3D Semantic Digital Twins: Data Streams and Ontologies for Risk Prediction in Heritage Contexts

Matteo Codiglione [1,2], Fabio Remondino [1],

[1] 3D Optical Metrology (3DOM) unit, Bruno Kessler Foundation (FBK), Trento, Italy - (mcodiglione, remondino)@fbk.eu
[2] University of Trento, Department of Information Engineering and Computer Science (DISI)

Keywords: 3DOnt, Ontology, Digital Twin, Semantic, Sensor Data Streams, Graph, 3D Point Cloud

Abstract

Sensor data are core for real-time monitoring in heritage contexts, allowing to ground the synchronization between the real asset and its Digital Twin (DT). Such synchronization can then be leveraged to perform risk prediction with a time aware reasoning process, which can be crucial for heritage prevention. Nonetheless, sensor data by itself lack both geometrical and semantic information which can also be pivotal in spotting such incipient degradations. For this reason, mature DT approaches need to tie these three main sources of knowledge and to conduct temporal reasoning in a form which can consider them altogether. This paper extends the *3DOnt Framework*, originally designed to link geometries and semantics, by integrating sensor data streams and enabling time-aware reasoning over live updates, thus turning it into a *Semantic-Geometric Digital Twins* framework. The extension includes an expansion of the 3DOnt backbone ontology to manage sensor data alongside raster-encoded and other static information, to model sensors as individuals within the 3D Graph and to link them to the macro-objects they monitor. Furthermore, 3DOnt's reasoning processes are enhanced to exploit dynamic sensor inputs, incorporating the temporal persistence of risk factors in prediction tasks. This extension is evaluated on the 3D Graph of a portion of Centro Culturale Santa Chiara, in Trento (Italy). With historical records simulating dynamic sensor data streams concerning temperature and humidity, we integrate this numerical information with semantic-geometric one encoded in the 3D Graph (i.e., proximity to windows, proximity to floors) to perform risk prediction assessing the formation of dew, which is often an early warning for more severe issues like indoor leaching or biological growth. More information about the 3DOnt Framework is available at: https://3dom.fbk.eu/projects/3DOnt.

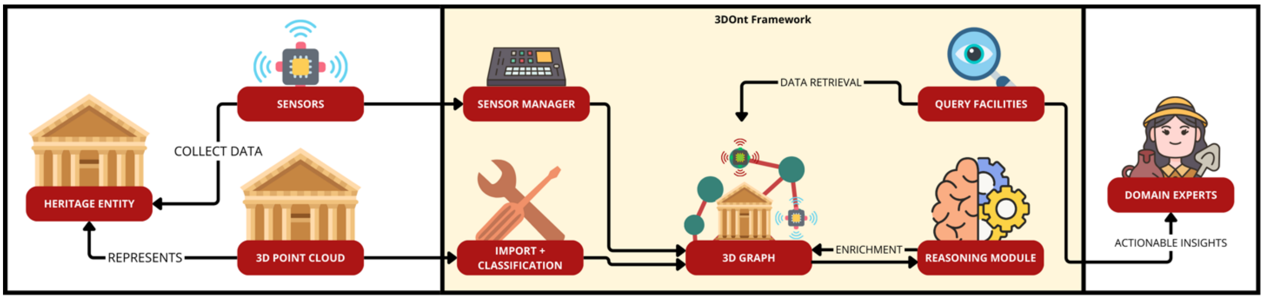

Figure 1: Graphical pipeline of the proposed extension of the 3DOnt Framework towards sensor-based digital twinning capabilities.

1. Introduction

In recent years, the concept of Digital Twins (DT) has evolved from a high-level metaphor to a practical framework widely applied across sectors like cultural heritage, urban monitoring, and infrastructure management (Tao et al., 2019; Fuller et al., 2020). DTs were originally conceived as dynamic virtual replicas of physical entities, updated in real-time through data streams and used for predictive diagnostics, simulation, and monitoring (Boschert Rosen, 2016). Following this track, early DT research focused on geometric fidelity, sensor synchronization, and simulation-based control. Limited attention has been paid to semantic richness, ontological structuring, or knowledge-based inference mechanisms (Boschert Rosen, 2016; Tao et al., 2019; Fuller et al., 2020; Attaran et al., 2023). Nonetheless, current research is incrementally shifting toward those initially ignored topics (Akroyd et al., 2021; Li et al., 2021; He et al., 2021; Zhiheng et al., 2022; Zheng et al., 2024). Indeed, semantic modelling plays a crucial role in enabling DTs to go beyond geometric mirroring and simulation-based control. By embedding machine-readable knowledge structures, semantic modelling enables automated reasoning, interoperability, and context-aware data interpretation. This is crucial in domains like cultural heritage, where complex relationships between structural components, materials, and environmental conditions play a pivotal role. In such contexts, semantic enrichment facilitates advanced inspection tasks, risk assessment, and long-term preservation. Moreover, semantics are key to support high-level queries, linking heterogeneous data sources, and enabling dynamic adaptation of the DT structure to support user needs and real-time conditions. Ontologies and knowledge graphs offer the necessary abstraction to unify data streams, geometric representations, and expert knowledge within a coherent, extensible framework.

We can identify four key requirements for a comprehensive Semantic Digital Twin framework which aims to provide effective and deployable support to heritage domain experts:

(i). semantic modeling,
(ii). fine-grained 3D spatial data,
(iii). real-time sensor integration, and
(iv). temporal reasoning.

