

DFBI 2025 aims to encourage the international exchange of innovative ideas between researchers from academia and industry. In addition to knowledge dissemination, the conference offers a valuable platform for professional networking, particularly benefiting university professors, graduate students, and postdoctoral researchers.

Research Article

Structuring Knowledge for Positive Energy Districts: An Ontology-Driven Knowledge Graph Approach

Medha Phatak¹, Mengjie Han^{1,2*}, Xingxing Zhang^{1,2}

1. School of Information and Engineering, Dalarna University, 79188 Falun, Sweden

2. Sustainable Energy Research Centre (SERC), Dalarna University, 79188 Falun, Sweden

*Correspondence: mea@du.se

Copyright: Copyright: © 2026 by the authors. ISSN: 2979-1677

DFBI is an open-access proceedings distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0). View this license's legal deed at <https://creativecommons.org/licenses/by/4.0/>



Abstract

Positive Energy Districts (PEDs) are discussed across a wide range of policy papers, research studies, and project reports. Yet the terminology and conceptual boundaries used to describe them remain fragmented and inconsistent. This study examines how ontology engineering and Knowledge Graph (KG) methods can bring structure and clarity to this diverse body of material by providing a transparent and reproducible representation of PED knowledge. A study-specific PED ontology was developed to capture recurring concepts in the literature and was aligned with SAREF4City to incorporate relevant smart-city semantics. Using this ontology as a foundation, a KG was constructed from a curated corpus of PED documents by extracting sentences, generating semantic triples, and aligning them with the ontology. The resulting KG links documents, sentences, entities, and conceptual relations, enabling structured retrieval across technical, governance, process, financial, environmental, legal, managerial, and social dimensions. Querying the KG revealed recurring patterns, highlighted dominant themes, and showed that Governance, Process, and Technical concepts appear far more frequently than other aspects, indicating an imbalance in how PEDs are documented. While the approach depends on sentence-level extraction quality and reflects uneven coverage in the underlying corpus, it offers a practical way to support PED planning and implementation. The framework can assist with tasks such as technology selection, governance design, and evidence-based decision-making, and contributes to more transparent, accessible, and socially meaningful PED development. By improving the clarity and accessibility of PED knowledge, the approach supports more inclusive and sustainable urban transformation.

Keywords: Positive Energy Districts; Knowledge Graph; Ontology Engineering; Information Extraction; SAREF4City.

Cite this article:
 Lastname, F. M., & Lastname, F. M.
 (Year). Title of the article. Title of the
 Journal, Volume (Issue), Page range.
<https://doi.org/10.66408/dfbi.2026.26>
 388

1 Introduction

As European cities advance toward climate neutrality targets, Positive Energy Districts (PEDs) have emerged as an integrated urban planning and technological approach for transforming district-level energy systems through demand reduction and decentralised renewable energy generation (Bossi, Gollner, & Theierling, 2020; Zhang et al., 2021). PEDs are generally defined as urban areas that produce more renewable energy than they consume, supported by integrated energy systems, energy-efficient buildings, and active stakeholder involvement. Despite growing interest, knowledge about PEDs remains fragmented across project deliverables, technical reports, and policy documents, often using inconsistent terminology and conceptual boundaries (Ahlers et al., 2023; Alpagut, Akyürek, & Mitre, 2019). This fragmentation makes it difficult to compare approaches, transfer lessons, or build a shared understanding of PED development.

Knowledge Graphs (KGs) offer a structured way to address this challenge by representing entities and relationships in a transparent, queryable format (Hogan et al., 2022; Kejriwal, 2022). While KGs are increasingly used in energy and urban research, their application to PEDs remains limited and often focuses mainly on technical aspects (Hofer, Obraczka, Saeedi, Köpcke, & Rahm, 2024; Lygerakis, Kampelis, & Kolokotsa, 2022). There is a need for a domain-specific approach that can organise heterogeneous PED knowledge and support cross-project learning. To address this fragmentation and enable a more structured understanding of PED practice, the study examines three research questions:

RQ1: What concepts and relationships most frequently appear in PED literature?

