MKP GmbH)

3.3. RDF Representation of Spatial Relations between Sensors, Observable Properties, and Components

The RDF representation uses the SOSA ontology to assign instances to the classes *sosa:Sensor*, *sosa:Observation, sosa:ObservableProperty*, or *sosa:FeatureOfInterest*.

To enable component-based queries, instances can also be assigned classes of the Bridge Component Ontology (BRCOMP) (Hamdan & Scherer, 2020), e.g., *brcomp:Girder*.

Space types are defined using the SPOT ontology (Göbels et al., 2024), which distinguishes between different spaces and decomposes spatial resources into primitive geometries (volumes, areas, lines, and points) defined by (qualitative or quantitative) boundary points and axes expressing orientation relative to other spaces.

The Relative Location Ontology (RELOC) (Göbels & Beetz, 2024) combines topological (containedIn, intersect, meet, and disjoint) and directional (e.g., top, left, longitudinalCenter) terms to predicates such as *reloc:containedInTop*. For this method, RELOC is used to qualitatively express spatial relationships between spaces.

Expressing a sensor, its observation, and its feature of interest as individual instances enables querying semantics and spatial relations across all three entity spaces. The following two examples illustrate distinct spatial mapping cases between these instances (Fig. 3, 4):

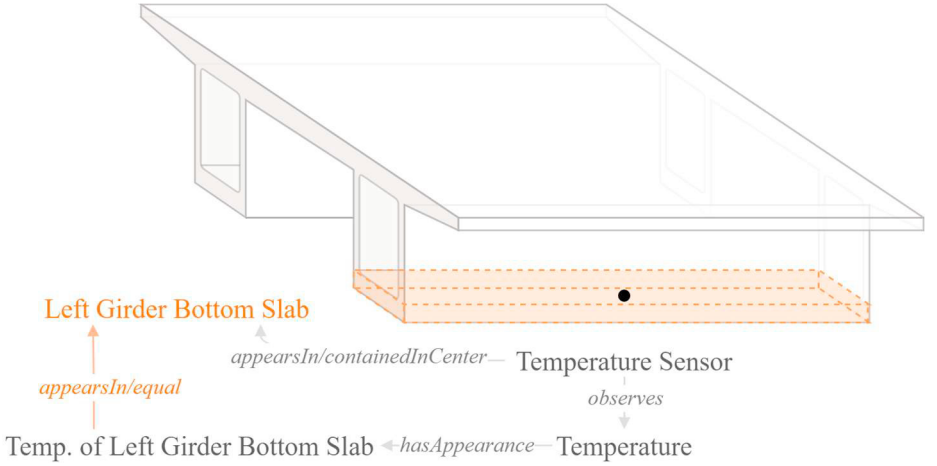


Fig. 3. *Example A:* Temperature sensor placed inside the bottom slab of a girder. The sensor observes an omnidirectional property that cannot be mapped to a component axis, but rather the component as a whole.

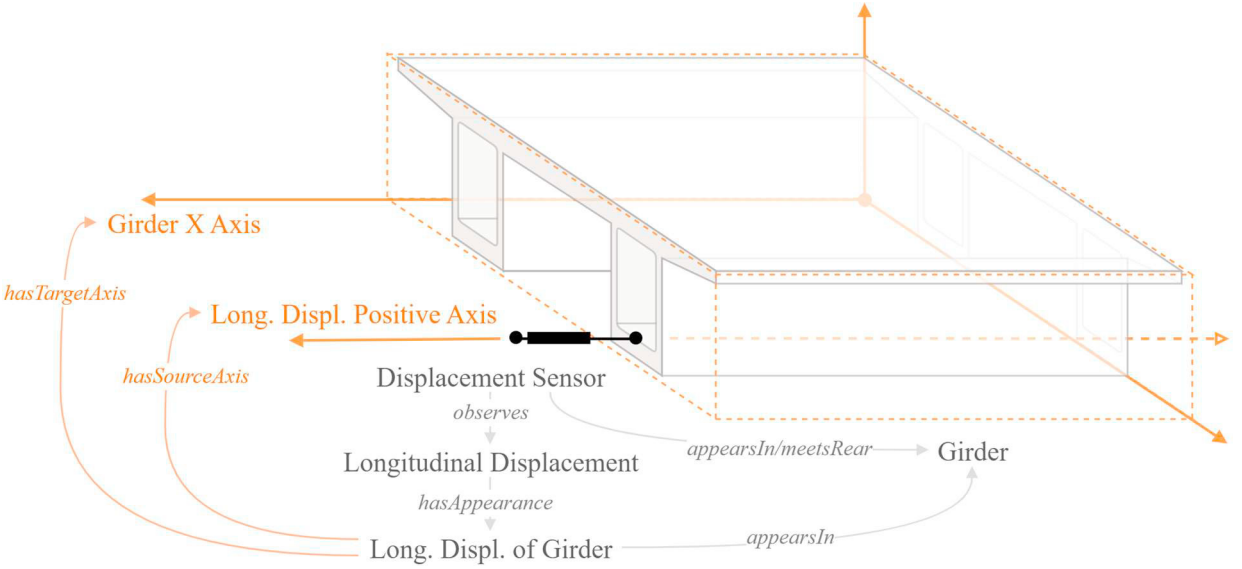


Fig. 4. *Example B:* Displacement sensor, placed at the rear of a girder, measuring its longitudinal displacement. The sensor observes a directional property that can be mapped to the corresponding component x-axis.

