CONSTRUCTING CURATED KNOWLEDGE GRAPH STRUCTURES WITH AI-ASSISTED SEMANTIC PROCESSING

Alex Dekin[1,3], Thomas M. Froese[1,4], Wilma Leung[2,5] and Phalguni Mukhopadhyaya[1,6]
[1] Department of Civil Engineering, Faculty of Engineering and Computer Science,
University of Victoria, Canada
[2] B.C. Housing Research Centre, Burnaby, Canada
[3] alexdekin@uvic.ca
[4] froese@uvic.ca
[5] wileung@bchousing.org
[6] phalguni@uvic.ca

Abstract: Housing and construction research increasingly depends on transforming diverse documents, policies, technical standards and stakeholder inputs into structured knowledge that supports analysis and decision making. Knowledge graphs offer a principled representation of this information, yet their curation remains a bottleneck: converting unstructured text into ontology-aligned, provenance-aware graph structures is labour-intensive and difficult to scale. This study addresses this challenge by designing and evaluating an AI-assisted pipeline that constructs curated knowledge-graph fragments from real-world documents. Building on the Best Practices in Building Systems (BPiBS) project and its CIV (Collaborative Intelligence Vision) frameworks, the work defines target ontologies and exemplar subgraphs as ground truth. The pipeline integrates large language models for entity, relation and attribute extraction, followed by ontology-aware alignment, canonicalization and provenance tracking. Evaluation uses a design-science, mixed-methods approach combining quantitative metrics, entity and relation precision, recall, structural overlap and competency-question coverage, with expert review of semantic fidelity, ontology conformance and usefulness decision-support tasks. Human-in-the-loop checkpoints are examined to identify where limited reviewer input provides the highest value. The contributions are threefold: (1) a reproducible pipeline with versioned prompts, alignment strategies and export to a Neo4j environment; (2) an evaluation protocol and benchmark artifacts for curated graph construction in housing systems; and (3) guidance on when to rely on automated extraction versus targeted human intervention. By reducing the distance between narrative evidence and computable structure, the approach supports more transparent and updateable system maps and is generalizable to other civil-infrastructure domains that require translating heterogeneous text into high-fidelity, queryable knowledge.

1 Introduction

Housing research spans such broad and interconnected realms of social and technical systems, that it must make sense of diverse bodies of evidence and information: policy and technical standards, program documents, case studies, stakeholder insights and academic literature. Knowledge graphs offer a principled way to represent this complexity: entities and relationships can be made explicit, provenance can be retained and cross-domain patterns can be queried rather than guessed (Hogan et al. 2021; Ji et al. 2022). Yet, the practical bottleneck is curation. Turning messy text into a stable, computable graph that conforms to a domain ontology remains labour-intensive and difficult to scale (Wang et al. 2012; Ji and Grishman 2011; Suchanek et al. 2013; Jaradeh et al. 2023). This paper addresses that bottleneck for the housing and construction context by investigating how curated knowledge-graph structures can be constructed from real-world documents with the aid of large language models (LLMs) for entity/relation extraction, semantic labelling and ontology alignment.

The work has arisen from an ongoing research project called Best Practices in Building Systems (BPiBS), which undertook to weave together technical, social, environmental and governance information in a reliable, queryable representation to support roadmapping, evaluation, and policy learning. Manual curation of an integrated knowledge repository struggled to keep pace with evolving guidance, local practices and emergent innovations. Recent advances in LLMs makes semi-automation plausible, but unproven: information extraction often drifts semantically, under-specifies relationships, or breaks when confronted with domain-specific ontologies (Dagdelen et al. 2024). In short, the field lacks a tested, reproducible pipeline that translates domain texts into high-fidelity graph structures while preserving provenance and aligning to curated schemas.

This paper therefore takes up the following problem: **how to design and evaluate an LLM-assisted process that constructs curated knowledge-graph structures from diverse documents with sufficient semantic reliability to be useful in practice.** The central challenge is not simply recognizing mentions in text, but producing graph fragments (nodes, typed relations, attributes, and constraints) that conform to an evolving ontology and answer competency questions typical of the domain. Addressing this challenge has direct value for knowledge mobilization in housing systems: higher-quality graphs can connect lived experience to standards and policy and support cross-sector coordination.

