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

ONTOLOGY-BASED SEMANTIC MANAGEMENT OF BRIDGE DRAWINGS

Xiaoli Song[1], Mengyan Peng[1], Steffen Marx[1], and Chongjie Kang[1]
[1]TUD Dresden University of Technology, Dresden, Germany

Abstract

As one of the primary information carriers throughout the bridge lifecycle, drawings contain rich semantic information. Traditional file-based management does not support efficient organization and retrieval. In this work, an ontology-based semantic management method for bridge drawings is proposed, leveraging information extracted by artificial intelligence (AI) models. After processing, AI outputs are transformed into structured semantic entities and relationships defined in the ontology. The resulting knowledge graph enables semantic search and helps engineers trace how structural components appear and evolve across drawings. It enhances the accessibility and reusability of bridge drawings and provides a practical foundation for intelligent asset management.

Introduction

Throughout the lifecycle of bridge engineering projects, a large number of heterogeneous files are generated (Kang et al., 2025). As a key part of this technical documentation, bridge drawings support all major stages, including design, construction, and maintenance. They carry essential engineering knowledge and form the basis for design decisions, construction guidance, and maintenance management.

As bridge projects continue to grow in scale and increase in complexity, a single project can generate a massive volume of drawings (Yang et al., 2022). However, traditional file-system-based drawing management treats each drawing as an independent file and organizes them through folder hierarchies and naming conventions. Although this approach is simple and intuitive, it has clear limitations when applied to large, heterogeneous drawing collections. Drawings from different project stages are sometimes stored in separate locations, even when they depict the same component at different points in time. As a result, engineers cannot efficiently track a component across views and stages, and the drawing content remains fragmented, lacking a systematic and explicit representation.

In recent years, AI has made significant progress in extracting information from engineering drawings (Moreno-García et al., 2018). AI models can identify component outlines, text, and symbols, providing a solid foundation for drawing digitization. However, most current approaches process individual drawings in isolation and cannot structure data across multiple project phases in a systematic manner. This limitation is particularly relevant in bridge engineering, where drawings are continuously produced throughout design, construction, and maintenance stages. Enabling cross-phase component traceability can therefore improve drawing reuse, inspection documentation, and lifecycle-oriented infrastructure management.

An ontology-based semantic management approach for bridge drawings can help address this gap and support the digitalization of bridge engineering. It enables semantic, content-based retrieval without relying on file names or folder paths, thereby improving accessibility. Explicitly modeling relationships enables the system to track the appearance of components across drawings from different project stages, thereby enhancing traceability for inspection and maintenance. Moreover, it can benefit Building Information Modeling (BIM) by supporting the reconstruction of BIM models from drawings, and by enabling model-to-drawing trace-back and collaborative viewing, thereby increasing the reuse value of drawing files.

Overall, this paper introduces a semantic pipeline that organizes bridge drawing information in a unified way. It further provides a new way of storing and organizing drawing content from the perspective of views. A case study is conducted to demonstrate how this pipeline can enhance the retrieval and traceability of drawing information. This work lays a foundation for ontology-based digital management and reuse of bridge engineering file assets.

Related Work

AI-based information extraction from engineering drawings

To support the digitalization of engineering drawings, various information extraction methods have been proposed, which can broadly be divided into traditional image processing approaches and deep learning–based methods. Early studies focused on text extraction and symbol recognition using techniques such as connected component analysis, Hough transforms, and rule-based databases (Fletcher and Kasturi, 2002; Or et al., 2005). These methods were later applied to interpret 2D drawings for 3D model generation, for example, through line detection and connectivity rules (Lewis and Séquin, 1998). In such approaches, geometric primitives extracted from 2D drawings are mapped to 3D elements based on predefined reconstruction rules, which often require strong assumptions about drawing conventions and completeness.

With the advent of deep learning, automatic feature learning has significantly improved the robustness and scalability of drawing interpretation. Recent studies have applied convolutional neural networks and related models for object detection, symbol classification, scene text detection, and OCR in technical drawings (Elyan et al., 2018; Vilgertshofer et al., 2020; Nguyen et al., 2021). These techniques enable the extraction of both geometric features and high-level semantic cues from 2D drawings, which are essential for bridging the gap between 2D representations and 3D models. Consequently, AI-based extraction is increasingly deployed to facilitate downstream applications, such as BIM and IFC-based bridge modeling. In these workflows, detected components, annotations, and spatial relationships serve as critical inputs for semi-automatic or rule-based 3D reconstruction (Lu et al., 2020; Akanbi and Zhang, 2022; Feist et al., 2024).

However, most existing studies focus on extracting information within individual drawings, while the organization and semantic management of extracted results across multiple drawings remain insufficiently addressed.