Nonetheless, we highlight in Section 2 how few existing solutions properly integrate these features. While many systems address one or two, complete integration remains rare. Moreover, the application of semantic technologies often remains superficial, not exploiting their full potential for advanced reasoning and dynamic interaction. In this paper, we aim to extend the *3DOnt framework* (Codiglione et al., 2024a, 2024b) - which already addresses features (i) and (ii) - to cover the remaining aspects (iii) and (iv). Specifically, we propose a method that allows *3D Graphs* (Section 2) to model sensors as

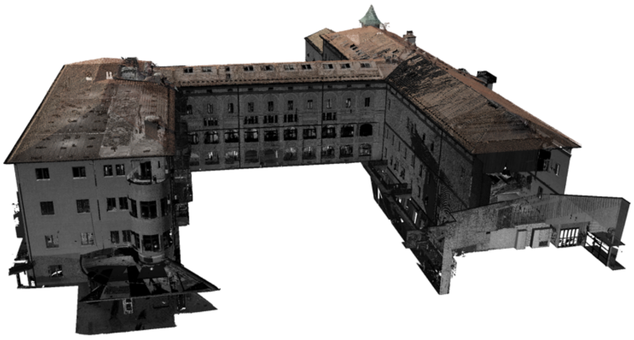

a)

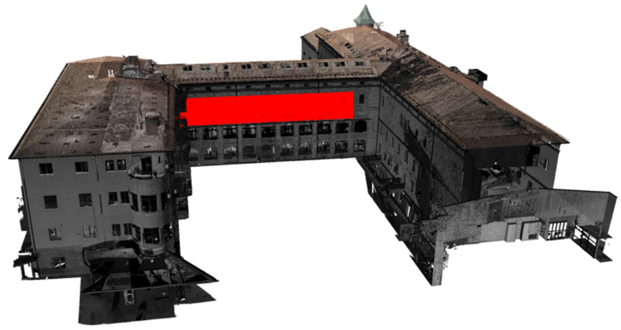

b)

Figure 2: Point cloud of Centro Culturale Santa Chiara in Trento, Italy. Original point cloud (a) and, in red, the portion employed in the present study (b).

ontological individuals, link them to monitored macro-objects, and reason over the temporal permanence of risk indicators. This enables the 3D Graph to serve as a fully semantic Digital Twin, fusing semantic-geometric representations with streamed numerical data for a multi-modal temporal reasoning (Figure 1). This integration is evaluated in a real-world scenario on the 3D dataset of the Centro Culturale Santa Chiara in Trento (Roman et al., 2023; Figure 2). Historic environmental data is used to simulate live sensor data streams and linked with 3D semantic structures to assess dew formation risk via spatial and temporal inference. Building on previous 3DOnt applications in heritage documentation and semantic spatial analytics, this work moves a step further - enabling 3D Graphs to set a new benchmark for *Semantic-Geometric Digital Twins.*

Related Works

Various approaches have been developed to integrate semantics into 3D models and DTs. Semantic modelling has initially been approached in the most basic form of semantic annotation of 3D models, which has long been a goal in digital heritage and architectural informatics (Poux et al., 2017; Grilli and Remondino, 2019; Murtiyoso and Grussenmeyer, 2019; Grilli and Remondino, 2020; Matrone et al., 2020; Teruggi et al. 2020; Moyano et al., 2021; Mazzacca et al. 2022). Nonetheless, this kind of semantic structuring is only able to encode for semantic information with the general logical shape of *This point/object belongs to this class*" (Codiglione et al., 2024a). Other approaches, also beyond the heritage context, aim for a more advanced semantic modelling. This is achieved by the application of formal ontologies (Beetz et al., 2009; Hiebel et al., 2017) and, in general, of semantic technologies, which offer a formal, method to describe 3D scenes. This include the description of spatial entities, their relationships, and associated metadata (Grüninger Fox, 1995). Nonetheless, this semantic scaffold is obtained by sacrificing the possibility to ground such semantics on fine-grained 3D data (e.g., point clouds). Together with this modelling-oriented usage, recent trends also highlight the increasing use of knowledge graphs and semantic web standards as data collectors for heterogeneous data sources, particularly real-time sensor streams (Steindl et al, 2021; Shi et al., 2023). Wang et al. (2022) enriched BIM models - natively addressing (i) and partially (ii) - with SensorML to support environmental monitoring - (iii). Zhiheng et al. (2022) proposed a dynamic spatial-temporal knowledge graph - (i) and (iii) - to support real-time reasoning in logistics - (iv). The Universal Digital Twin by Akroyd et al. (2021) models sensor updates through agent-driven dynamic graphs, tackling (i), (iii), and (iv). Similarly, Mandal O'Connor (2024) introduced a microservice architecture that dynamically updates knowledge graphs with IoT data. Li et al. (2024) applied a hybrid graph-ontology approach to infrastructure DTs, showing how real-time status tracking can be aligned with domain ontologies. A relevant contribution by Xie et al. (2022) presents an ontology-based information management framework aimed at enabling Digital Twins through multi-source data integration. Their core focus is on requirement (iii): managing and aligning heterogeneous datasets - including BIM, sensor, and operational data - within a unified structure. Ontologies are used as an integration tool to provide interoperability across data sources, but the semantic layer remains primarily structural and organizational. Such layer is never properly dedicated to domain-knowledge-driven modelling or semantic reasoning. As such, their use of semantic technologies supports data unification more than semantic enrichment or formal inference, and requirement (i) is only superficially addressed. Regarding 3D spatial detail - (ii) - the framework interfaces with BIM data, offering macro-object spatial modelling but without emphasis on fine-grained geometries or per-point semantics. Temporal reasoning (iv) is not part of their scope: while data streams may include time-stamped inputs, the framework does not perform reasoning over temporal dynamics, thus not leveraging them for knowledge extraction.