3.4. Use Case – Nibelungen Bridge

To test the described method, it was applied to a documentation of sensors installed on the Nibelungen Bridge in Worms, Germany.

The documentation PDF consists of 42 pages, containing 165 reference spaces (drawings, photographs, text), most of which hold spatial information about the position, type, orientation, and spatial extent of 20 sensors and 4 sensor network components.

While most sensor positions were determinable through cross-referencing of complementary technical drawings, ambiguous edge cases required additional cross-referencing with textual or photographic resources. These included combi-sensors (Fig. 5) and cases where a sensor's precise position relative to a bridge component, either inside or on its surface, was visually not distinguishable (Fig. 6).

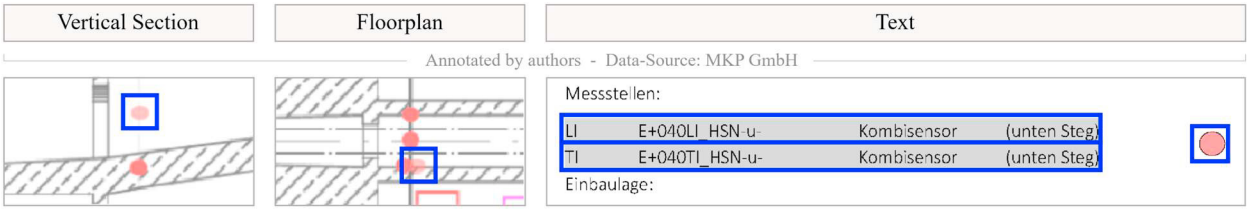


Fig. 5. Edge case in which a single dot represents two sensors (combi-sensor). Cross-referencing drawings with text is required to classify the dot as appearances of two distinct sensor instances.

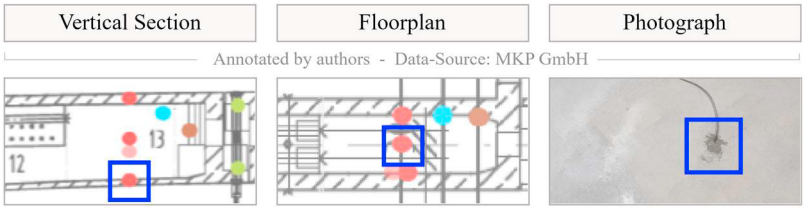


Fig. 6. Edge case in which determining the sensor installation position (surface-mounted or embedded) requires cross-referencing drawings with photographic documentation.

4. RDF Output and Query Results

The full RDF representation of the Nibelungen Bridge sensor documentation, exemplary SPARQL queries, and the custom annotation application are available online (Rolke et al., 2026).

SPARQL query results can be used to navigate to appearances of a sensor inside a documentation resource based on the spatial attributes of the device, its observation, or the monitored component/space (Fig. 7). Alternatively, results can be used to construct 3D information from complementary 2D resources, such as the position in 3D space, or an axis mapping across the x, y, and z axes relative to the asset space (Fig. 7).

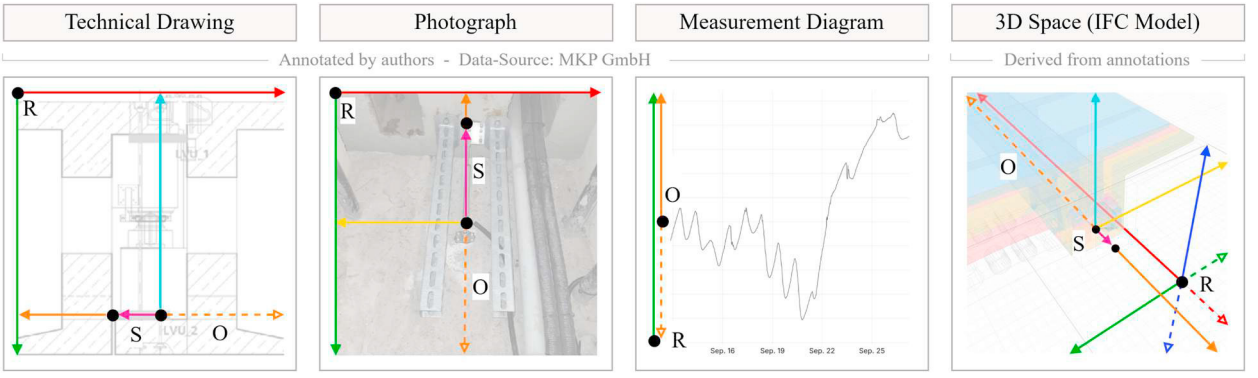


Fig. 7. Query results used to display the boundary points and axis alignments of a sensor (S) and its observation (O) relative to 2D reference spaces (R) (drawing, photograph, diagram) or 3D space (IFC model).