The purpose of the study is applied and methodological: to design, implement, and assess a reproducible pipeline for AI-assisted semantic processing that yields graph-level outputs aligned to a curated ontology. Guiding research questions are framed as follows:

- To what extent can an LLM-assisted pipeline construct curated subgraphs from domain documents with acceptable fidelity at the entity, relation, and subgraph levels?
- Which prompting, filtering, and alignment strategies most improve ontology conformance while retaining provenance and minimizing manual edits?
- Under what conditions (document types, subdomains, or task types) does human-in-the-loop review measurably improve utility—as evidenced by competency-question coverage and task-based usefulness for roadmap/decision-support activities?

The results will be measured through a combination of quantitative and qualitative evaluations. Quantitative assessment will focus on precision and recall at the entity, relation, and subgraph levels, along with coverage of predefined competency questions. Complementing this, expert review will provide qualitative validation of semantic accuracy, ontology alignment, and practical usefulness for roadmap and decision-support activities.

The scope of this study is field-proximal yet computational. It focuses on constructing graph structures from textual materials into a working graph database environment while maintaining source provenance. Ethical considerations are observed where human-sourced materials are involved (e.g., de-identification of listening-session content and adherence to consent/REB requirements).

This study is timely. Stakeholders across the housing ecosystem increasingly expect transparent, navigable representations of how strategies, enabling activities, and evidence connect, yet they also need agility: the ability to ingest new documents and regenerate consistent views without starting from scratch. A validated pipeline that yields semantically disciplined graph fragments can shorten the distance between narrative evidence and structured knowledge, making it easier to explore alternatives, compare interventions, and connect policy to practice.

2 Research Paradigms

This study adopts a pragmatic, critical-realist paradigm suited to a socio-technical problem in which curated domain structures (ontologies, standards, practices) have real organizing properties, yet are only partially accessible through text and data. The research is applied and methodological, following a Design Science Research strategy focused on creating and evaluating an AI-assisted pipeline for entity and relation extraction, semantic alignment and ontology-conformant graph construction. Reasoning proceeds through an abductive–deductive–inductive cycle: hypothesizing representational solutions, testing them against explicit metrics and competency questions and generalizing lessons from iterative build–evaluate loops. A mixed-methods approach combines quantitative performance measures with qualitative expert assessments, using triangulated evidence across models, document types and reviewers. Conducted on real domain materials but within controlled computational conditions, the study emphasizes construct, external and reliability validity through versioned prompts, reproducible pipelines and ethical handling of human-sourced content. Its contributions include the design artifact itself, empirical evaluation evidence and prescriptive guidance for producing high-fidelity, semantically consistent knowledge graphs from heterogeneous textual sources.

3 Literature Review

Automated construction of knowledge graphs (KGs) from unstructured text has progressed through several methodological waves, from early rule-based extraction systems and probabilistic pipelines to contemporary, large language model (LLM), assisted frameworks. Foundational work on automated knowledge base construction (Wang et al., 2012; Suchanek et al., 2013) established enduring principles for scalable extraction, uncertainty representation, and human-in-the-loop validation. These studies argued that knowledge should not be treated as deterministic output but as probabilistic, evolving information requiring iterative refinement. Their emphasis on provenance, confidence estimation, and hybrid reasoning continues to shape current research concerned with balancing automation and human curation in complex domains.

Subsequent advances in deep learning reframed information extraction as a joint learning problem. The DeepKE framework (Zhang et al., 2023) exemplifies this shift by integrating named entity recognition, relation extraction, and attribute identification under a unified architecture that supports both standard and low-resource scenarios. This modularity aligns with design principles of transparency and replicability in research systems, offering reusable components that can be adapted across domains. Complementary developments in document-level relation extraction further extend this logic. (Yuan et al. 2024) introduced hierarchical tree–graph representations and relation segmentation mechanisms that explicitly model entity interactions across sentences, addressing a longstanding challenge of cross-sentence reasoning and contextual dependency in document-level extraction.