In general, as noted by Listl et al. (2024), most of semantic DT efforts do not fully leverage the Semantic Web Stack (Berners-Lee et al., 2001), the set of semantic technologies comprising ontologies and knowledge graphs. This results in such approaches exploiting only a limited portion of its capabilities. Furthermore, while the presented approaches effectively tackle (i), (iii), and often (iv), they tend to operate either in abstract, non-geometric spaces or limit their 3D spatial resolution to the macro-object level, failing to properly address (ii). No current solution deeply integrates all four features. In general, they do not support hybrid capabilities such as explicit 3D spatial semantics combined with ontology-based temporal reasoning. Combining these two latter considerations, we can argue that the current state of the art lacks integration at two different levels. The first level is articulated within (i) itself, since semantic technologies such as ontologies, knowledge graphs, and reasoning and query mechanisms are often used in isolation without mutual reinforcement. The second level occurs between (i) and the other core features, with semantics rarely embedded into or interacting with fine-grained spatial data (ii), real-time sensor integration (iii), or temporal reasoning (iv) in a cohesive framework.

The 3DOnt framework (Codiglione et al., 2024a, 2024b), following this semantic-oriented interest, emerged as a response to the need for semantic expressivity in large-scale point cloud datasets. It introduced the 3D Graph, a data structure that integrates geometry, semantics and domain knowledge via RDF/OWL ontologies.

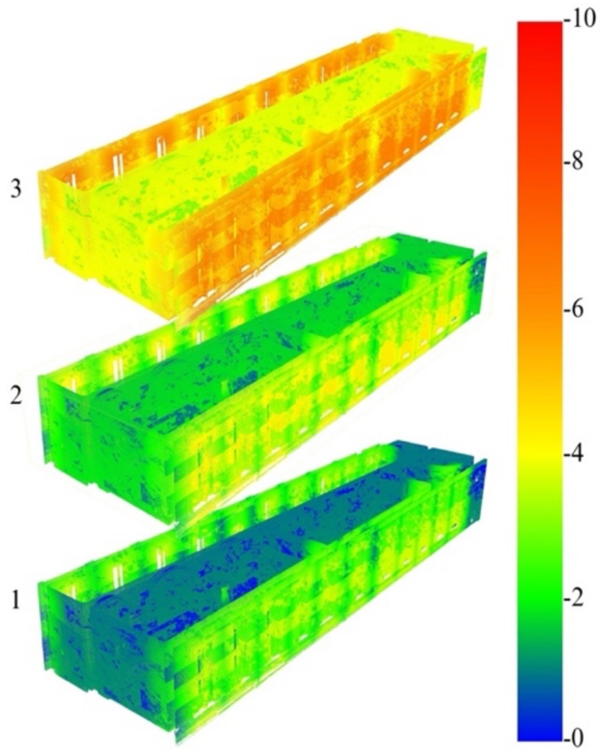

Figure 3: Visual Results of the dew formation risk index for the 3 analyzed data streams concerning temperature and humidity. 3D view.

The core idea is to leverage domain ontologies as a semantic backbone for a graph-based representation of 3D point clouds, which can thus natively leverage structured semantics and reasoning/querying possibilities. This graph-based representation is centered on modelling both points of the cloud and the macro-objects they constitute as individuals in a knowledge graph. Such graph also embeds their geometrical features (e.g., Coordinates and Verticality for points, Dimensions and Surface Value for the objects). Finally, the graph-based representation embeds the relationships tying together the point-level and the object-level (e.g., Constitutes, Is Constituted By) and the mereological - but not only - relationships intercurrent between objects (e.g., Is Part Of, Has As Part, Is On). In such a way the graph can also be visualized and analyzed like a 3D model (*Spatial Consistency*), yet presenting complex semantics which can be leveraged for both reasoning and querying. Furthermore, the 3DOnt framework offers advances inspection possibilities like SPARQL (Codiglione et al., 2024a) or natural language (Codiglione et al., 2025) visual querying possibilities. These tools aid semantically rich and user-friendly inspections of the 3D Graph, supporting its practical deployability.
Nonetheless, while the 3DOnt framework in its current setting effectively addresses features (i) and (ii), it only partially tackles (iv) - by performing multi-temporal analysis over static data. Moreover, it leaves (iii) completely uncovered, only focusing on static data integration.
With the proposed modifications, the 3DOnt framework is extended to also include the modelling of sensors and real-time data streams - thus accounting for (iii). Furthermore, reasoning processes are updated to account for the temporal dimension - (iv) - essential when dealing with dynamic input data streams to leverage their information value. Compared to the state of the art, the 3D Graph-based DTs show a unique combination of the four core features presented at the beginning of the section. In summary, these include semantic modelling (with full deployment of the Semantic Web Stack), 3D spatial data (with proper spatial consistency), real-time sensor integration, and temporal reasoning (Figure 3). Together, they allow for knowledge extraction and insight retrieval processes based both on geometric, semantic, and sensor information.

2. Experimental Materials

2.1 Technological Stack

Main technological components of the presented DT approach are the 3DOnt Framework, its signature data format (3D Graph) and the associated 3DOnt Ontologies, as described in section 2. Nonetheless, we had to modify some aspects of the 3DOnt Framework and its ontologies, as reported in Section 3.3 and in Section 3.2, respectively.
For what concerns the implementation of the framework we maintained the original usage of:

- Owlready2 (v0.46) (Owlready2, 2025) for the generation and management of the 3D Graph,
- Virtuoso with default configuration (Virtuoso Open-Source Edition, 2025) for 3D Graph storage and querying
- CloudCompare (CloudCompare, 2025) for visualization.

Everything has been conducted locally on a MacBook Pro with Apple M3 Max chip and 48GB of RAM.

2.2 Dataset

To evaluate this 3DOnt DT extension (Section 3.1) we employed a portion (10 million points) of the unclassified point cloud relative to the Centro Culturale Santa Chiara in Trento, Italy. Additionally, we employed historical records of temperature and humidity relative to a time span from 03/24 to 05/24 (Roman et al., 2023). Specifically, we extracted 3 portions of the historical sensor records concerning 3 6-hour long data streams (Table 1) with an update every minute.