5. Discussion

This section reflects on insights from the annotation process and the usability and limitations of the resulting RDF representation of the sensor documentation.

In its current state, the proposed method's information acquisition largely relies on manual annotation effort. Automating the detection and classification of sensors in heterogeneous documentation resources through object detection networks could significantly reduce this overhead. More challenging is the automated cross-referencing of sensor appearances across resources to their corresponding RDF instance, as this requires concurrent semantic and spatial reasoning, a task complicated by ambiguous edge cases, the absence of shared identifiers, and the heterogeneity of representation formats.

The query results demonstrate the RDF graph's ability to make previously implicit 3D spatial information, fragmented and distributed across separate 2D resources, explicit (Fig. 7). The combined use of SPOT, RELOC, and SOSA is therefore sufficient to express the spatial types and relations of sensors, their observations, and the monitored components documented in drawings, photographs, and text, addressing RQ-1. Furthermore, SPOT, in conjunction with SOSA, enables the formal distinction between a sensor instance and its observation, each of which can be mapped to monitored components as individual spatial entities via distinct RELOC predicates, thereby addressing RQ-2.

As sensors and their observations in this use case are primarily aligned orthogonally to their feature of interest, expressing their directional relations qualitatively is acceptable. Where qualitatively representing the topological relation of a sensor to its monitored component does not provide sufficient mapping precision, a subspace of a component can be declared (e.g., the bottom slab of a girder) to which the sensor and its observation can be mapped.

By expressing directional and topological information qualitatively, query results are human-readable. This means sensor measurements can be directly related to specific axes or subparts of monitored components, with no further processing or viewer required. However, this spatial expressiveness comes at the cost of computational efficiency. Deducing spatial relations between implicitly related entities requires recursive calculations that quickly exceed SPARQL's native pattern-matching capabilities as the number of calculation steps increases. This limitation suggests the need to store spatial sensor information both qualitatively, for direct human interpretability, and quantitatively, for machine-based transformation calculations. Future work could explore integrating them into a combinatorial framework that leverages the advantages of each representation in spatially contextualizing sensor measurement data.

A minor limitation is the dataset's direct accessibility to domain experts unfamiliar with formal query languages. A natural-language interface based on an LLM could address this by translating user questions into SPARQL queries, using context derived from the graph and the ontologies used. This would lower the barrier for non-expert users and support wider dissemination of the method.

6. Conclusion

This paper examined the use of the RELOC, SPOT, and SOSA ontologies to describe the spatial semantics of sensors, their observations, components, and the heterogeneous resources that document them, thereby enabling sensor instances and their observations to be explicitly related to the structural components they monitor.

Sensor installation documentation from the Nibelungen Bridge served as the use case for transforming spatially implicit information distributed across heterogeneous resources into a queryable RDF dataset. The results demonstrate that the three ontologies are jointly sufficient to explicitly express semantic and spatial relations between sensors, their observations, monitored components, and the heterogeneous resources that document them.

The approach contributes to the broader goal of predictive maintenance through Digital Twins for condition assessment and prognosis models by making distributed heterogeneous documentation resources spatially accessible.

Future work should address the automation of sensor detection and cross-referencing across heterogeneous resources, a combinatorial qualitative-quantitative spatial framework, and LLM-based natural-language querying for non-expert users.

Acknowledgements

This research is part of the Spatio-B-RAG project funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 562882930. The project is part of the DFG-funded Priority Program SPP 2388, "Hundred plus - Extending the Lifetime of Complex Engineering Structures through Intelligent Digitalization".