RQ2: How can PED documents be transformed into a structured KG?

RQ3: What patterns and conceptual gaps emerge from analysing the KG?

The work contributes a study-specific PED ontology, a reproducible KG pipeline, and insights into dominant and underrepresented PED dimensions, supporting clearer and more systematic PED conceptualisation.

Section 2 outlines the theoretical foundations, Section 3 describes the methodology, Section 4 presents the results, Section 5 discusses the implications, and Section 6 concludes the paper.

2 Theoretical and Conceptual Foundations

This section provides the conceptual basis for the study by defining core terms, reviewing relevant theories and frameworks, and identifying gaps in existing PED and knowledge-modelling research. It concludes by presenting the conceptual model that underpins the study's approach to structuring and analysing PED knowledge.

2.1 Defining Key Concepts

PEDs are commonly described as urban areas that aim to produce more renewable energy than they consume, supported by energy-efficient buildings, integrated energy systems, and active stakeholder engagement (Bossi et al., 2020). Definitions vary across EU projects, with some emphasising technical performance and others highlighting governance, social participation, or financial models. For this study, PEDs are understood as socio-technical systems that combine technological, organisational, and regulatory elements.

KGs are structured representations of entities and their relationships, expressed as machine-readable triples (Hogan et al., 2022). They support semantic querying, reasoning, and transparent knowledge organisation. Ontologies provide the conceptual backbone for KGs by defining classes, properties, and constraints (Kejriwal, 2022). In the context of PEDs, KGs can represent concepts such as energy systems, governance processes, stakeholders, and planning activities.

Information extraction refers to methods for identifying entities and relationships in unstructured text. Modern approaches combine rule-based techniques, statistical models, and pretrained language models to extract meaningful concepts from documents. In this study, extraction is used to transform PED reports and deliverables into structured triples that populate the KG.

These operational definitions establish the conceptual scope for modelling PED knowledge and reflect the study's focus on structured, interpretable representations that support comparison and reuse across projects.

2.2 Existing Theories and Frameworks

PED development is shaped by several major policy and planning frameworks, such as the SET-Plan Action 3.2, the Smart Cities and Communities (SCC) projects, and national energy-transition strategies. These frameworks emphasise coordinated planning, integrated energy systems, and long-term monitoring, but they differ in how they define PED boundaries and priorities.

In parallel, semantic technologies offer established methodologies for modelling domain knowledge. Ontology engineering frameworks such as Methontology (Gómez-Pérez, Fernández-López, & Corcho, 2004) and semantic web standards (RDF, OWL) provide structured approaches for defining concepts and relationships. KG construction pipelines typically include corpus preparation, entity extraction, ontology alignment, triple generation, and storage.

NLP-based extraction frameworks, including transformer-based models, have improved the ability to process large text corpora. However, generative models may introduce errors or hallucinations, making them less suitable for tasks requiring traceability and explicit structure.

Existing PED-related KGs focus mainly on technical or energy-system aspects (Hofer et al., 2024; Lygerakis et al., 2022), with limited integration of governance, social, or financial dimensions. This limits their usefulness for holistic PED planning. In addition to ontology-based approaches, urban-systems knowledge has also been organised using simpler methods such as rule-based classification, dictionary-driven structures, and manual thematic coding (Hogan et al., 2022; Nickel, Murphy, Tresp, & Gabrilovich, 2016; Singh, 2018). However, these methods offer limited semantic structure and interoperability compared to ontology-driven KGs, which is essential for representing PED complexity (Hogan et al., 2022).

Despite these advances, several gaps remain in how PED knowledge is represented and analysed.

2.3 Knowledge Gaps and Research Opportunities

Three gaps emerge from the literature.

First, PED knowledge is dispersed across heterogeneous documents, making it difficult to synthesise or compare insights. Second, semantic modelling efforts connected to PEDs typically emphasise energy technologies, offering only partial coverage of the wider socio-technical aspects of PEDs. Third,

there is no domain-specific, reproducible methodology for extracting and structuring PED knowledge in a way that supports cross-project learning.

