

Knowledge Graph Integration for Enhancing Semantic Search and Intelligent Information Retrieval

Dr. Noëlle van Dijk

Associate Professor, Lowlands Institute of Intelligent Information Systems, Netherlands

Abstract

The increasing availability of digital information has created significant challenges for search systems required to interpret complex queries and retrieve contextually relevant documents. Conventional information retrieval techniques commonly depend on lexical similarity between user queries and indexed content. Although these methods provide efficient mechanisms for identifying documents containing relevant keywords, they may encounter limitations when search queries involve ambiguous terminology, alternative entity names, implicit relationships, or multiple connected concepts. Knowledge graphs offer a structured approach to representing entities, attributes, and semantic relationships, creating opportunities to improve information retrieval through explicit contextual knowledge. The present paper develops a hypothetical methodological framework for investigating knowledge graph integration in semantic search and intelligent information retrieval. Four retrieval configurations are considered: BM25 lexical retrieval, entity-aware query expansion, knowledge graph-based relationship expansion, and hybrid semantic retrieval with graph-informed reranking. The corresponding knowledge graph contains 2,400 entities and 7,200 relationship triples representing conceptual connections among creators, artworks, historical periods, artistic movements, and collection objects.

The proposed methodology incorporates textual preprocessing, entity recognition, entity disambiguation, graph construction, ontology alignment, relationship-based query processing, and comparative retrieval evaluation. A hypothetical collection of 240 information-seeking queries is organized into exact-entity, alias-based, relationship-oriented, and multi-hop categories. Two illustrative analytical tables describe the evaluation dataset and comparative ranking performance. Normalized Discounted Cumulative Gain, Recall, and Mean Reciprocal Rank are considered principal evaluation measures. All numerical values are manually constructed and do not represent completed retrieval experiments. The methodological discussion examines semantic ambiguity, knowledge graph completeness, relationship accuracy, computational complexity, and evaluation reliability. Particular attention is given to distinguishing valid knowledge graph relationships from inferred or unsupported connections.

Keywords: Knowledge Graph, Semantic Search, Intelligent Information Retrieval, Entity Linking, Ontology, Linked Data, Graph-Based Query Expansion, Hybrid Retrieval, Semantic Relationships, Information Ranking.

1. Introduction

Digital information systems have become fundamental resources for storing, organizing, and retrieving knowledge across academic, cultural, commercial, and institutional environments. The continuous expansion of digital collections has increased the amount of information available to users while also introducing challenges associated with efficient discovery. Conventional search engines commonly identify relevant content by comparing query terms with indexed document representations. These approaches are useful when the vocabulary of a query corresponds closely to the vocabulary of relevant documents. However, information needs frequently involve conceptual relationships, alternative terminology, and references to entities whose meanings cannot be established through isolated keywords.

Semantic search represents an approach to information retrieval that seeks to interpret the intended meaning of a query rather than relying exclusively on literal word matching. A user searching a digital cultural heritage repository may request artworks associated with a particular movement, objects created by an artist, or collections connected to a historical period. Such queries depend on relationships among concepts and entities. Relevant documents may not explicitly contain every term used in the query, even when their underlying content satisfies the intended information need. Consequently, retrieval effectiveness may depend on the ability to represent relationships among objects, creators, collections, and historical contexts.

Knowledge graphs provide a structured representation of information through entities and the relationships connecting them. An entity may represent a person, location, organization, object, event, or abstract concept, while relationships describe how entities are associated. Knowledge graphs commonly use structured assertions in which a subject is connected to an object through a defined predicate. These representations provide opportunities to support entity disambiguation, relationship traversal, semantic query expansion, and knowledge-based information organization. However, the usefulness of a knowledge graph depends on the completeness, accuracy, and interpretability of its underlying relationships.

Integrating knowledge graphs with information retrieval systems introduces additional computational and methodological possibilities. Entity linking can connect textual references with canonical graph entities, while graph traversal can identify relationships not explicitly expressed within an individual document. Semantic query expansion may incorporate verified aliases or related concepts, and graph-informed reranking can prioritize documents satisfying the relational structure of a query. Nevertheless, these procedures can introduce errors when entity identification is inaccurate or when the graph contains incomplete or unreliable relationships. Consequently, semantic enrichment should be evaluated according to its contribution to relevance rather than the quantity of additional terms or connections generated.