A parallel trajectory has emerged around leveraging LLMs to perform end-to-end structured information extraction. (Dagdelen et al. 2024) demonstrated that fine-tuned LLMs can convert scientific text directly into machine-readable JSON representations, substantially reducing the need for complex, multi-stage pipelines. Their results show that human evaluators often rate such outputs higher than traditional exact-match metrics would suggest, indicating that conventional evaluation methods may underestimate the real-world utility of LLM-based extraction. Similar findings appear across emerging surveys of document-level information extraction (Zheng et al., 2024), which emphasize the potential of generative models to unify entity recognition, relation identification, and reasoning under a single architecture, albeit with continuing issues around long-range coherence, label noise, and factual consistency.

Ontology alignment and schema evolution represent another frontier. Traditional approaches rely on lexical and structural matching, yet recent work explores LLMs for more complex correspondence discovery. (Amini et al. 2024) show that incorporating modularized ontology context enables GPT-4 to perform non-trivial, many-to-many alignments without shared instances. This shift from symbolic heuristics to contextual reasoning broadens possibilities for integrating heterogeneous data sources—a capability increasingly critical in domains with overlapping but non-identical taxonomies. (Zhang and Soh 2024) extend this paradigm through their Extract–Define–Canonicalize (EDC) framework, which introduces a three-stage LLM workflow for open knowledge graph construction. By separating extraction, schema definition, and canonicalization, the method preserves both flexibility and interpretability, avoiding prompt overload while ensuring terminological consistency.

The convergence of these strands indicates a transition toward hybrid architectures where LLMs function not as black boxes but as adaptive components embedded in transparent workflows. In biomedicine, (Cao et al. 2024) combine ontology-enhanced prompting with chain-of-thought reasoning to extract rare disease information, demonstrating that structured domain knowledge can constrain generative models and improve factual reliability. Collectively, such approaches suggest that high-performing KG systems will integrate symbolic ontologies, probabilistic reasoning, and LLM-based semantic processing, with measurable performance on accuracy, consistency, and human-editing efficiency.

Despite rapid progress, persistent gaps remain. Most existing systems optimize for extraction precision rather than curator effort or maintenance cost; few measure the stability of machine-generated assertions over iterative updates. Provenance tracking and explainability are still underdeveloped, limiting accountability in dynamic, multi-source environments. Moreover, transferability across domains is uncertain, models fine-tuned for one corpus often degrade when confronted with new linguistic registers or schema variations. Addressing these gaps requires methodological frameworks that integrate rigorous evaluation, transparent data governance, and continuous feedback between automated inference and expert review.

Building on these insights, current research seeks to operationalize LLM-assisted semantic processing within structured, design-driven workflows for knowledge graph construction. The objective is to align the generative capabilities of modern language models with the precision and traceability demanded by curated data infrastructures. By situating LLM extraction, ontology alignment, and post hoc canonicalization within reproducible experimental designs and well-defined measurement protocols, this work contributes to a broader shift, from static information extraction toward adaptive, curator-aware systems that evolve with both data and context.

4 Methods

This research employs a design science and mixed-method approach to develop, test, and evaluate an AI-assisted framework for constructing curated knowledge graph (KG) structures from domain-specific documents. The methodology integrates qualitative ontology engineering, automated language model processing and quantitative performance evaluation to ensure both conceptual validity and empirical rigor. The overall process follows an iterative build–evaluate cycle, in which system components are developed, tested against ground-truth subgraphs, refined based on expert review, and revalidated through competency-question coverage.

4.1 Research Design and Framework

The study adopts a system-oriented methodology grounded in three interconnected phases (Jaradeh et al. 2023):

- Knowledge Preparation and Ontology Contextualization – defining the data foundation and representational schema.
- AI-Assisted Construction Pipeline – extracting, aligning and generating graph data from unstructured text using language models.
- Evaluation and Validation – assessing construction accuracy, semantic fidelity and practical applicability.

This methodological structure directly supports the research objectives introduced in the paper: (1) to automate the construction of curated graph structures from textual sources, (2) to evaluate the semantic alignment of extracted entities and relations and (3) to measure the reliability and interpretive usefulness of AI-assisted graph representations.