Addressing these gaps is essential for enabling evidence-based planning, improving transparency, and supporting replication of PED approaches across cities.

2.4 Proposed Conceptual Model

This study adopts an ontology-driven KG model that integrates PED concepts, extracted entities, and document-level evidence. The model consists of four components: a study-specific PED ontology aligned with SAREF4City (ETSI, 2026), a curated document corpus, an extraction pipeline for generating triples, and a KG that supports semantic querying and pattern discovery.

Together, these components provide a structured lens for analysing PED knowledge, address the study’s research questions, and fill the gaps identified in Section 2.3 through a structured, reproducible approach to representing PED knowledge.

3 Methodology

This section outlines the research design, data sources, analytical procedures, and validation steps used to construct the PED KG. The methodology follows a structured, multi-phase workflow integrating ontology development, document processing, information extraction, and KG construction. The overall process is summarised in Figure 1.




Figure 1. Conceptual Diagram of the PED KG Construction Process (adapted from Phatak, 2026)

Figure 1 illustrates the workflow consisting of two stages: (1) ontology development, including domain vocabulary preparation and SAREF4City alignment, and (2) document processing, information extraction, alignment, and KG construction.

3.1 Research Design and Link to Research Questions

The study adopts a qualitative, document-based research design supported by computational methods for information extraction and graph construction. This approach is suitable because PED knowledge is largely embedded in unstructured text such as project deliverables, reports, and policy documents. The workflow links to the research questions: ontology development addresses RQ1,

information extraction and NLP processes address RQ2, and KG construction and analysis address RQ3.

A hybrid, iterative strategy was used. The pipeline was first tested on a single document to refine extraction steps and resolve issues related to consistency and accuracy, following recommendations from (Ahlers et al., 2023; Hofer et al., 2024). Once stabilised, it was scaled to the full corpus in line with established KG development practices (Hogan et al., 2022; Kejriwal, 2022). All development was carried out in Google Colab, using integrated NLP and data-science tools.

3.2 Ontology Construction

Ontology development formed the conceptual backbone of the KG. A study-specific PED ontology was created based on literature, project documentation, and domain understanding. SKOS was imported to support standardised labelling through *prefLabel* and *altLabel* (Miles & Bechhofer, 2009). To ensure semantic interoperability, the ontology was aligned with SAREF4City (ETSI, 2026) using annotation-level mappings, without modifying the external ontology (Euzenat & Shvaiko, 2013; Noy & McGuinness, 2001). The ontology includes classes for PED components such as energy systems, governance processes, stakeholders, planning activities, and measurement properties.

Validation involved consistency checking, manual review, and iterative refinement based on extracted document triples. The ontology served both as a conceptual model and as a target schema for aligning extracted entities and relationships.

3.3 Data Acquisition and Pre-processing

A curated corpus of PED-related documents was assembled, including project deliverables, technical reports, and publications from EU-funded initiatives. Documents were selected based on relevance, availability, and representation of diverse PED contexts. The final 56-document corpus was selected because these documents explicitly addressed PED planning, implementation, or evaluation, and collectively represented diverse project types, geographical contexts, and thematic areas, ensuring broad coverage of PED-related knowledge.

Text was extracted, cleaned, and standardised to ensure compatibility with downstream processing.

A domain vocabulary was prepared by compiling key PED terms from the ontology and literature. This vocabulary supported entity recognition and improved extraction accuracy.

3.4 Information Extraction

Information extraction was conducted in two stages.

Single-document extraction

A pilot document was processed to identify sentences containing PED-related concepts. Entities and relationships were extracted using a hybrid approach combining rule-based patterns, statistical methods, and pretrained language models. This stage enabled refinement of filtering rules, predicate handling, and alignment thresholds.