Semantic Management of Engineering Information

To manage the massive, fragmented data generated by modern projects, the concept of semantic file systems provides a crucial theoretical foundation (Gifford et al., 1991). Rather than treating files merely as "items stored under a rigid path," subsequent semantic file systems incorporated ontologies to model files as explicit objects with semantic metadata, enabling concept-based associative retrieval and reasoning (Ngo et al., 2007; Mashwani and Khusro, 2018). This ontology-driven paradigm is essential for organizing heterogeneous files.

Product Lifecycle Management (PLM) and Product Data Management (PDM) systems are used in mechanical engineering to manage data across the entire lifecycle of standardized, mass-produced items (Ji and Abdoli, 2023). However, they cannot be directly applied to bridge engineering because construction projects are unique, site-specific, and non-standardized, whereas manufacturing typically involves standardized products (Zhang et al., 2024). In the Architecture, Engineering, and Construction (AEC) industry, ontologies are more extensively deployed as a unified semantic layer (Noy et al., 2001; Elshani et al., 2025). Specifically, they are utilized to enhance information integration (Wu et al., 2021) and achieve interoperability across different systems (Hagedorn et al., 2023; Lorvão Antunes et al., 2024). These are particularly beneficial for the growing adoption of infrastructure digital twins (DT) (Zinke et al., 2023; Ramonell et al., 2023).

However, these macro-level approaches often treat engineering drawings as "black-box" attachments. By focusing on physical assets and broad business processes, they overlook the rich geometric and spatial semantics embedded within the drawings. To address this, specialized ontologies have emerged. Early work explored semantic drawing classification (Lu et al., 2008), and recently, Schönfelder and König (2025) proposed an ontology to standardize AI-extracted visual information, successfully introducing a semantic layer for individual drawings.

While such drawing-specific ontologies enhance the queryability of standalone files, they still lack the relational mechanisms required for complex, multi-drawing projects. Specifically, existing research cannot differentiate view types or trace identical components across drawings. To bridge this gap, this paper proposes a semantic management ontology for bridge engineering drawings that supports unified view representation and cross-view component tracking in multi-phase projects.

Methodology

Semantic Processing Pipeline

As shown in Figure 1, the overall pipeline can be divided into three main parts.

First, a set of bridge drawings is processed by trained AI models to extract basic information. The models produce segmented subviews for each drawing and JSON files containing OCR text and other detected elements. These outputs serve as the input for subsequent semantic processing.

Second, a Bridge Engineering Drawing Ontology (BEDO) is constructed. It defines the classes and relationships needed to represent drawing semantics. Where possible, existing domain ontologies are reused to avoid redefining common concepts and to maintain compatibility with established standards. In parallel, a semantic encoding module processes the AI-generated outputs and organizes information such as view types, components, and scales into a better structured encoding JSON file. This file is then used to automatically populate the ABox of the ontology via scripts. On top of these assertions, a set of rules is applied to further enrich the semantics. Since AI models extract unstructured text without understanding project-specific organizational logic, users can define custom rules tailored to specific naming conventions and operational needs. Applying these rules on top of the initial assertions further enriches the semantics and explicitly expands the relational structure of the final knowledge graph.

Finally, in the query stage, the pipeline supports semantic retrieval and visual localization. Users can issue queries based on components, view types, scales, and their relationships, and the system retrieves related drawings and views from the drawing collection folder. For selected results, the corresponding subviews are highlighted directly on the original full drawing to support positioning. In this way, the pipeline turns semi-structured AI outputs into computable knowledge and visualizable query results.

AI-Based Extraction of Drawing Information

2D engineering drawings contain both geometric and non-geometric information, including graphical representations of structural views as well as textual annotations and title blocks. In this work, information extraction follows an