The present paper develops a hypothetical comparative framework for examining knowledge graph integration in intelligent information retrieval. A fictional digital cultural heritage repository is selected as the application environment because it contains naturally structured relationships among artworks, creators, historical movements, and collections. Four retrieval configurations are proposed, ranging from conventional lexical ranking to hybrid semantic retrieval with graph-informed reranking. The methodology emphasizes entity resolution, relational query interpretation, independent evaluation, and the reliability of retrieved evidence. No actual catalog records were indexed, and no retrieval

experiments were executed. The paper provides a structured methodological foundation for future empirical investigations.

2. Objectives of the Study

The primary objective of this study is to examine the methodological potential of knowledge graph integration for improving semantic search and intelligent information retrieval. The investigation focuses on entity representation, relationship-aware retrieval, comparative ranking performance, and analytical reliability. Three specific objectives guide the proposed research framework.

2.1 Examination of Structured Knowledge and Semantic Relationships

The first objective is to examine how entities and their relationships can be organized within a structured knowledge graph to support information retrieval. Digital repositories often contain related information distributed across textual descriptions, catalog metadata, and classification systems. The proposed investigation considers how these sources could be integrated through canonical entity identifiers and clearly defined relationship types. Particular attention is given to representing creators, collection objects, artistic movements, and historical periods without losing the distinctions among different kinds of relationships.

The investigation further aims to identify how structured semantic representations may address limitations associated with keyword-based information retrieval. A catalog record may contain an alternative spelling of an artist's name or describe a work through terminology different from that used in a search query. Knowledge graph integration could connect these expressions when reliable entity mappings exist. However, relationships should not be introduced solely because two concepts appear semantically similar. This objective emphasizes accurate entity identification, relationship validation, and the preservation of meaningful distinctions within the represented information.

2.2 Comparative Evaluation of Knowledge-Enhanced Retrieval Methods

The second objective focuses on establishing a comparative framework for evaluating lexical retrieval and knowledge-enhanced search configurations. BM25 provides the conventional search baseline, while entity-aware query expansion introduces verified entity aliases. Relationship expansion incorporates selected graph connections, and hybrid retrieval combines textual relevance with graph-informed reranking. The proposed comparison examines nDCG@10, Recall@10, and MRR@10. These measures provide complementary information about the quality and ordering of retrieved results.

A further objective involves examining whether knowledge graph integration offers benefits for queries requiring relationships among multiple entities. Exact-name searches may be handled effectively through conventional indexing, whereas relation-oriented and multi-hop queries can require additional contextual information. The proposed framework therefore distinguishes query categories during evaluation. This objective emphasizes determining where structured knowledge contributes useful retrieval information rather than assuming that graph integration improves every search task. Consistent datasets and evaluation procedures are necessary for defensible comparison.

2.3 Assessment of Semantic Retrieval Reliability and Integration Constraints

The third objective is to identify methodological challenges that may influence knowledge-enhanced search reliability. Entity ambiguity, incomplete graph relationships, inconsistent ontology definitions, and errors in metadata integration can affect retrieval outcomes. The proposed study considers how

entity disambiguation, graph validation, and controlled expansion procedures could reduce selected difficulties. Particular attention is given to preventing erroneous relationships from increasing the ranking of irrelevant documents.

The investigation also considers computational and evaluation requirements associated with intelligent retrieval systems. Graph traversal and relation-based reranking may introduce processing costs beyond those required for conventional lexical ranking. Relevance assessment can also become more complicated when queries involve several connected concepts. Consequently, this objective emphasizes transparent retrieval logic, suitable evaluation metrics, and the assessment of graph-related errors. These considerations establish a foundation for evaluating semantic search beyond general assumptions about contextual understanding.

3. Review of Literature

3.1 Development of Semantic Web Technologies and Structured Knowledge

Berners-Lee et al. (2001) described the Semantic Web as an approach for organizing information in forms that computational systems can interpret through explicit meaning and relationships. Horrocks et al. (2003) examined the development of the Web Ontology Language and the formal representation of semantic relationships. Auer et al. (2007) introduced DBpedia as a structured knowledge resource derived from Wikipedia. These contributions established important foundations for machine-readable semantic representations and demonstrated how structured information could be connected across heterogeneous digital resources.