To ensure that alignment thresholds were not chosen arbitrarily, they were calibrated through an empirical grid search that tested multiple combinations of predicate, subject, and object similarity scores. Thresholds were selected based on configurations that produced stable and semantically coherent alignments, following established ontology-matching practices (Euzenat & Shvaiko, 2013;

Gómez-Pérez et al., 2004). This provided a balanced and defensible set of thresholds for aligning extracted triples to the ontology.

Corpus-level extraction

Extracted triples were aggregated across the corpus, normalised, and aligned with ontology classes and properties. Triple classification distinguished valid domain relations from noise, and ontology alignment ensured semantic consistency.

3.5 KG Construction and Access

The cleaned and aligned triples were transformed into a graph structure and imported into Neo4j, a graph database suited for semantic querying and visual exploration. Data preparation included entity normalisation, relationship mapping, and the creation of document-level provenance links to support traceability. The resulting KG integrates documents, sentences, entities, and relationships, enabling multi-level analysis and evidence-based querying. Representative queries demonstrated the KG's usefulness for pattern discovery, evidence tracing, and identifying conceptual gaps.

3.6 Evaluation and Methodological Limitations

Evaluation focused on three dimensions aligned with the research questions

- Conceptual coverage (RQ1): assessed through ontology alignment and entity frequency analysis.
- Extraction accuracy (RQ2): evaluated through manual correctness assessment, alignment quality, and stability checks.
- Representational usefulness (RQ3): assessed through query-based demonstrations.

Limitations include dependency on document quality, extraction errors inherent to NLP methods, and the evolving nature of PED terminology. Ethical considerations were limited, as the study relied exclusively on publicly available documents and did not involve personal or sensitive information.

4 Results – Key Findings

This section presents the key findings of the ontology development, knowledge extraction, and KG construction processes. Results are organised according to the three research questions and focus on the main outputs of the pipeline.

4.1 Results for RQ1: Identification and Structuring of PED Concepts

Ontology development produced a Core study-specific PED Ontology containing 444 classes, 369 object properties, and 25 data properties. After integrating corpus-derived extensions, the Final Integrated study-specific PED Ontology expanded to 1,099 classes, 782 object properties, and 8,783 axioms. Table 1 summarises the structure of the ontology before and after corpus-driven extension. A total of 76 alignments were created with SAREF4City (45 classes, 23 object properties, 8 data properties). The ontology was validated using the HermiT reasoner, with no inconsistencies detected.

Axiom count increased because ontology extension introduced new classes and object properties, whereas data properties remained constant. Measurement-related data properties were designed as a stable schema aligned with SAREF4City, and instances are generated dynamically in Neo4j from extracted triples rather than being manually created in Protégé (Rijgersberg, van Assem, & Top, 2013).

Table 1. Structure and Size of the PED Ontology Before and After Extension (adapted from Phatak, 2026)

Metric	Core PED Ontology (study-specific)	Final integrated PED Ontology (study-specific)
Classes	444	1099
Object properties	369	782
Data properties	25	25
Axioms	4082	8783
Imported ontologies	SKOS, SAREF4City	SKOS, SAREF4City, Corpus-extended

Note: The increase reflects corpus-driven ontology extension

Corpus-driven extension added 692 new classes and a controlled number of new properties derived from 1,212 canonicalized triples. A domain vocabulary of 839 PED-specific terms and 24 instance terms was extracted to support ontology-guided information extraction.

Frequency-based analysis showed that concepts such as energy, city, project, partner, and stakeholder appeared most frequently across the corpus, while predicates such as *ENSURES*, *DEVELOP*, *COMMUNICATES*, and *INVOLVES* were the most recurrent.

These frequency patterns show that the ontology and KG capture the dominant conceptual structures of PED literature.

4.2 Results for RQ2: Extraction of Entities and Relationships

The 56-document corpus yielded 18,850 clean sentences after preprocessing. The extraction pipeline generated 25,716 initial triples, reduced to 2,162 cleaned triples after filtering and consolidation. Following quality assurance, 2,085 triples were retained for ontology alignment.

A corpus-level precision assessment based on a random sample of 100 triples produced a correctness rate of 0.73 (95% CI: 0.645–0.815). Triple classification showed 1,238 concept-level triples, 412 instance-level triples, 331 contextual triples, and 104 noise triples.