4.2 Corpus and Ground-Truth Data

The corpus for experimentation consists of policy documents, codes and standards, research reports and listening-session transcripts derived from the BPiBS project. All files are converted to normalized UTF-8 text and .md format, with structural elements such as titles, headings and tables preserved as embedded markers (Utkucu et al. 2025). To support evaluation, we created small "ground-truth" examples of the graph, subgraphs that were manually checked by domain experts. These curated subgraphs, built in Neo4j, represent typical themes in the database such as enabling activities, governance processes, stakeholder roles and technical strategies. These subgraphs serve as both alignment targets and evaluation benchmarks, supporting quantitative measures of precision and recall as well as qualitative judgments of semantic relevance. The target schema follows a domain ontology modeled in Neo4j with class hierarchies, relation types and attribute (properties and key-value pairs) constraints.

4.3 AI-Assisted Construction Process

The proposed construction process transforms unstructured text into semantically aligned knowledge graphs through a series of automated and semi-automated steps. Each stage is designed to preserve data provenance, interpretability and alignment with the reference ontology.

4.3.1 Text Ingestion and Structuring

Documents are batch-imported, deduplicated and segmented by section markers. Each text unit is stored with metadata (document ID, name, description, source, publication date, etc.) and embedded for semantic retrieval.

4.3.2 Chunking and Retrieval

Text segments are divided into coherent chunks using structure-aware rules based on headings, sentence boundaries and metadata hierarchy. The system retrieves contextually relevant evidence packs to support focused extraction prompts.

4.3.3 Entity, Relation and Attribute Extraction

A large language model (LLM) is prompted with schema hints and few-shot examples to extract entities, relationships and key attributes in structured JSON format. Each output fragment is linked to the original source span to preserve interpretability and traceability.

4.3.4 Coreference and Normalization

Mentions of equivalent entities across multiple documents are merged through string similarity matching, ontology alias lists and embedding-based clustering. The process generates canonical labels while retaining surface forms for provenance.

4.3.5 Ontology-Aware Alignment

Extracted sematic information pieces are aligned to the domain ontology using a two-stage hybrid approach: lexical and embedding similarity for initial matching, followed by an LLM-assisted verification layer.

4.3.6 Constraint Checking and Graph Assembly

Shape rules and cardinality constraints are applied to validate graph consistency. The resulting information pieces are transferred into a Neo4j instance, producing a versioned graph with diff reports for iterative comparison across runs.

4.3.7 Human-in-the-Loop Validation

Curators review uncertain or conflicting entries through a web-based dashboard (UI). Relevant decisions (accept, modify, reject) feed back into refining prompt templates, ontology structures and extraction rules, creating an iterative learning loop between machine generation and expert, humans-in-the-loop correction.

4.4 Evaluation Strategy

The evaluation framework combines quantitative accuracy metrics with qualitative expert assessment to capture both structural and interpretive dimensions of graph quality.

- Quantitative Evaluation. Performance is measured against ground-truth subgraphs using entity- and relation-level precision, recall, and F1 scores, as well as subgraph overlap metrics such as graph edit distance and path recall. Competency-question (CQ) coverage evaluates whether the constructed graphs can answer domain-specific queries correctly. Ablation studies test variations in prompt design, schema context size and alignment strategy to determine their influence on model performance.

- Qualitative Evaluation. Subject-matter experts (knowledge graph curators) assess semantic fidelity, ontology compliance, and interpretive usefulness. Their feedback informs error categorization and subsequent methodological refinements.

- Baselines and Comparison. Results are benchmarked against three reference systems: (1) a rule-based extraction baseline, (2) an Open Information Extraction (OpenIE) model, and (3) an LLM-only pipeline without ontology alignment.

4.5 Reproducibility and Ethical Considerations

All scripts, prompts and schema files are version-controlled, and inference parameters are logged to ensure reproducibility. Random seeds are fixed, and each pipeline execution produces deterministic outputs. Ethical protocols are followed for human-derived documents such as listening-session transcripts. Personal identifiers are anonymized, and sensitive information is masked in publicly released datasets. Provenance metadata remains intact to preserve accountability and traceability within the research process.