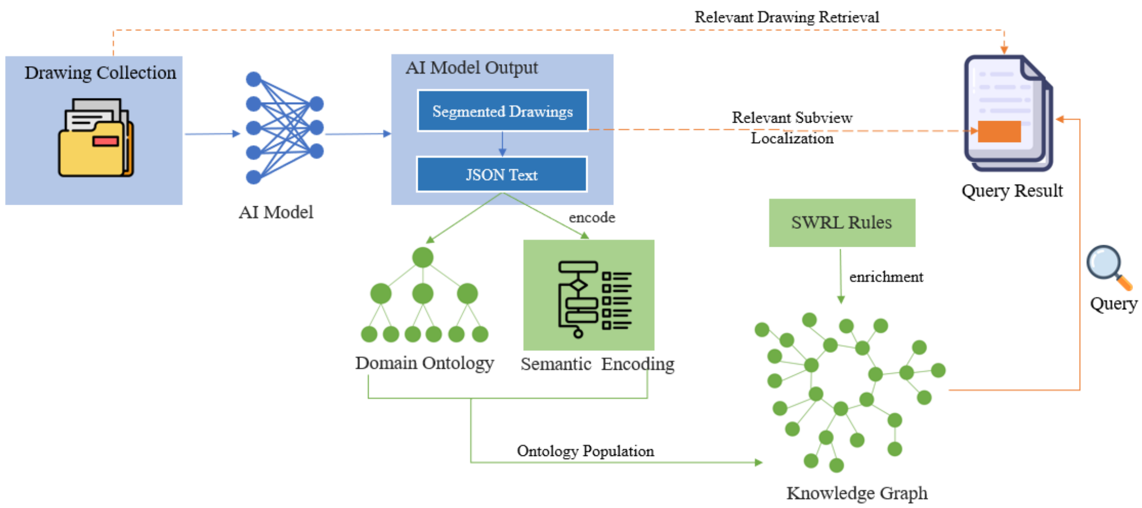

Figure 1: Overall Pipeline: from AI output to ontology-based query and visualization

integrated workflow that combines view localization, text recognition, and semantic post-processing, largely building upon the methodology proposed by Peng et al. (2025). This workflow enables the structured extraction of heterogeneous information from engineering drawings as a foundation for subsequent semantic modeling.

Graphical views in engineering drawings are localized using an AI-based object detection model. In this step, YOLOv7 is applied to identify view regions such as top and side views, providing bounding boxes that describe the spatial extent of each view. These localized view regions establish the geometric reference framework for organizing and associating extracted information. Textual information is obtained through OCR, which detects and recognizes text elements from annotations and title blocks. The OCR outputs include recognized text strings and precise bounding box coordinates, enabling spatial alignment between textual elements and graphical views within the drawings.

To transform the raw outputs of object detection and OCR into semantically meaningful information, post-processing is applied, including title-block metadata extraction, font-size-based view-title identification, and spatial matching between titles and detected views (Peng et al., 2025, 2026). In the current prototype, title-block metadata are extracted using an interactive semi-automatic workflow in which a reference region defined on the first sheet is transferred to subsequent drawings. The results are further filtered using clustering-based analysis and heuristic rules, while uncertain cases are retained for manual verification.

Furthermore, engineering drawings produced by different companies across various project phases exhibit heterogeneous layouts, terminology, varying scales, and geometric distortions from scanning. To handle these inconsistencies, the proposed workflow completely avoids fixed templates and absolute geometric dimensions. Instead, view localization relies on AI-based object detection using normalized bounding-box coordinates, which naturally mitigates the effects of geometric distortion. Similarly, text recognition and view-title identification are driven by relative font-size clustering and spatial relationships rather than predefined title-block positions or fixed pixel thresholds. Finally, explicitly stated scale information is extracted independently and encoded as a semantic attribute. This relative, feature-driven approach ensures that the extraction pipeline remains robust across highly diverse drawing collections.

However, the extracted outputs at this stage still lack a unified relational structure, necessitating further semantic encoding to align with the core ontology concepts.

Semantic Encoding

To transform the semi-structured AI outputs into ontology representations, a semantic encoding module is introduced. Based on the JSON files produced in the information extraction stage, the module identifies and normalizes elements such as view types, components, and scales. The result is an encoding JSON file that is ready for automatic ABox population in the subsequent step.

Figure 2 shows the encoding workflow. To enhance the robustness of the encoding method, the script implements three layers of matching. The matching conditions are progressively relaxed across these layers, so that exact matches are used whenever possible, while fuzzy matches can still be obtained when the available information is incomplete or noisy.

The first layer is exact matching, which uses direct string comparison with word-boundary checks and is suitable for standardized, error-free text. The second layer is substring matching, which relies on substring inclusion to handle compound words (e.g., German compound words). The third layer is fuzzy matching, which compensates for OCR and image segmentation errors, such as missing initial characters or spelling mistakes, by applying a similarity threshold (set to 90% by default).

These three layers are applied on top of the encoding keyword lists, ensuring that semantically relevant terms can still be detected even when the textual inputs deviate from

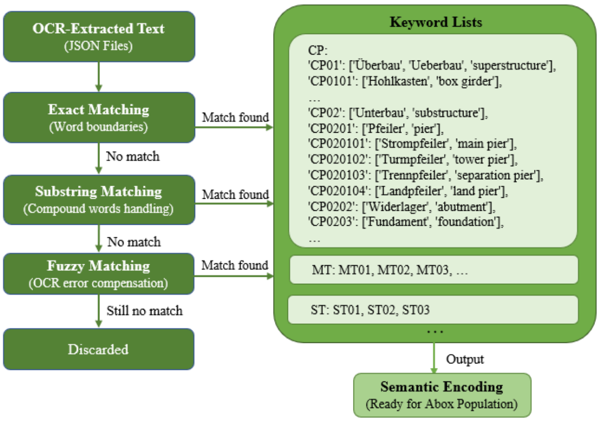

Figure 2: Semantic encoding workflow for OCR-extracted text

their canonical forms. Building on this, the script iterates through seven predefined keyword categories to generate semantic encodings for the extracted text: Bridge Type (BT), Component (CP), Document (DC), View (VW), Material (MT), Lifecycle Stage (ST), and Others (OT).