References

Beetz, J., van Leeuwen, J., de Vries, B., 2009. IfcOWL: A case of transforming EXPRESS schemas into ontologies. Artificial Intelligence for Engineering Design, Analysis and Manufacturing 23(1), 89–101. https://doi.org/10.1017/S0890060409000122

Bonduel, M., Wagner, A., Pauwels, P., Vergauwen, M., Klein, R., 2019. Including widespread geometry formats in semantic graphs using RDF literals. Computing in Construction 1, 341–350. https://doi.org/10.35490/EC3.2019.166

Borrmann, A., König, M., Koch, C., Beetz, J. (Eds.), 2018. Building Information Modeling: Technology Foundations and Industry Practice. Springer International Publishing. https://doi.org/10.1007/978-3-319-92862-3

Compton, M. et al., 2012. The SSN ontology of the W3C semantic sensor network incubator group. Journal of Web Semantics 17, 25–32. https://doi.org/10.1016/j.websem.2012.05.003

Eisermann, C., Voigt, C., Marx, S., Kang, C., 2024. The initial structural health monitoring system of the Nibelungen Bridge Worms. E-Journal of Nondestructive Testing 29(07). https://doi.org/10.58286/29620

Eneyew, D.D., Capretz, M.A.M., Bitsuamlak, G.T., 2022. Toward Smart-Building Digital Twins: BIM and IoT Data Integration. IEEE Access 10, 130487–130506. https://doi.org/10.1109/ACCESS.2022.3229370

Glisic, B., Yarnold, M.T., Moon, F.L., Aktan, A.E., 2014. Advanced Visualization and Accessibility to Heterogeneous Monitoring Data. Computer-Aided Civil and Infrastructure Engineering 29, 382–398. https://doi.org/10.1111/mice.12060

Göbels, A., Beetz, J., 2024. Relative Location Ontology: An ontological model for representing directional topological relationships between spatial entities in oriented space. Presented at the LDAC 2024 – 12th Linked Data in Architecture and Construction Workshop, Bochum, Germany.

Göbels, A., Schulz, O., Beetz, J., 2024. Towards a Common Digital Space: Proposing a Schema for Spatially Linking Heterogeneous Resources. In: Proceedings of the 41st International Conference of CIB W78, Marrakech, Morocco.

Göbels, A., Schulz, O., Beetz, J., 2025. Space Types Ontology (SPOT). https://w3id.org/spot#

Hamdan, A.-H., Scherer, R., 2020. Integration of BIM-related bridge information in an ontological knowledgebase. In: Proceedings of the 8th Linked Data in Architecture and Construction Workshop. https://ceur-ws.org/Vol-2636/06paper.pdf

Janowicz, K., Haller, A., Cox, S.J.D., Le Phuoc, D., Lefrançois, M., 2019. SOSA: A lightweight ontology for sensors, observations, samples, and actuators. Journal of Web Semantics 56, 1–10. https://doi.org/10.1016/j.websem.2018.06.003

Kang, C. et al., 2025. Intelligente digitale Methoden zur Verlängerung der Nutzungsdauer der Nibelungenbrücke. Beton- und Stahlbetonbau, best.70070. https://doi.org/10.1002/best.70070

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

Preidel, C., Borrmann, A., Mattern, H., König, M., Schapke, S.-E., 2018. Common Data Environment. In: Borrmann, A., König, M., Koch, C., Beetz, J. (Eds.), Building Information Modeling: Technology Foundations and Industry Practice. Springer International Publishing, pp. 279–291. https://doi.org/10.1007/978-3-319-92862-3_15

Rasmussen, M.H., Lefrançois, M., Schneider, G.F., Pauwels, P., 2020. BOT: The building topology ontology of the W3C linked building data group. Semantic Web 12(1), 143–161. https://doi.org/10.3233/SW-200385

Rolke, F., Becks, H., Göbels, A., Schulz, O., Beetz, J., 2026. spot-sensor-dataset. Design Computation, RWTH Aachen University. https://github.com/Design-Computation-RWTH/spot-sensor-dataset

Törmä, S., 2013. Semantic Linking of Building Information Models. In: 2013 IEEE Seventh International Conference on Semantic Computing, pp. 412–419. https://doi.org/10.1109/ICSC.2013.80

Törmä, S., 2014. Web of building data – integrating IFC with the Web of Data. eWork and eBusiness in Architecture, Engineering and Construction: ECPPM, pp. 141–147.

Valinejadshoubi, M., Bagchi, A., Moselhi, O., 2017. Managing Structural Health Monitoring Data Using Building Information Modelling. In: Proceedings of the 4th International Conference on Smart Monitoring, Assessment and Rehabilitation of Civil Structures (SMAR 2017), Zurich, Switzerland.

World Wide Web Consortium (W3C), 2013. SPARQL 1.1 Query Language. W3C Recommendation. https://www.w3.org/TR/sparql11-query/

World Wide Web Consortium (W3C), 2014. RDF 1.1 Concepts and Abstract Syntax. W3C Recommendation. https://www.w3.org/TR/rdf11-concepts/

Zhang, C., Beetz, J., de Vries, B., 2018. BimSPARQL: Domain-specific functional SPARQL extensions for querying RDF building data. Semantic Web 9(6), 829–855. https://doi.org/10.3233/SW-180297