5 Results and Discussion

The following section presents the empirical outcomes of the AI-assisted construction pipeline in relation to the conceptual and methodological foundations previously established. Building on the design objectives, ontology structures and evaluation criteria outlined earlier, the results reported here examine how the pipeline performs when applied to domain-relevant materials from the BPiBS project context. A representative use case is employed to demonstrate the pipeline's behaviour under realistic conditions, illustrating its capacity to construct semantically consistent graph structures and its practical relevance within both the BPiBS and the CIV (Collaborative Intelligence Vision) environments. The CIV framework is a hybrid human–AI system called to help researchers and practitioners navigate complex, interconnected knowledge. It integrates curated graph data, large language models and human expertise to support evidence-informed analysis, decision making and roadmap development. CIV enables users to explore relationships across technical, social, environmental and governance themes through a transparent, navigable knowledge-graph interface.

5.1 Representative Use Case: "Circularity to Material Reuse" Pathway

To ground evaluation in a real BPiBS context, the pipeline was applied to a curated use case centred on the frequently discussed problem: "How do we move from theoretical circularity to actual material reuse in construction?" This use case contains rich, cross-domain semantics (steel reuse, selective deconstruction), governance mechanisms (procurement levers, certification), enabling activities (data standardization), and stakeholder roles (industry associations, regulators, fabricators). It therefore provides a realistic test of the pipeline's ability to construct interconnected graph structures rather than isolated facts (Lu et al. 2024).

5.1.1 Input Materials

The construction experiment was conducted using a consolidated evidence set assembled to reflect the diversity of information sources characteristic of the BPiBS domain. This evidence pack included curator-developed descriptions of enabling activities, policy excerpts addressing circularity and embodied-carbon reduction, and technical guidance documents related to material passports and deconstruction practices. To capture practitioner-oriented perspectives, the corpus also incorporated a synthetic persona-based narrative ("David Rousseau") designed within BPiBS to illustrate industry–research collaboration in the context of steel reuse. Finally, a set of curator-verified subgraphs representing established pathways toward construction-material reuse served both as alignment targets and as ground-truth structures for subsequent evaluation.

5.1.2 Extracted Entities and Relations

The construction of the circularity-focused use case produced a structured subgraph centred on the synthetic persona *David Rousseau*, represented in the database as a "Use Case" node linked to his primary inquiry: *"How do we move from theoretical circularity to actual material reuse in construction?"* The pipeline successfully identified this question as a focal concept and aligned it with a set of thematically relevant pathways and domain topics present in the BPiBS knowledge base.

The system connected the research question to a diverse set of entities spanning technical strategies, environmental considerations and material-life-cycle processes. These included *Ecosystem Integration and Biodiversity*, *Sustainable Housing Foundations*, *Sustainable Building Materials*, *Innovative Construction Materials*, and *CleanBC Advancing Sustainable Housing*, alongside more construction-specific themes such as *End-of-Life and Deconstruction*, *Deconstruction Techniques*, *Deconstruction vs. Demolition*, *Salvage and Reuse of Materials*, and *Recycling of Construction and Demolition Waste*. Additional nodes such as *Waste Management and Recycling*, *Waste Minimization and Recycling*, *Green Construction Materials* and the *Climate Help Desk* were also correctly identified as semantically associated with the reuse challenge.

The relational structure generated by the pipeline therefore reflects a coherent mapping between the user's inquiry and the broader topical environment of the BPiBS ontology. The links capture both high-level sustainability themes and operational strategies that modulate material reuse in practice. Importantly, this constructed subgraph aligns closely with the curator-validated structure stored in the database: most associations between the research question and its related topics appear in the ground-truth model, and the pipeline consistently recovered the multi-domain spread of concepts that define this use case. This illustrates the system's capacity to integrate dispersed textual evidence into a unified, semantically structured representation within the CIV environment (Chen et al. 2025).

5.1.3 Human-AI Collaboration for Semantic Precision

Human reviewers assessed ambiguous alignment cases identified by the pipeline, revealing several consistent patterns. Reviewer intervention was most valuable where documents expressed implicit causal relationships or described actors with multiple roles; resolving these ambiguities notably improved relation-level precision. In contrast, routine elements such as simple attributes, dates and provenance information required minimal correction, indicating stable automated performance in low-complexity areas. Prompt refinement also produced measurable gains (Wang et al. 2012; Ji and Grishman 2011). Overall, the results confirm that selective, strategically placed human oversight yields significant improvements in semantic accuracy and graph consistency while keeping curator workload manageable.