The keywords used in these categories were developed based on existing domain ontologies reused in BEDO, which are described in detail later in this paper. For each category code, the associated keyword lists include terms in multiple languages, such as German and English, allowing multilingual matching of textual outputs. Figure 2 shows an example encoding keyword list for the Component (CP) category. The design of the encoding keyword lists is extensible. By appending additional language variants to the keywords, the same encoding logic can be directly applied to drawings in new languages. Furthermore, new semantic categories can be added to the keyword lists as AI models become capable of extracting richer information from drawings. In this way, the semantic encoding can be incrementally extended while reusing the same overall pipeline.

Ontology Design

Figure 3 shows the overall structure of the proposed BEDO ontology. The ontology design consists of three main parts: classes, object properties, and data properties.

Classes

To describe engineering drawings, several core classes were designed. The `bedo:Drawing` class, which is a subclass of the `bedo:Document` class, represents technical drawings. `bedo:DrawingView` represents the corresponding views within a drawing. To ensure the defined ontology remains simple and practically applicable, established international standards DIN EN ISO 10209 (2022) and DIN EN ISO 128-3 (2022) were explicitly taken into account when defining the basic nomenclature of drawing views. However, because the definitions in these standards include many view types rarely encountered in the bridge domain, the view classes were selectively filtered and structured to incorporate only those practically utilized in bridge engineering. Additionally, the `bedo:DrawingSeries` class enables higher-level abstraction and organization of drawings when semantic information is more abundant.

The `bedo:Stage` class defines different phases of bridge engineering projects. Throughout a bridge's lifecycle, various types of drawings are produced at different stages, including design drawings, construction drawings, and maintenance drawings. These are represented by the `bedo:Design`, `bedo:Construction`, and `bedo:Maintenance` classes, respectively.

Object Properties

The object properties define the key relationships between classes. The `bedo:hasDrawingView` property links a drawing to its different views, as a single drawing can

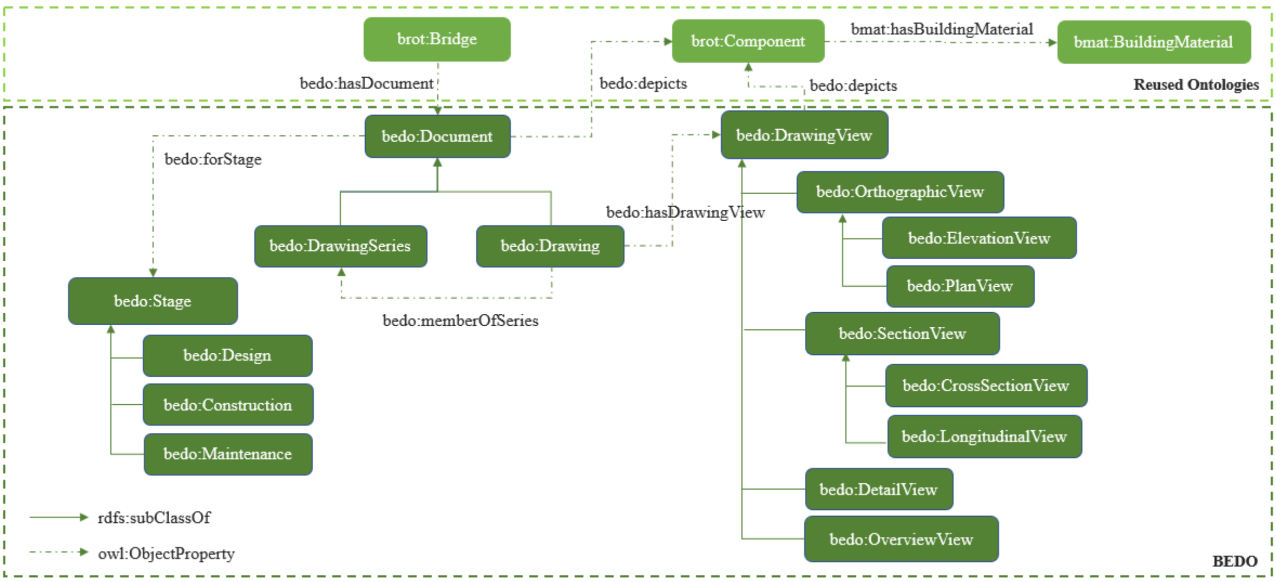

Figure 3: Bridge Engineering Drawing Ontology

contain multiple views. The `bedo:forStage` property describes which engineering stage a bridge document belongs to. The `bedo:memberOfSeries` property indicates that a drawing belongs to a specific drawing series. A single drawing can be included in multiple series based on different classification criteria. The corresponding inverse properties are also defined but are omitted here for brevity.