Ontology alignment using pre-extension ontology resulted in 1,117 ontology-related triples and 255 fully aligned triples. After extension, alignment improved to 1,459 ontology-related triples and 801 fully aligned triples, with 1,411 Neo4j-ready edges generated. Frequent aligned relations included *ENSURES* (133), *DEVELOP* (100), *COMMUNICATES* (68), and *INVOLVES* (57).

Alignment thresholds were selected through a semantic grid search, and the final configuration demonstrated stable behaviour under threshold perturbation. Sensitivity diagnostics showed invariant counts for ontology-related triples and moderate, expected variation in fully aligned triples. Bootstrap resampling introduced natural sampling variation but preserved the overall stability pattern. Overall alignment accuracy was 79.83%, manual triple-correctness evaluation yielded 73%, and ontology concept coverage reached 53.44%, indicating that the pipeline produced a reliable and representative semantic structure.

4.3 Results for RQ3: Representational Completeness and Usefulness of the KG

The final KG contained 1,382 sentence nodes, 1,770 triple nodes, and 602 entity nodes (263 curated PED entities and 339 corpus-derived entities). All triples were linked to their supporting sentences, and concept-level triples were grounded through DOMAIN and RANGE links.

PED concepts were distributed across nine dimensions: Governance (58), Process (51), Technical (50), Managerial (35), Social (23), Environmental (20), Financial (14), and Legal (12). Figure 2 illustrates this

distribution of the curated PED entities mapped to each of the PED dimensions. Governance, Process, and Technical layers were the most represented.




Figure 2. Distribution of Curated PED Concepts Across PED Layers

KG queries demonstrated that the graph supports multi-layered retrieval, pattern identification, and evidence tracing. Query outputs highlighted recurring themes such as data-driven modelling, stakeholder engagement, and structured planning processes. Gap-related concepts (e.g., *barrier*, *challenge*, *constraint*) were present primarily at the instance level and were analysed through combined triple- and sentence-level evidence.




Figure 3. Asset and Data Sub-Graph Extracted from the PED Knowledge Graph (adapted from Phatak, 2026)

A representative portion of the populated knowledge graph is illustrated in Figure 3, highlighting semantic relationships supported by document-level evidence. For example, Figure 3 shows that the multinational technology company ABB is linked to both data and asset-related concepts through the relation DEVELOP, illustrating how stakeholder-relevant information is represented within the graph.

The stability of the KG structure, combined with consistent alignment behaviour, confirms that the graph provides a coherent and traceable representation of PED knowledge suitable for practical querying and analysis.

5 Discussion

This study set out to address the conceptual fragmentation that characterises current PED literature. As noted in the introduction, PEDs are defined inconsistently across projects and policy documents, making it difficult to compare approaches or build cumulative knowledge. The three research questions were therefore designed to examine how ontology engineering and KG methods can provide a structured and explainable representation of PED concepts, terminology, and relationships.

The ontology developed in this study extends SAREF4City with PED-specific concepts related to energy flows, planning processes, and stakeholder interactions, directly responding to the need for clearer conceptual grounding. Although the Smart City Ontology 2.0 (SCO 2.0) could not be integrated due to the absence of a formal OWL version, the resulting ontology provides a coherent and extensible foundation for future semantic modelling in the PED domain.

The KG complements this ontology by extracting entities, relations, and evidence from 56 PED documents. Its layered structure, consisting of documents, sentences, and triples, balances abstraction with traceability and reflects established practices in KG construction. The two-stage extraction strategy, beginning with a single-document pilot, enabled refinement of filtering rules and alignment thresholds. Manual correctness checks and confidence-interval diagnostics confirmed that the pipeline behaves consistently under parameter variation, strengthening confidence in the resulting semantic structure. This finding is consistent with prior research on knowledge graphs in urban and energy domains, which has primarily focused on technical representations (Hofer et al., 2024; Hogan et al., 2022). However, the present study extends these approaches by incorporating governance, social, and process-related dimensions specific to PEDs, thereby supporting a more holistic representation of PED knowledge.