6 Summary

This study demonstrates the feasibility and value of an AI-assisted pipeline for constructing curated knowledge graph structures from heterogeneous, domain-specific documents in the housing and construction sector. By combining structured ontology context, large language models and iterative human-in-the-loop validation, the approach narrows the long-standing gap between narrative evidence and computable graph representations. The design-science framework used in this work shows that higher-fidelity graph fragments can be generated consistently when extraction, canonicalization and alignment are guided by explicit ontological constraints and evaluated against ground-truth subgraphs. The findings indicate that semantically reliable construction is achievable without full manual curation and that the resulting graph structures can meaningfully support roadmap development, decision-support tasks and cross-domain learning and organization of information.

6.1 Significance and Predicted Impacts

The current research addresses a persistent and consequential bottleneck in knowledge-intensive domains: the challenge of translating heterogeneous textual evidence into structured, interoperable and queryable knowledge graph formats. By operationalizing an AI-assisted rconstruction pipeline grounded in ontology alignment, provenance preservation and selective human oversight, the work provides a methodological foundation for scaling knowledge-graph curation beyond what is achievable through manual processes alone (Wang et al. 2012; Ji and Grishman 2011). The significance of the pipeline lies in its ability to accelerate iteration cycles while maintaining semantic fidelity. Automated extraction and alignment, combined with targeted expert intervention, enable more frequent updates to curated graph structures, thereby increasing the responsiveness of knowledge systems to emerging policy guidance, technical innovations and stakeholder insights. This directly improves the transparency, traceability and reproducibility of decision-support processes that rely on complex, multi-source evidence.

Embedded within the BPiBS project, the approach enhances the capacity to map interdependencies across technical, social, environmental and governance domains (Hu et al. 2021; Lilis et al. 2025). It supports the identification of systemic gaps, surfaces cross-domain interactions that are difficult to detect through document review alone and enables the maintenance of living, navigable representations of the housing ecosystem. These structured representations can be seamlessly integrated into tools such as the CIV environment, strengthening roadmap development, scenario exploration and cross-sector alignment.

Beyond the immediate context, the methods are broadly generalizable. Many civil-infrastructure fields, transportation, energy systems, environmental planning and public asset management, face similar demands to integrate large volumes of narrative evidence into computable knowledge systems (Chen et al. 2025; Lu et al. 2024). The design principles and evaluation protocol demonstrated here offer a transferable template for building semantically disciplined, scalable and curator-aware information infrastructures. As such, the predicted impacts extend beyond efficiency gains to include enhanced knowledge governance, improved analytical coherence and strengthened institutional capacity for evidence-informed decision-making.

6.2 Limitations

Several limitations emerged through the development and evaluation of the construction pipeline. First, the system is optimized for text-based inputs and provides only partial support for complex tables, figures or other multimodal content. Information embedded in these formats often requires manual interpretation or preprocessing, which constrains full automation. Second, long-range reasoning across documents remains a substantive challenge. Relationships that depend on multi-step inference, dispersed evidence or implicit contextual cues are more prone to omission or misalignment, particularly when semantic cues are weakly expressed.

Variation in document style and genre also affects extraction performance. Policy briefs, technical reports and narrative descriptions exhibit distinct linguistic patterns, and model behaviour can drift when transitions between genres are not explicitly managed. In addition, ontology evolution introduces non-trivial complexity: when class hierarchies or relation types change, alignment rules and prompt templates must be updated to maintain consistency and recall typically decreases until recalibrated.

Finally, although human-in-the-loop mechanisms substantially reduce semantic drift, they introduce dependencies on reviewer judgment. Variability in interpretation can affect alignment outcomes unless supported by clear review protocols and standardized decision criteria.

6.3 Acknowledgements

This research was conducted within the Best Practices in Building Systems (BPiBS) initiative, a collaboration between BC Housing, the University of Victoria and project partners under the Housing, Infrastructure and Communities Canada (HICC) Research and Knowledge Initiative (RKI). The authors gratefully acknowledge the financial, institutional and logistical support provided through this partnership. We also thank the CIV (Collaborative Intelligence Vision) development team at UVic for their sustained guidance during the design and testing of the semantic processing workflow, as well as the domain experts who contributed to ontology development, curation and validation. The authors note that AI-enabled editorial tools were used to support language refinement and formatting during the preparation of this research paper.