Data Stream	1	2	3
Temporal Frame	8:00–14:00 [24/03/2024]	3:00–9:00 [15/05/2024]	8:30–14:30 [28/05/2024]
Min Temperature	16.80°C	22.69°C	19.50%
Max Temperature	17.28°C	25.93°C	22.62%
Mean Temperature	17.07°C	24.12°C	21.51%
Min Humidity	50.61%	57.89%	86.17%
Max Humidity	51.19%	70.21%	94.61%
Mean Humidity	50.91%	62.75%	89.53%

Table 1: Temporal and numerical specifics of the analyzed data streams.

3. Methodology

3.1 Extending the 3DOnt Ontology

First step in extending the 3Dont Framework towards a proper Digital Twin direction has been the extension of the backbone ontology. In the first place, we substituted the multiple domain-specific ontologies with a unique and multi-comprehensive one, which we called 3DOntCore. It displays the same structure described in Codiglione et al. (2024b) but it presents changes under the previously called *Acquisition* class branch, originally meant to only represent static rater data acquisitions. Such atomic branch has been replaced by the following class hierarchy:

Data Import
----*Static Import*
--------*Raster Import* (& other sibling classes)
----*Dynamic Import*
--------*Sensor Data Stream* (& other sibling classes)

Such a structure introduces into data collection the concept of resolution layer. This allows reasoning and querying processes to refer and to consider all imported data homogeneously, to discriminate between static and dynamic data imports or to discriminate even further and choose a specific source. Moreover, this allows measured properties (e.g., Temperature) to be easily modelled in an RDF reification setting as just having the Data Import superclass as their domain. This respects the intuition that temperature values measured by a raster or by a sensor, while differently characterized as static or dynamic, are both instances of the same *Temperature*.
Another structural ontological change regarded object relationships linking Data Import instances with the Macro Entities they describe. Relationships going from the formers to the latter are structured as follows (for each one, an inverse relationship is present):

Contains Data Relative To
----*Contains Static Data Relative To*
--------*Contains Static Original Data Relative To*
--------*Contains Static Aggregated Data Relative To*
--------*Contains Static Disaggregated Data Relative To*
----*Contains Dynamic Data Relative To*
--------*Contains Dynamic Original Data Relative To*
--------*Contains Dynamic Aggregated Data Relative To*
--------*Contains Dynamic Disaggregated Data Relative To*

The distinction between Original, Aggregated and Disaggregated data is meant to express the different source of data in three main cases. Respectively: (i) data directly collected from a sensor, (ii) data which has been inferred following a part → whole direction and (iii) data which has been inferred following a whole → part direction. An example for (ii) is provided by the temperature of a room inferred by the values of multiple sensors placed over different walls. Similarly, an example for (iii) comes from the temperature of a certain point inferred from the temperature of the room and, optionally, from other factors such as its height.
Each Data Import instance is then described by an Import Time and an Import Tag (Data Properties) for metadata storage.

3.2 Extending the 3DOnt Framework

To allow for a seamless management of sensor data streams, also the 3DOnt Framework functionalities needed an update. Firstly, we introduced the possibility to add a sensor in the 3D Graph by only specifying its name, the properties it measures, and the object (or objects) it describes. Then, we added an automatic update system which:

1. Retrieves the actual values from each sensor (we also developed a beta version managing real sensors and their connection. Nonetheless, in the present work only a simulated retrieval from the historical record has been effectively employed and tested).
2. Updates the ontologically encoded value of the concerned data stream.
3. Updates the relative historic, allowing for storage of historical records.

Given a user-specified number of hours, a temporal reasoning is run and the obtained information is stored back into the ontology.

3.3 Approach Evaluation: Dew Formation Risk Index

First step towards real-time, temporal reasoning-based monitoring of dew formation has been the generation of a 3D Graph moving from the unclassified point cloud. Initially, the point cloud is classified employing the method presented in Alami and Remondino (2024) for open-vocabulary 3D point cloud classification (integrated within the 3DOnt Framework's import facility).

