2026 European Conference on Computing in Construction
Corfu, Greece
July 12–15, 2026

A CROSS-DOMAIN FRAMEWORK FOR EXPLAINABLE STRUCTURAL INSPECTION

Vedant Girish Dalvi, Cornelius Preidel, and Simon Vilgertshofer
HM Munich University of Applied Sciences, Munich, Germany

Abstract

The digital transformation of structural inspection has led to the widespread use of image-based artificial intelligence for automated damage detection. Deep learning models, particularly convolutional neural networks, have demonstrated strong capability in identifying surface-level defects such as cracks from photographic data. However, although these systems provide pixel-level predictions, they lack mechanisms for semantic interpretation, cross-domain data integration, and structured knowledge representation. For deployment in civil engineering workflows, AI-based inspection must move beyond detection toward contextual reasoning and interoperable knowledge capturing.

Introduction

This paper presents a conceptual cross-domain framework that connects four layers of structural damage intelligence: (1) damage detection and quantification, (2) semantic damage captioning, (3) ontology-based data integration, and (4) knowledge graph population, as shown in Fig. 1. The proposed framework aims to bridge geometric, semantic, and relational representations of damage within digital construction environments.

Methodology

The proposed framework facilitates not only the detection and quantification of pixel-level damages and the explanation of contextual relationships in overlapping damage scenarios, but also the standardization of these damages through the modification of an existing ontology. This enables the population of a knowledge graph and, ultimately, the development of a digital twin, as discussed in the following sections.

Damage Detection and Quantification

The first layer focuses on multi-label damage detection from inspection images belonging to three different damage domains: concrete, steel, and glass. These domains differ in terms of damage type, causality, severity, complexity, and relationships. Custom multi-label semantic segmentation models are developed to detect damages across these domains, with domain-specific models later combined via deterministic orchestration or AI agents. Currently, models have been trained and successfully inferred on pre-existing concrete damage datasets like the DACL10k dataset introduced in Flotzinger et al. (2024) and new benchmark dataset created for this purpose, as described by Krefft et al. (2026).

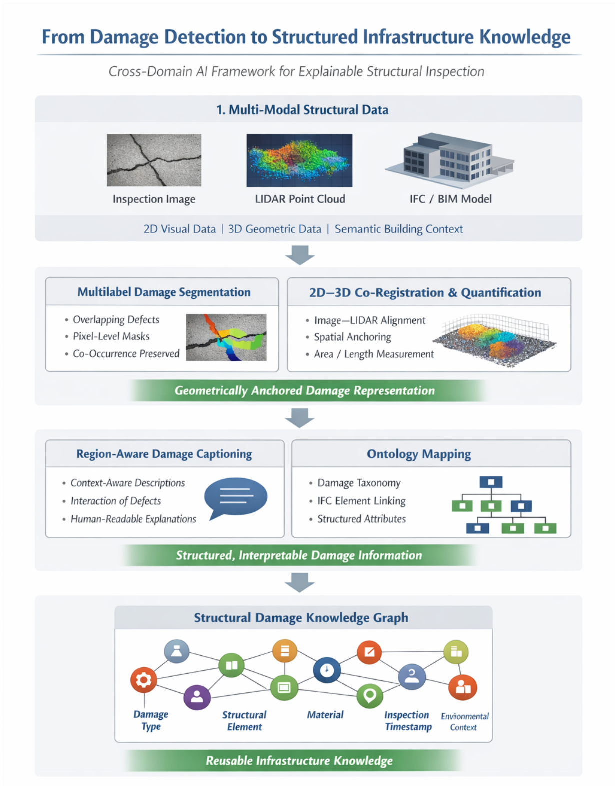

Figure 1: Proposed Methodology

Beyond 2D image analysis, the framework incorporates co-registration between image data and LiDAR-based geometric models. By aligning segmented damage regions with three-dimensional point clouds, damage localization can be anchored in spatially accurate structural coordinates. This integration enables the quantification of defect dimensions, and spatial orientation, as demonstrated by Krefft and Hoegner (2025). The result is a damage mask and parametric damage metadata, which can then be linked to a 3D model of the structure.