7 References

Towards complex ontology alignment using large language models.

Cao, L., Sun, J., & Cross, A. (2024). An automatic and end-to-end system for rare disease knowledge graph construction based on ontology-enhanced large language models: Development study. JMIR Medical Informatics, 12, e60665

Chen, C., Zhang, J., Zhao, H., & Li, Y. (2025). Digital Twin-Based and Knowledge Graph-Enhanced Real-Time Risk Assessment and Emergency Response Framework for Tunnel Construction. Applied Sciences, 15(11), 6009.

Dagdelen, J., Dunn, A., Lee, S., Walker, N., Rosen, A. S., Ceder, G., Persson, K. A., & Jain, A. (2024). Structured information extraction from scientific text with large language models. Nature Communications, 15, 1418.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. MIS Quarterly, 28(1), 75–105.

Hogan, A., Blomqvist, E., Cochez, M., D'Amato, C., de Melo, G., Gutiérrez, C., Kirrane, S., Gayo, J. E. L., Navigli, R., Neumaier, S., Ngomo, A.-C. N., Polleres, A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., & Zimmermann, A. (2021). Knowledge Graphs. ACM Computing Surveys, 54(4), 71.

Hu, Z.-Z., Leng, S., Lin, J.-R., Li, S.-W., & Xiao, Y.-Q. (2021). Knowledge Extraction and Discovery Based on BIM: A Critical Review and Future Directions. Archives of Computational Methods in Engineering, 29(1), 335–356.

Jaradeh, M. Y., Singh, K., Stocker, M., Both, A., Auer, S., Demidova, E., Prinz, M., D'Souza, J., Kismihók, G., & Oelen, A. (2023). Information extraction pipelines for knowledge graphs. Knowledge and Information Systems, 65(5), 1989–2016.

Ji, H., & Grishman, R. (2011). Knowledge Base Population: Successful Approaches and Challenges. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (pp. 1148–1158).

Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2022). A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. IEEE Transactions on Neural Networks and Learning Systems, 33(2), 494–514.

Lilis, G. N., Wang, M., Katsigarakis, K., Mavrokapnidis, D., Korolija, I., & Rovas, D. (2025). BIM-based semantic enrichment and knowledge graph generation via geometric relation checking. Automation in Construction, 173, 106081.

Lu, Z., Sun, C. T., Hu, Y., & Kumar, A. (2024). BIM and Knowledge Graph-Based Building Material Recycle and Reuse Assessment Framework. In Computing in Civil Engineering 2023: Visualization, Information Modeling, and Simulation. American Society of Civil Engineers (ASCE).

Suchanek, F. M., Fan, J., Hoffmann, R., Riedel, S., & Talukdar, P. P. (2013). Advances in Automated Knowledge Base Construction. SIGMOD Record, 42(4), 13–24.

Utkucu, D., et al. (2025). Ontology for Holistic Building Performance Modeling and Evaluation. Automation in Construction (in press).

Wang, D. Z., Chen, Y., Goldberg, S., Grant, C., & Li, K. (2012). Automatic knowledge base construction using probabilistic extraction, deductive reasoning, and human feedback. In AKBC-WEKEX 2012.

Yuan, J., Zhang, F., Qiu, Y., Lin, H., & Zhang, Y. (2024). Document-level biomedical relation extraction via hierarchical tree graph and relation segmentation module. Bioinformatics, 40(7), btae418.

Zhang, B., & Soh, H. (2024). Extract, Define, Canonicalize: An LLM-based framework for knowledge graph construction. In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (pp. 9820–9836).

Zhang, N., Xu, X., Tao, L., Yu, H., Ye, H., Qiao, S., Chen, H. (2023). DeepKE: A deep learning–based knowledge extraction toolkit for knowledge base population.

Zheng, H., Wang, S., & Huang, L. (2024). A comprehensive survey on document-level information extraction. In Proceedings of the Workshop on the Future of Event Detection (FuturED) (pp. 58–72). ACL.