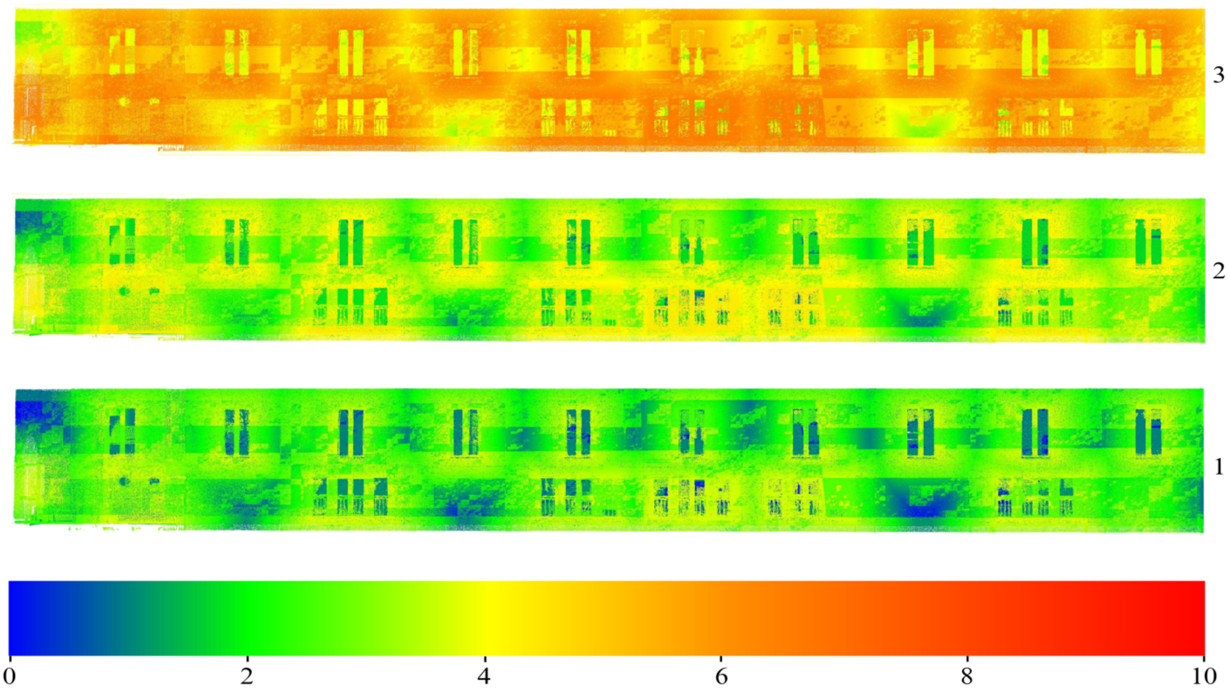

Figure 4: Visual Results of the dew formation index for the 3 analyzed data streams concerning temperature and humidity. View from the side.

https://doi.org/10.5194/isprs-archives-XLVIII-2-W12-2026-105-2026

This solution is based on performing a 2D semantic segmentation on images with known poses by mean of 2D vision models. The segmented images then get reprojected on the 3D model and an XGBoost-guided region growing process occur. Due to our lack of 2D data, we integrated the generation process with a 2D rendering step. At the end, the generation followed the following pipeline:

1. 2D Images extraction from the point cloud
2. Open-vocabulary classification
3. Instance Segmentation
4. 3D Graph Population
5. First Reasoning

In this specific case, we considered Window as the only class employed for segmentation, since it is the only one considered during dew risk assessment.

Then, we have inserted a Temperature and a Humidity sensor as individuals within the Sensor Data Stream class, simulating 3 real-time data streams with the extracted portions of the historic records.

For each temporal frame, after 6 hours – 360 updates – we ran the 3DOnt temporal reasoning process to predict dew formation. In designing the custom inferential rule, we considered the following factors as impactful in determining the final Dew Formation Risk Index by mean of a weighted sum. The weights have been defined by educated guesses for demonstration purposes.

Temperature and Humidity (weight: 0.60): Ambient temperature and relative humidity are the primary thermodynamic drivers of dew formation. Dew develops when the surface temperature of a material falls below the dew point of the surrounding air, which depends directly on both temperature and humidity. High relative humidity combined with air or surface temperatures near the dew point significantly increases the likelihood of condensation. In heritage buildings, where thermal insulation is often minimal and ventilation limited, these conditions can persist for extended periods - especially at night or in unheated spaces (ASHRAE, 2017). Therefore, the presence and persistence of dew-compatible conditions are the most critical factors in assessing condensation risk and receive the highest weight in the model. Temporal reasoning has only been considered in assessing the humidity and temperature contributions, given the static nature of the other factors. Being T the temperature and RH the humidity, we employed the following function to provide a smooth score between 0 and 1 for each set of temperature-humidity values (a set for each minute):

$$score(\delta) \begin{cases} min\,(1, 0.9 \cdot e^{-\left(\frac{\delta}{3}\right)^{0.8}} + 0.1 \cdot \left(1 - e^{\delta}\right) & if\ \delta \leq 0 \\ 0.9 \cdot e^{-\left(\frac{\delta}{3}\right)^{0.8}} & if\ 0 < \delta \leq 10 \\ 0 & if\ \delta \geq 10 \end{cases} \tag{1}$$

Where $\delta = T - T_d$, and T_d (the dew point) is calculated by mean of the Magnus-Tetens approximation:

$$T_d = \frac{b \cdot ln\left(\frac{RH}{100}\right) + \frac{a \cdot T}{b+T}}{a - ln\left(\frac{RH}{100}\right) + \frac{a \cdot T}{b+T}} \tag{2}$$

Where a = 17.62 and b = 243.12. The scoring function ranges between 0 and 1, assigning 0.9 to cases with $\delta = 0$. We then divided the sum of the scores by the number of considered minutes. In this way, the temporal persistence of critical factors is considered and represented in the index.

Proximity to windows (weight: 0.25): Proximity to windows is a strong architectural predictor of surface cooling and condensation risk. In traditional or poorly insulated buildings, window zones often act as thermal bridges, leading to lower interior surface temperatures - especially near frames, sills, or stained glass. Additionally, window zones are often subject to reduced air circulation due to framing or recess geometry, which limits moisture evaporation. For these reasons, proximity to a window is treated as a geometric-thermal risk amplifier in the condensation index. Proximity to windows is calculated geometrically leveraging the semantic classification.

Proximity to floors (weight 0.10): Lower wall regions (near the floor) are particularly susceptible to condensation for both thermal and fluid-dynamic reasons. Cold air stratifies near the ground, making lower surfaces more likely to cool below the dew point, especially in unheated or poorly ventilated rooms. In addition, these areas may receive less convective drying and can be affected by moisture rising from the subfloor or capillary action in porous materials. From a building physics perspective, floor proximity is thus a minor but reliable contributor to condensation likelihood and is incorporated as a supporting factor in the risk model (Hall Hoff, 2002). Proximity to floors is calculated geometrically on the basis of the Z coordinate of points.

Surface Verticality (weight 0.05): Surface orientation plays a role in condensation behaviour. Vertical surfaces, such as walls, may experience greater thermal bridging with exterior walls, causing them to cool more rapidly, especially if they are poorly insulated. Moreover, vertical wall surfaces often receive less direct convective drying compared to upward-facing horizontal surfaces. While verticality alone is not sufficient to cause condensation, it influences the surface's thermal behaviour and moisture retention, making it a meaningful geometric factor in the overall dew risk calculation.