As a use case, the KG can provide quick information retrieval without cross-checking multiple documents. In the PED-ACT project¹, for example, the city Umeå is considered as the best-performing case city, reaching the high-performance level and demonstrating strong PED alignment across eight aspects. In addition to quickly retrieving which aspect Umeå performs best, the KG can return what this specific aspect aims for in another PED manifesto document², only three queries are required. This method will significantly save manual tasks and avoid errors when the number of documents is large. Compared to LLMs, KG outputs are more accurate, relevant, and traceable. The KG can help explain why Umeå performs strongly and highlight where its documentation is thin, making it easier to benchmark Umeå against other PED cases and to identify gaps across different PED aspects.

Beyond this example, the KG can be applied in planning and policy work. Planners can use it to locate relevant precedents, compare how different cities address PED challenges, and trace the evidence behind concepts without reviewing dozens of documents. Policymakers can use the KG to check alignment with EU frameworks, identify governance or stakeholder elements, and ensure consistent

¹ <https://ped-act.com/>

² <https://wonderland.cx/the-ped-manifesto/>

terminology across strategies. Because triples link to sources, the KG supports transparent decision justification.

Within its scope, the study is limited by its focus on European PED projects, the lack of numerical data extraction, and reliance on lightweight NLP models. However, the measurement-related data properties in the ontology are not yet populated, even though the schema supports quantitative values. Integrating SCO 2.0, enriching socio-technical concepts, adding temporal data, and applying more advanced extraction models would further enhance the KG's analytical potential. Overall, the study demonstrates that an ontology–KG approach can help bring greater clarity, consistency, and analytical depth to PED research and practice.

As a result, the KG reflects EU-specific policy frameworks, governance models, and terminology. Its application in other regions would require adaptation of vocabulary, ontology extensions, and inclusion of region-specific documents.

6 Conclusions

This study set out to address the conceptual fragmentation that characterises current PED literature by examining how ontology engineering and KG methods can provide a structured and transparent representation of PED knowledge. By developing a study-specific ontology aligned with SAREF4City and constructing a corpus-derived KG, the research addressed all three objectives. The ontology captured recurring PED concepts identified across the literature, the extraction pipeline produced a stable set of entities and semantic relations, and the resulting KG supported meaningful querying across tools, data practices, engagement processes, planning activities, and gap-related themes.

These findings contribute to the field of architecture and urbanism by demonstrating that PED knowledge, which is often dispersed across heterogeneous reports and project documents, can be organised into a coherent semantic structure. This structured representation offers practical value for planners, researchers, and policymakers by enabling faster retrieval of relevant information, clearer terminology, and more consistent comparison across PED initiatives. This has direct implications for urban planning and policy development, where structured and comparable knowledge is essential for designing, evaluating, and scaling PED initiatives. More broadly, the study illustrates how semantic technologies can support evidence-based decision-making in emerging urban sustainability domains.

The scope of the study introduces some limitations: numerical data were not extracted, SCO 2.0 could not be fully integrated, and the corpus was limited to European sources. Future work should incorporate temporal and quantitative data, apply more advanced extraction models, and expand the ontology as new standards become available.

Overall, the ontology–KG framework developed here provides a strong foundation for advancing semantic modelling in PED research. It has the potential to support more transparent, comparable, and effective PED planning and implementation.

Acknowledgements

The authors thank Mia Xiaoyun Zhao, Dalarna University, for providing valuable feedback that helped improve the clarity and consistency of the manuscript.

Funding

This research was funded the Swedish Energy Agency (Project ID: P2024-02691 and P2024-02856), under (1) the Clean Energy Transition Partnership, co-funded by the European Commission (GA N°101069750) and (2) the Driving Urban Transition Partnership, co-funded by the European Commission (GA N° F-DUT-2023-0033).

Data Availability Statement

Data is available on request.

Conflicts of Interest

The authors declare no conflict of interest.