Semantic Damage Captioning

While detection provides spatial precision and quantification, engineering interpretation requires contextualization. The second layer introduces vision-language modeling for damage captioning. These captions can explain the overlapping patterns, and relative pixel positions among defects. Visual Language Models (VLMs) have been fine-tuned on a curated image captioning dataset to explain the context in overlapping concrete damages, and initial results have been presented by Dalvi et al. (2025). This layer enhances interoperability by transforming geometric predictions into human-readable contextual descriptions.

Ontology-based Data Integration

The third layer addresses the challenge of interoperability. AI outputs, whether masks or parametric data, must be aligned with construction-specific information standards to be usable. The proposed framework, therefore, incorporates ontology-based structuring of detected damages. The existing Damage Topology Ontology (DOT), proposed in Hamdan et al. (2019), is modified, and the detected damage categories are mapped to standardized taxonomies via ontology-aligned JSON templates and APIs and then linked to structural elements within digital building models, as described by Dalvi et al. (2026).

By structuring semantic outputs into ontologies, damage information becomes machine-readable and interoperable. This step enables alignment with Industry Foundation Classes (IFC)-based models and facilitates integration into a digital twin, a continuously coupled virtual representation of a physical structure that integrates geometry, damage observations, and structured knowledge to support inspection, reasoning, and long-term maintenance. The workflow assumes availability of a building model (e.g., IFC), but it can work with alternative structured data, such as classified point clouds or other spatial representations.

Knowledge Graph Population

The RDF (Resource Description Framework) graph explicitly models multiple relation types, representing information as a network of interconnected nodes and edges. These include spatial relationships (e.g., adjacency, overlap, containment), inspection timestamps, and hierarchical associations between damages and structural elements. Temporal evolution across inspections is represented by connecting damage instances over time, enabling tracking of progression, propagation, or mitigation effects.

In addition, the graph incorporates expert-defined causal relations that link observed damage patterns to potential causes, including environmental exposure, loading conditions, and material degradation mechanisms. Through this relational modeling, inspection data, digital building model information, and domain expertise converge into a unified structure. The resulting knowledge graph supports querying, and cross-project knowledge capturing, enabling analysis of recurring defect patterns, condition evolution trends, and intervention effectiveness.

Conclusion

This approach aims to transform visual damage detection into structured, explainable, and reusable infrastructure knowledge within digital construction workflows. Future work will focus on AI model orchestration, empirical validation of the integrated pipeline, refinement of ontology mappings, exploration of reasoning capabilities within the constructed knowledge graph, and development of a front-end application for visualizing all the detected damages and attributes.

Acknowledgments

The authors gratefully acknowledge the support of the Nemetschek Innovation Foundation as part of the AInspect2Twin project, which enabled our research and its practical application.

References

Dalvi, V. G., Martin, J., and Weilbach-Eyüboglu, T. (2025). Image captioning for building damage using deep learning: Explaining context in overlapping structural defects. In Proceedings of the 36th Forum Bauinformatik 2025, pages 349–357, Aachen, Germany.

Dalvi, V. G., Preidel, C., Fröhlich, D., and Dauberschmidt, C. (2026). A semantic workflow linking image-based concrete damage segmentation to ontology-driven knowledge graphs for digital twin integration and automated structural diagnostics. Manuscript in preparation.

Flotzinger, J., Rösch, P., Benz, C., Ahmad, M., Cankaya, M., Mayer, H., Rodehorst, V., Oswald, N., and Braml, T. (2024). DACL challenge: Semantic segmentation during visual bridge inspections. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision Workshops, pages 716–725.

Hamdan, A.-H., Bonduel, M., and Scherer, R. (2019). An ontological model for the representation of damage to constructions. In Proceedings of the 7th Linked Data in Architecture and Construction Workshop (LDAC 2019), volume 2389 of CEUR Workshop Proceedings, pages 64–77, Lisbon, Portugal.

Krefft, L., Fröhlich, D., Hoegner, L., and Dauberschmidt, C. (2026). Crackquant: Multimodal dataset for the detection and quantification of cracks in concrete using image and spatial data. Manuscript in preparation.

Krefft, L. and Hoegner, L. (2025). A method for crack detection and quantification in masonry using neural network-based image analysis. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, X-4/W6-2025:113–119.