Data Properties

The data properties primarily capture attribute information from drawing files and processed model outputs. For drawing files, `bedo:drawingFilePath` specifies the file path location within the system, while `bedo:drawingTitle` stores the title from the drawing title block. For subviews generated after AI processing, `bedo:boundingBox` records the bounding box parameters used during AI-based object detection, which can be utilized for reverse view localization during queries. The `bedo:drawingViewTitle` property captures the titles of subviews and thus helps determine the semantic information of each view. Additionally, `bedo:viewFilePath` stores the file path of AI-generated view files. The `bedo:viewJsonFilePath` property records the file path of the OCR output JSON files. Other data properties include `bedo:scale` for scale information and `bedo:seriesType` for classifying drawing series types.

Reusing Existing Bridge Domain Ontologies

To improve interoperability and reusability, some existing bridge domain ontologies were reused. These include the Bridge Topology Ontology (brot), the Bridge Component Ontology (brcomp), and the Bridge Material Ontology (bmat) (Hamdan and Scherer, 2020). While BEDO models drawing- and view-level semantics, it relies on these external ontologies to represent the physical assets depicted in the drawings. Extracted components and materials are mapped directly to classes in `brcomp` and `bmat`, respectively, while topological relationships among components are aligned with `brot`. The reused ontologies and their IRIs are summarized in Table 1.

Table 1: Reused Ontologies in BEDO

Prefix	URI
brot	https://w3id.org/brot#
bmat	https://w3id.org/bmat#
brcomp	https://w3id.org/brcomp#

Application and Results

A case study was conducted using a set of inspection drawings of the Nibelungen Bridge. The Nibelungen Bridge is a historic prestressed concrete bridge in Worms, Germany, constructed in 1953 (Kang et al., 2024).

Ontology Population

The drawing dataset consists of 22 drawing files and was first processed using the model proposed by Peng et al. (2025). Note that the drawings used in this case study are not included in the original training dataset of that model. After semantic encoding and ABox population, a knowledge graph representing this drawing set was generated.

Rule-based Semantic Enrichment

Before being used for downstream tasks, the knowledge graph was further semantically enriched by applying a set of rules. As introduced in the methodology, these rules capture project-specific organizational logic. As shown in Figure 4, the inferred series instances, such as *Drawingseries_003* and *Drawingseries_Hessen*, are automatically linked to their corresponding member drawings, forming a more structured semantic network. Two illustrative examples are as follows:

- Rule 1 (Series by filename): Any drawing whose filename starts with ”003.” is assigned to the drawing series *Drawingseries_003*, where the *seriesType* is set to ”filename”.
- Rule 2 (Series by region): Any drawing whose drawing title contains the word ”Hessen” is assigned to the drawing series *Drawingseries_Hessen*, where the *seriesType* is set to ”region”.

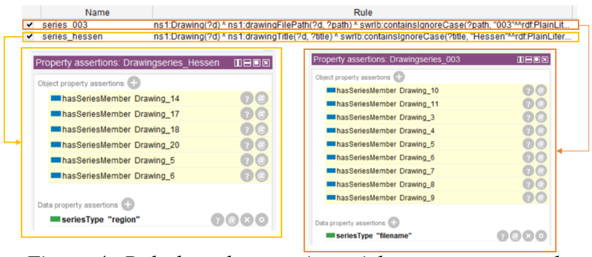

Figure 4: Rule-based semantic enrichment: two example drawing series defined by filename (right) and region (left)

Semantic Retrieval and Visualization

With the knowledge graph, users can perform semantic queries at different levels of granularity. For example, based on the aforementioned enrichment rules, users can directly retrieve all documents belonging to the *Drawingseries_Hessen*. Moreover, the system supports fine-grained, component-level queries, such as requesting all views related to ”pier” components. As shown in Figure 5, when querying for pier components, the system retrieves a list of matching subviews on the left panel, accompanied by their view titles. For each result, its exact location within the original drawing can be further visualized. In this specific example, the red mask effectively highlights the bounding box of the queried pier view.