After the weighted sum, the resulting index has been multiplied by 10 for better visualization. The index has been then attributed to every point of the 3D Graph as a data property called Dew Formation Risk Index. Finally, it has been showed (Fig. 3,4,5) directly on the 3D Graph, through the scalarfield-like 3D visualization possibility provided by the 3DOnt Framework (Codiglione et al., 2024b).

4. Results and discussion

Figures 3,4 and 5 showcase the per-point dew formation risk index. Beyond considering the single isolated factors, these figures present the intensity of the risk as a unified and distributed feature. Since only one temperature and one humidity sensor have been employed, in each moment all the values of such properties are equal across all the points. Thus, risk intensity differences within a single data stream-related image are entirely to be imputed to geometric (distances, verticality) and semantic (windows) features embedded in the 3D Graph. Notably, both verticality and proximity to windows, which we identified as supporting factors for the insurgence of dew formation, are strongly represented by the visual results. This confirms the ability of the approach to reflect and consider known conservation concerns. The comparison between visuals depicting different temporal frames shows instead the impact of sensor derived information, given the identity of geometric and semantic features across the various simulated temporal frames. The differences in the Dew Formation Risk Index across temporal frames correctly mirror the expected trends in dew formation, increasing as the humidity increases and the temperature decreases. These considerations highlight how 3D Graph-based DTs can meaningfully leverage both acquired geometric values, structured semantic content and directly measured properties. They can tie together data of different nature - static vs dynamic, numeric vs semantic - into a single, unified and temporal-aware new higher-level property. Such a property, encoded in the 3D Graph just as for the others, could then be used as a starting point for new inferences – e.g., for biological growth prediction. Considering the concerned case study, results show an overall low risk of dew formation, with average values ranging between 2 and 4 and peaks close to 6 in the highest-humidity day (data stream 3). Although this case study is more of a toy example than a real-world use case, its successful application to dew formation risk is nonetheless significant. It highlights the broader value of *3D graph-based digital twins*, particularly their flexibility in integrating heterogeneous data sources and supporting custom time-aware rules for user-specific tasks. Significantly, these tasks are performed while retaining spatial consistency which allows for 3D visualization and analysis. Achieved results also specifically showcase the new 3DOnt capabilities to ontologically model sensor data streams, to manage their updates and to perform a temporal reasoning for a time-aware knowledge extraction. With these modifications, the 3DOnt framework can successfully assess all the four features described in the Introduction which characterize a full exploitation of Semantic-Geometric DTs. This enforces the claim that the 3D Graph format is now more than a semantically structured, spatially consistent yet static representation of a real-world entity. Indeed, it can be rightfully defined as a Semantic-Geometric DT approach which can also undergo a temporal reasoning process.

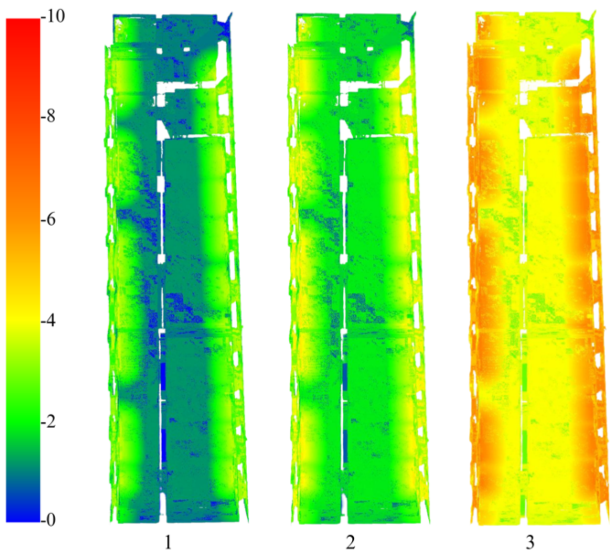

Figure 5: Top-view visualization of the results of the dew formation risk index for the analyzed data streams concerning temperature and humidity.

Moreover, all the features previously displayed by the 3DOnt Framework (Codiglione et al., 2024a; Codiglione et al., 2024b, Codiglione and Remondino, 2025) can still empower the DT management approach proposed by the present work. For instance, 3D graph-based DTs can be queried using both SPARQL or natural language, supporting user-driven inspections.

Before conclusion, there are still two tightly interleaved concerns about performances and scalability which need to be further discussed. We demand to previous work (Codiglione et al., 2024a; Codiglione et al., 2024b) for the discussion about the performances of the 3DOnt Framework in its generality. Instead, we hereby focus on the performances of the Owlready2-based process of the 3D Graph update with incoming sensor values and of the time-aware reasoning. Owlready2 works on an in-memory object representing the graph, so that, once such an object is generated, no performance issues arise in conducting standard operations on it. Sensor values updates, since only consist in the assignment of a Data Property value, fall within this *safe* category. The in-memory object generation can instead occur in two different ways: directly from a local file (*.ttl* or *.rdf*), and in this case the loading process is slow and not scalable (~ 15 minutes for a 3D Graph with 10 million points), or with a lazy-loading from a Virtuoso-stored 3D graph by mean of queries. In the latter approach an initial query generates an in-memory and Owlready2-compatible object comprising only points and colors for visualization. Then, when needed, additional queries are performed to extract from Virtuoso the desired information from the 3D Graph. This approach leverages Virtuoso's indexing capabilities, so that all the queries are sped up (~10 seconds for the initial query on a 3D Graph with 10 million points). Moreover, the initial – and more substantial - query only occurs upon 3DOnt Framework opening and does not need to be repeated anymore during the DT monitoring process. Possible approaches – depending on the use-case - could involve keeping the 3DOnt Framework running, if real-time notifications of risks are needed. Alternatively, sensor values could be collected in the historics and the framework could be opened only when time-aware inference is required.