References

- Ahlers, D., Riedesel, K., Dommerholt, T., Amin, S., Wyckmans, A., Junqueira de Andrade, E., & Kvaal, B. (2023). *How to PED – The +CityxChange cookbook: Experiences and guidelines on Positive Energy Districts*. Zenodo. <https://doi.org/10.5281/zenodo.8372848>
- Alpagut, B., Akyürek, Ö., & Mitre, E. M. (2019). Positive Energy Districts methodology and its replication potential. *Proceedings*, 20(1), 8. <https://doi.org/10.3390/proceedings2019020008>
- Bossi, S., Gollner, C., & Theierling, S. (2020). Towards 100 Positive Energy Districts in Europe: Preliminary data analysis of 61 European cases. *Energies*, 13(22), 6083. <https://doi.org/10.3390/en13226083>
- ETSI. (2026). SAREF4CITY: An extension of SAREF for the Smart City domain. Retrieved February 27, 2026, from <https://saref.etsi.org/saref4city/v2.1.1/>
- Euzenat, J., & Shvaiko, P. (2013). *Ontology Matching*. Berlin: Springer. <https://doi.org/10.1007/978-3-642-38721-0>
- Gómez-Pérez, A., Fernández-López, M., & Corcho, O. (2004). *Ontological engineering: With examples from the areas of knowledge management, e-Commerce and the semantic web*. Springer.
- Hofer, M., Obraczka, D., Saeedi, A., Köpcke, H., & Rahm, E. (2024). Construction of knowledge graphs: Current state and challenges. *Information*, 15(8), 509. <https://doi.org/10.3390/info15080509>
- Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., Melo, G. de, Gutierrez, C., ... Zimmermann, A. (2022). Knowledge graphs. *ACM Computing Surveys*, 54(4), 1–37. <https://doi.org/10.1145/3447772>
- Kejriwal, M. (2022). Knowledge graphs: A practical review of the research landscape. *Information*, 13(4), 161. <https://doi.org/10.3390/info13040161>
- Lygerakis, F., Kampelis, N., & Kolokotsa, D. (2022). Knowledge graphs' ontologies and applications for energy efficiency in buildings: A review. *Energies*, 15(20), 7520. <https://doi.org/10.3390/en15207520>
- Miles, A., & Bechhofer, S. (2009). SKOS Simple Knowledge Organization System Reference. Retrieved December 28, 2025, from <https://www.w3.org/TR/2009/REC-skos-reference-20090818/>
- Nickel, M., Murphy, K., Tresp, V., & Gábrilovich, E. (2016). A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1), 11–33. <https://doi.org/10.1109/JPROC.2015.2483592>
- Noy, N. F., & McGuinness, D. L. (2001). What is an ontology and why we need it. Retrieved September 16, 2025, from https://protege.stanford.edu/publications/ontology_development/ontology101-noy-mcguinness.html
- Phatak, M. (2026). *Information extraction for constructing a knowledge graph of Positive Energy District*. Master's thesis, Dalarna University. Retrieved April 10, 2026, from <https://urn.kb.se/resolve?urn=urn:nbn:se:du-53284>
- Rijgersberg, H., van Assem, M., & Top, J. (2013). Ontology of units of measure and related concepts. *Semantic Web*, 4(1), 3–13. <https://doi.org/10.3233/SW-2012-0069>
- Singh, S. (2018, July 6). *Natural language processing for information extraction*. arXiv:1807.02383. <https://doi.org/10.48550/arXiv.1807.02383>
- Zhang, X., Penaka, S. R., Giriraj, S., Sánchez, M. N., Civiero, P., & Vandevyvere, H. (2021). Characterizing Positive Energy District (PED) through a preliminary review of 60 existing projects in Europe. *Buildings*, 11(8), 318. <https://doi.org/10.3390/buildings11080318>

Disclaimer/Publisher's Note

The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and do not reflect the views of the Architecture, Buildings, Construction and Cities (ABC2) Journal and/or its editor(s). DFBI Journal and/or its editor(s) disclaim any responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.