Figure 6 shows the underlying knowledge graph of the drawing in Figure 5. The graph explicitly represents nine distinct drawing views and their containment relationships. By comparing the two figures, the mapping between visual representations and semantic instances be-

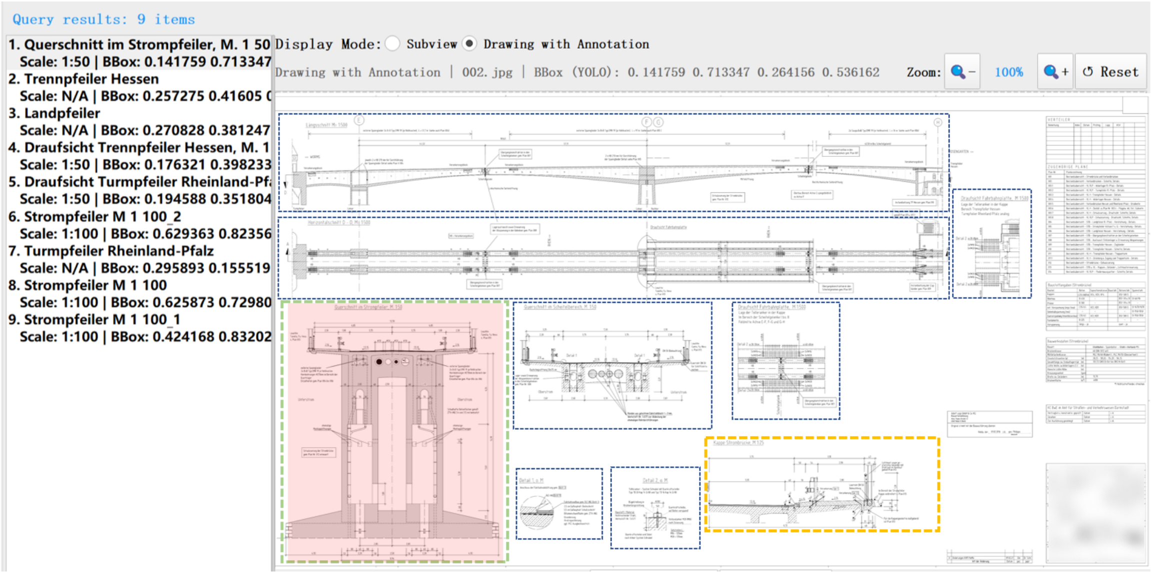

Figure 5: Query and localization of a pier subview within the original drawing

comes transparent. For example, the pier view enclosed in the green dashed box in Figure 5 corresponds to the *DrawingView_23* node. As highlighted by the green dashed box in Figure 6, this node is explicitly marked as a *CrossSectionView* depicting a *Pier* component. Similarly, the view enclosed in the yellow dashed box in the bottom right of Figure 5 maps to *DrawingView_18*, which depicts a *Cap* component. This direct linkage demonstrates how unstructured visual layouts are successfully transformed into a highly organized and queryable semantic network.

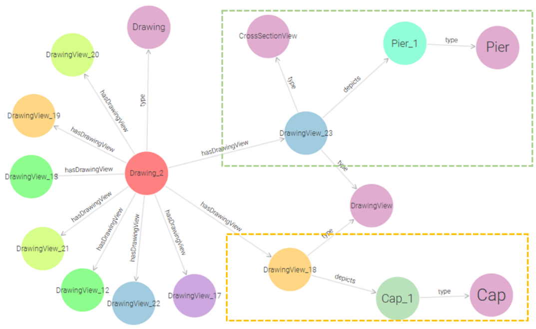

Figure 6: Queried drawing represented in the knowledge graph

Discussion

In the previous section, the proposed pipeline was demonstrated on a set of inspection drawings of the Nibelungen Bridge. The results show that, based on the semantics parsed from subview titles, components can be successfully traced back to their corresponding view locations in the original drawings. As long as the relevant information is explicitly encoded in the drawings, it can be recognized and linked through the ontology, enabling consistent cross-view and cross-stage traceability.

Although the case study uses German bridge drawings, the proposed semantic management approach is not limited to a specific national context. Since the ontology captures generic relationships among drawings, views, components, and project stages rather than country-specific regulations, it can be adapted to drawing collections produced under different standards and documentation conventions. Its applicability is also not constrained by bridge span size. Since BEDO operates at the semantic level of drawings, views, and component occurrences rather than on structural dimensions, it can in principle be applied to bridges with small, medium, or large spans, provided that the relevant information is represented in the drawings. Additionally, the pipeline shows potential for scalability, although this has not yet been validated on a large-scale dataset. Because the AI extraction processes individual drawings independently, the workload can be parallelized to handle much larger numbers of files. Furthermore, the generated semantic data consists of lightweight RDF triples, which modern graph databases can generally manage efficiently at large scale.

However, some limitations remain. In situations where view headings are missing due to overlapping elements, or where titles provide little semantic information on their own (e.g., “Schnitt A - A, M. 1:25”), the current pipeline struggles to assign complete semantics. In these cases, additional context from other parts of the drawing is required to interpret the meaning of the view, which is not yet systematically captured. Consequently, the semantic representation for such views remains incomplete. To further improve the richness of semantic information, additional extraction strategies are needed. These may include exploiting more diverse textual annotations, leveraging layout and graphical context, or incorporating prior domain knowledge to infer missing semantics beyond what is explicitly stated in the view titles. Furthermore, since a manually labeled benchmark dataset for systematically validating metadata extraction is currently unavailable, establishing one remains an important direction for future work.