For what concern the time-aware reasoning process it is not possible to provide a general performance evaluation, due to the possibility to freely model customized inferential rules. Each rule, according to its algorithmic structure, will involve different operations and as such will present a different computational cost. The inferential rule for the Dew Formation Risk Index, employed in the presented work, took ~30 minutes to compute the showed results. Nonetheless, we conducted no algorithm optimization, so that this performance could be improved. It is really a matter of use-case if the computational cost of a certain inferential rule is or not acceptable, depending on the urgency of the predictions. In the presented case, we consider 30 minutes to be appropriate for the prediction of the Dew Risk Formation Index in a heritage context. In general, the inference time only needs to be lower than the user-defined time interval between two subsequent inference processes.

The issue of scalability, by its side, is directly dependent of the size growth of the 3D Graphs as the original point clouds grow in points number. The employed section (10 million points) of the Santa Chiara point cloud presents a size of 717 MB as a point cloud (*.las*) and of 4,1 GB as a 3D Graph (*.ttl*). Assuming a number of macro-objects linearly dependent on the number of points (fixed resolution), and a fixed number of properties, the 3D Graph grows linearly with respect to the size of the cloud. The linear coefficient for this growth strongly depends on the number of additional imported data for each point. Despite 3D Graph format being a structurally bulky format, the linearity of this growth does not giving rise to particular scalability concerns. Moreover, while we conducted the described experiments in a local setting, the proposed methods could be applied by leveraging an online Virtuoso server instance. This would reduce the issue of the in-memory footprint of the 3D Graph format, which is by itself a memory-intensive data representation format. While a dedicated work could be done in quantitatively assessing the scalability of the approach as rule-complexity, graph size and/or sensor number grow, we can still hereby assess general deployability of the approach. We already showed how:

1. Rule-complexity can be adjusted freely according to case-specific hardware possibilities and task-required update frequencies.
2. Graph size and computational load scale linearly (with a factor dependant on the number of properties, relationships and data imports).

On the basis on such considerations, we can claim scalability-related concerns to be not critical nor urgent to assess in determining the validity of the proposed approach.

5. Conclusions

This work extends the capability of the 3DOnt Framework to evolve from a static semantic–geometric representation into a fully operational Digital Twin tool, able to integrate real-time sensor data streams and perform time-aware reasoning. By merging geometric, semantic, and environmental information into a single 3D Graph structure, the framework enables the extraction of higher-level, temporally dependent properties that can inform monitoring and predictive analyses. The case study exemplifies how semantic and geometric relationships can interact with sensor-derived values to generate meaningful spatial-temporal insights. Importantly, this holds even for limited sensing configurations, as the one considered in the present work. This confirms the potential of 3D Graph–based DTs to combine static and dynamic, numeric and symbolic data within a coherent reasoning environment. Notably, this feature-aggregation possibility sets this method apart from purely sensor-based approaches to alert systems. In the first place, the latter need to entirely rely on the available numerical data, without being able to suffice for the lack of some of them by leveraging expert or common knowledge. In our case study, for instance, we had no material information about windows, so that we could not quantitatively assess their actual heat dispersion metrics. A purely sensor-based approach would have had to just ignore this potential information. In contrary, our semantic-based approach could still partially leverage it in the form of an educated (and reasonable) guess that proximity to windows could still be considered as a relevant factor to carry on the risk prediction. In the second place, purely sensor-based approaches need to deal with a set of challenges which are instead natively solved by the proposed approach. These include:

1. Tracking the link between sensors and described objects/areas.
2. Considering in a properly integrated way both geometrical and measured properties.
3. Writing inferential systems which are more advanced than simple thresholding of sensor-measured values without being able to rely on a human-readable data scaffolding.

In the last place, as already mentioned, this approach locates itself in the broader context of the 3DOnt Framework, inheriting all its other functionalities. Using 3D Graphs as the base structure for digital twinning implies the compatibility of such DTs with the whole 3DOnt set of facilities. Overall, the results highlight the flexibility, expressiveness, and applicability of the proposed approach, particularly for cultural heritage contexts where interpretability, spatial consistency, and extensibility are often more critical than real-time performance.

References

Alami, A., Remondino, F., 2024. Querying 3D Point Clouds Exploiting Open-Vocabulary Semantic Segmentation of Images. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W8-2024, 1-7.

Akroyd, J., Mosbach, S., Bhave, A., Kraft, M., 2021. Universal Digital Twin - A Dynamic Knowledge Graph. *Data-Centric Engineering*, *2*.

American Society of Heating, Refrigerating and Air-Conditioning Engineers., 2017. *2017 ASHRAE Handbook: Fundamentals (Inch-pound Edition)*. ASHRAE.

Attaran, M., Celik, B., 2023. Digital Twin: Benefits, use cases, challenges, and opportunities. *Decision Analytics Journal*, *6*, 100165.

Beetz, J., van Leeuwen, J., Vries, de., 2009. IfcOWL: A case of transforming EXPRESS schemas into ontologies. *AI EDAM*, *23*, 89-101.

Berners-Lee, T., Hendler, J., Lassila, O., 2001. The semantic web. Scientific *American* 284(5), 28-37.

Boschert, S., Rosen, R., 2016. *Digital Twin - The Simulation Aspect*, pp. 59-74.

Codiglione, M., Beber, R., Remondino, F., 2024a. Leveraging Ontology for Enhanced Queries and Analyses of Urban Point Clouds. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *XLVIII-2/W8-2024*, 101-108.

Codiglione, M., Mazzacca, G., Remondino, F., 2024b. Dimensional Discoveries: Unveiling the Potential of 3D Heritage Point Clouds with a Robust Ontology Framework. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W4-2024, 119-125.

Codiglione, M., Remondino, F., 2025. NL-2-SPARQL: Ontology-Based Natural Language Querying over 3D Point Cloud Knowledge Graphs. Proc. *Digital Heritage*. The Eurographics Association.