Another limitation of the proposed pipeline is that the semantic accuracy of the populated ABox depends on the performance of the upstream AI models used for view detection and OCR. In practical bridge engineering projects, many drawings exist only as scanned documents with heterogeneous quality. Although the pipeline can be applied to such data, low resolution or incomplete scans may reduce extraction accuracy and lead to missing or incorrect semantic assertions. Nevertheless, the ontology-based representation itself remains independent of drawing format once sufficient information has been extracted.

The proposed approach is positioned not as an isolated solution for 2D document processing, but as a complementary mechanism for bridging unstructured legacy drawing data with modern BIM ecosystems. While new infrastructures are increasingly designed natively in 3D, existing assets still rely on vast archives of 2D engineering drawings. To ensure the proposed approach remains highly relevant in a BIM-dominated future, a critical avenue for future work is the direct integration of the proposed ontology with BIM frameworks. This can be achieved by establishing explicit semantic mappings between BEDO and IFC-based bridge models, such as IFCBridge. This integration would enable the transfer and alignment of information between legacy 2D drawings and 3D BIM models. Such a connection would support more comprehensive analyses across multiple data sources and facilitate downstream tasks such as lifecycle management, structural assessment, and digital-twin applications.

Conclusions

This paper presented BEDO, an ontology for representing bridge engineering drawings, together with a pipeline for automatically populating the ontology from the fragmented data produced by AI-based processing of the drawings. By parsing textual information and encoding it as ontology concepts, the approach links drawing content to a formal semantic model. Based on the instantiated ABox assertions, it then enables semantic querying of drawings at the view level. The case study on the Nibelungen Bridge demonstrated that, whenever sufficient semantic cues are explicitly encoded in the drawings, the proposed method can reliably trace component-related views back to their locations in the original documents. At the same time, several limitations remain, such as the incomplete extraction of relationships between views and the dependence on the quality of upstream AI components. Future work will address these limitations and further link BEDO with BIM models to support richer cross-modal analyses and more integrated lifecycle management of bridge infrastructures.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used Claude 4.5 and DeepL in order to improve the readability and language of the work. After using these services, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Acknowledgments

The authors gratefully acknowledge the financial support of the Deutsche Forschungsgemeinschaft (DFG) within the priority program SPP 2388 “Hundred plus – extending the service life of complex structures through intelligent digitalization” (project number 501804558) and of the Bundesministerium für Verkehr (BMV) within the project “HyBridGen – Hybrid Bridge Generator: AI-based bridge generator with knowledge and experience data and early public participation” (project number 01FV2067A).

References

Akanbi, T. and Zhang, J. (2022). Semi-automated generation of 3d bridge models from 2d pdf bridge drawings. In Construction Research Congress 2022, pages 1347–1354.

DIN EN ISO 10209 (2022). Technische produktdokumentation – vokabular – begriffe für technische zeichnungen, produktdefinition und verwandte dokumentation (ISO 10209:2022). Standard, Deutsches Institut für Normung.

DIN EN ISO 128-3 (2022). Technische produktdokumentation (TPD) – allgemeine grundlagen der darstellung – teil 3: Ansichten, schnitte und schnittansichten (ISO 128-3:2022). Standard, Deutsches Institut für Normung.

Elshani, D., Lombardi, A., Hernandez, D., Staab, S., Fisher, A., and Wortmann, T. (2025). Aec co-design workflow for cross-domain querying and reasoning using semantic web technologies. Automation in Construction, 176:106226.

Elyan, E., Garcia, C. M., and Jayne, C. (2018). Symbols classification in engineering drawings. In 2018 International joint conference on neural networks (IJCNN), pages 1–8. IEEE.

Feist, S., Jacques de Sousa, L., Sanhudo, L., and Poças Martins, J. (2024). Automatic reconstruction of 3d models from 2d drawings: A state-of-the-art review. Eng, 5(2):784–800.

Fletcher, L. A. and Kasturi, R. (2002). A robust algorithm for text string separation from mixed text/graphics images. IEEE transactions on pattern analysis and machine intelligence, 10(6):910–918.

Gifford, D. K., Jouvelot, P., Sheldon, M. A., and O’Toole, J. W. (1991). Semantic file systems. In Proceedings of the thirteenth ACM symposium on Operating systems principles, SOSP91, pages 16–25. ACM.

Hagedorn, P., Liu, L., König, M., Hajdin, R., Blumenfeld, T., Stöckner, M., Billmaier, M., Grossauer, K., and Gavin, K. (2023). Bim-enabled infrastructure asset management using information containers and semantic web. Journal of Computing in Civil Engineering, 37(1).

Hamdan, A.-H. and Scherer, R. J. (2020). Integration of bim-related bridge information in an ontological knowledgebase. In LDAC.

Ji, X. and Abdoli, S. (2023). Challenges and opportunities in product life cycle management in the context of industry 4.0. Procedia CIRP, 119:29–34.

Kang, C., Voigt, C., Eisermann, C., Kerkeni, N., Hegger, J., Hermann, W., Jackmuth, A., Marzahn, G., and Marx, S. (2024). Die nibelungenbrücke als pilotprojekt der digital unterstützten bauwerkserhaltung. Bautechnik, 101(2):76–86.

Kang, C., Walker, M., Bartels, J.-H., Marzahn, G., and Marx, S. (2025). Digital twin technologies for bridge lifecycle management—literature insights and a pilot study on the nibelungen bridge. Results in Engineering, 28:108288.

Lewis, R. and Séquin, C. (1998). Generation of 3d building models from 2d architectural plans. Computer-Aided Design, 30(10):765–779.

Lorvão Antunes, A., Barateiro, J., Marecos, V., Petrović, J., and Cardoso, E. (2024). Ontology-based bim-ams integration in european highways. Intelligent Systems with Applications, 22:200366.

Lu, Q., Chen, L., Li, S., and Pitt, M. (2020). Semi-automatic geometric digital twinning for existing buildings based on images and cad drawings. Automation in Construction, 115:103183.

Lu, T., guan, F., gu, N., and wang, F. (2008). Semantic classification and query of engineering drawings in the shipbuilding industry. International Journal of Production Research, 46(9):2471–2483.

Mashwani, S. R. and Khusro, S. (2018). The design and development of a semantic file system ontology. Engineering, Technology & Applied Science Research, 8(2):2827–2833.

Moreno-García, C. F., Elyan, E., and Jayne, C. (2018). New trends on digitisation of complex engineering drawings. Neural Computing and Applications, 31(6):1695–1712.

Ngo, H. B., Bac, C., Silber-Chaussumier, F., and Le, T. Q. (2007). Towards ontology-based semantic file systems. In 2007 IEEE International Conference on Research, Innovation and Vision for the Future, pages 8–13. IEEE.

Nguyen, M. T., Pham, V. L., Nguyen, C. C., and Nguyen, V. V. (2021). Object detection and text recognition in large-scale technical drawings. Proceedings of the 10th International Conference on Pattern Recognition Applications and Methods (ICPRAM).

Noy, N. F., McGuinness, D. L., et al. (2001). Ontology development 101: A guide to creating your first ontology.

Or, S.-h., Wong, K.-H., Yu, Y.-k., Chang, M. M.-y., and Kong, H. (2005). Highly automatic approach to architectural floorplan image understanding & model generation. Pattern Recognition, pages 25–32.

Peng, M., Cui, A., Kang, C., and Marx, S. (2025). AI-based extraction and management of text and view information from 2D bridge engineering drawings. University of Strathclyde Publishing.

Peng, M., Qian, H., Marx, S., and Kang, C. (2026). Optimizing 2d bridge engineering drawing digitization: A comparative study of text recognition tools and development of lightweight post-recognition structured information extraction methods. Results in Engineering, page 110186.

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.

Schönfelder, P. and König, M. (2025). Ontology-based reasoning in automatic floor plan analysis. Advanced Engineering Informatics, 68:103761.

Vilgertshofer, S., Stoitchkov, D., Borrmann, A., Menter, A., and Genc, C. (2020). Recognising railway infrastructure elements in videos and drawings using neural networks. Proceedings of the Institution of Civil Engineers-Smart Infrastructure and Construction, 172(1):19–33.

Wu, C., Wu, P., Wang, J., Jiang, R., Chen, M., and Wang, X. (2021). Ontological knowledge base for concrete bridge rehabilitation project management. Automation in Construction, 121:103428.

Yang, J., Xiang, F., Li, R., Zhang, L., Yang, X., Jiang, S., Zhang, H., Wang, D., and Liu, X. (2022). Intelligent bridge management via big data knowledge engineering. Automation in Construction, 135:104118.

Zhang, S., Tang, Y., Zou, Y., Yang, H., Chen, Y., and Liang, J. (2024). Optimization of architectural design and construction with integrated bim and plm methodologies. Scientific Reports, 14(1).

Zinke, T., Reymer, S., Kosse, S., Hagedorn, P., König, M., Wedel, F., Schneider, S., Marx, S., Nieborowski, S., and Windmann, S. (2023). Digital twins for bridges – concept of a modular digital twin based on the linked data approach, pages 1745–1752. CRC Press.