CloudCompare, 2025. https://www.cloudcompare.org

Fuller, A., Fan, Z., Day, C., Barlow, C., 2020. Digital Twin: Enabling Technologies, Challenges and Open Research. *IEEE Access*, *8*, 108952–108971.

Grilli, E., Remondino, F., 2019. Classification of 3D Digital Heritage. *Remote Sensing*, *11*(7), 847.

Grilli, E., Remondino, F., 2020. Machine Learning Generalization across Different 3D Architectural Heritage. *ISPRS International Journal of Geo-Information*, *9*(6), 379.

Grüninger, M., Fox, M. S., 1995. Methodology for the design and evaluation of ontologies. *IJCAI Workshop on Basic Ontological Issues in Knowledge Sharing*.

Hall, C., Hoff, W. D., 2002. *Water Transport in Brick, Stone and Concrete* (1st ed.). CRC Press.

Hiebel, G., Doerr, M., Eide, Ø., 2017. CRMgeo: A spatiotemporal extension of CIDOC-CRM. *International Journal on Digital Libraries*, 18.

Li, T., Rui, Y., Zhu, H., Lu, L., Li, X., 2024. Comprehensive digital twin for infrastructure: A novel ontology and graph-based modelling paradigm. *Advanced Engineering Informatics*, 62, 102747.

Li, X., He, B., Wang, Z., Zhou, Y., Li, G., Jiang, R., 2021. Semantic Enhanced Digital Twin System for Robot-Environment Interaction Monitoring. *IEEE Transactions on Instrumentation and Measurement*, 1-1.

Listl, F., G., Dittler, D., Hildebrandt, G., Stegmaier, V., Jazdi, N., Weyrich, M., 2024. Knowledge Graphs in the Digital Twin: A Systematic Literature Review About the Combination of Semantic Technologies and Simulation in Industrial Automation. *IEEE Access*, vol. 12, pp. 187828-187843,

Mandal, S., O'Connor, N. E., 2024. OntoSmartDCU: A Multi-Agent Microservice Based Automated, Dynamic Ontology and Knowledge Graph Creation Framework From Real-Time Sensor Data for the Smart DCU Digital Twin. *IEEE 10th World Forum on Internet of Things (WF-IoT)*, 784-789.

Matrone, F., Grilli, E., Martini, M., Paolanti, M., Pierdicca, R., Remondino, F., 2020. Comparing Machine and Deep Learning Methods for Large 3D Heritage Semantic Segmentation. *ISPRS International Journal of Geo-Information*, *9*, 535.

Mazzacca, G., Grilli, E., Cirigliano, G. P., Remondino, F., Campana, S., 2022. Seeing among Foliage with LIDAR and Machine Learning: Towards a Transferable Archaeological Pipeline. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *46*, 365-372.

Moyano, J., León, J., Nieto-Julián, J. E., Bruno, S., 2021. Semantic Interpretation of Architectural and Archaeological Geometries: Point Cloud Segmentation for HBIM Parameterisation. *Automation in Construction*, *130*, 103856.

Murtiyoso, A., Grussenmeyer, P., 2019. Point Cloud Segmentation and Semantic Annotation Aided by GIS Data for Heritage Complexes. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *XLII-2/W9*, 523–528.

Owlready2, 2025. https://owlready2.readthedocs.io/en/v0.48/

Poux, F., Neuville, R., Van Wersch, L., Nys, G.-A., Billen, R., 2017. 3D Point Clouds in Archaeology: Advances in Acquisition, Processing and Knowledge Integration Applied to Quasi-Planar Objects. *Geosciences*, *7*(4), 96.

Roman, O., Farella, E. M., Rigon, S., Remondino, F., Ricciuti, S., Viesi, D., 2023. From 3D surveying data to BIM to BEM: the Incube dataset. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *XLVIII-1/W3-2023*, 175-182.

Shi, J., Pan, Z., Jiang, L., Zhai, X., 2023. An ontology-based methodology to establish city information model of digital twin city by merging BIM, GIS and IoT. *Advanced Engineering Informatics*, *57*, 102114.

Steindl, G., 2021. *Digital Twinning for Industrial Energy Systems Utilizing Semantic Web Technologies* [Dissertation, Technische Universität Wien].

Tao, F., Zhang, H., Liu, A., Nee, A., 2019. Digital Twin in Industry: State-of-the-Art. *IEEE Transactions on Industrial Informatics*, *15*, 2405-2415.

Teruggi, S., Grilli, E., Russo, M., Fassi, F., Remondino, F., 2020. A Hierarchical Machine Learning Approach for Multi-Level and Multi-Resolution 3D Point Cloud Classification. *Remote Sensing*, *12*(16), 2598.

Virtuoso Open-Source Edition, 2025: https://vos.openlinksw.com/owiki/wiki/VOS

Wang, T., Gan, V., Hu, D., Liu, H., 2022. Digital twin-enabled built environment sensing and monitoring through semantic enrichment of BIM with SensorML. *Automation in Construction*, *144*, 104625.

Xie, X., Moretti, N., Merino, J., Chang, J. Y., Pauwels, P., Parlikad, A. K., 2022. Enabling building digital twin: Ontology-based information management framework for multi-source data integration. *IOP Conference Series: Earth and Environmental Science*, *1101*(9), 092010.

Zheng, Y., Barazi, A. A., Seppänen, O., Abou-Ibrahim, H., Görsch, C., 2025. Semantic digital twin framework for monitoring construction workflows. *Automation in Construction*, *176*, 106301.

Zhiheng, Z., Zhang, M., Chen, J., Qu, T., Huang, G. Q., 2022. Digital Twin-enabled Dynamic Spatial-temporal Knowledge Graph for Production Logistics Resource Allocation. *Computers Industrial Engineering*, *171*, 108454.