Suchanek et al. (2008) presented YAGO as a large ontology derived from Wikipedia and WordNet. Bollacker et al. (2008) described Freebase as a collaboratively developed structured knowledge repository. Bizer et al. (2009) subsequently examined Linked Data principles and the interconnection of distributed information resources. Vrandečić and Krötzsch (2014) introduced the design and collaborative knowledge organization principles of Wikidata. Together, these contributions established important models for constructing interconnected knowledge resources and supporting applications that depend on explicit entity relationships.

3.2 Entity Linking and Relationship-Aware Information Retrieval

Mihalcea and Csomai (2007) investigated automated linking between textual documents and encyclopedic knowledge resources. Their approach demonstrated how important concepts within documents could be associated with structured background information. Milne and Witten (2008) examined relationships among Wikipedia concepts and their potential usefulness for entity linking. Navigli (2009) reviewed word-sense disambiguation methods, emphasizing the difficulties associated with interpreting ambiguous linguistic expressions. These contributions established foundations for connecting unstructured text with semantic representations.

Hoffart et al. (2011) investigated named-entity disambiguation through contextual similarity and coherence among linked entities. Navigli and Ponzetto (2012) developed BabelNet by combining lexical and encyclopedic information across languages. Shen et al. (2015) reviewed entity linking methods and their relevance to information retrieval and knowledge-base applications. These studies demonstrated that successful semantic retrieval depends substantially on identifying the correct entities represented within both documents and queries. They also highlighted the importance of contextual disambiguation when several entities share similar names or descriptions.

3.3 Knowledge Graph Representation and Retrieval Evaluation

Bordes et al. (2013) investigated translation-based embeddings for representing multi-relational knowledge. Wang et al. (2014) developed a related approach using relation-specific hyperplanes, while Lin et al. (2015) examined separate entity and relation spaces for knowledge graph representation learning. Nickel et al. (2016) reviewed relational machine learning techniques for knowledge graphs, and Ristoski and Paulheim (2016) examined RDF graph embeddings for data mining. These contributions expanded the computational possibilities for representing structured knowledge through numerical features.

Robertson and Zaragoza (2009) reviewed probabilistic relevance methods associated with BM25 retrieval. Manning et al. (2008) provided a comprehensive treatment of indexing, document ranking, and information retrieval evaluation. Järvelin and Kekäläinen (2002) developed cumulative gain-based ranking measures, while Voorhees (2000) investigated the influence of relevance judgment variation on retrieval evaluation. These studies established methodological foundations for comparing lexical and knowledge-enhanced retrieval systems through controlled experiments and independently assessed document relevance.

4. Research Methodology

4.1 Research Design and Knowledge Graph Integration Framework

The present paper adopts a hypothetical comparative quantitative research design for examining knowledge graph integration within a digital cultural heritage information retrieval environment. Four configurations are proposed: conventional BM25 retrieval, BM25 with entity-alias query expansion, BM25 with relationship-aware graph expansion, and hybrid semantic retrieval with graph-informed reranking. The intended comparison uses a common set of catalog records and information-seeking queries. No actual retrieval models or graph-processing systems were implemented for this illustrative study.

The proposed analytical framework involves catalog data preparation, entity identification, graph construction, query interpretation, candidate retrieval, graph-based expansion, and ranking evaluation. The conventional configuration would use indexed textual descriptions. Entity-aware retrieval would incorporate verified alternative names associated with canonical entities. Relationship expansion would examine relevant one-hop and selected two-hop graph connections. The hybrid system would combine lexical relevance with graph-supported candidate scoring. All retrieval configurations would use the same independent evaluation queries.

4.2 Hypothetical Digital Repository and Graph Construction

For illustrative purposes, the study assumes a fictional repository containing 4,800 cultural heritage catalog records. The corresponding knowledge graph includes 2,400 distinct entities and 7,200 relationship triples. Entity categories include artworks, creators, artistic movements, historical periods, and collections. Relationship predicates may represent connections such as *createdBy*, *associatedWithMovement*, *heldInCollection*, and *associatedWithPeriod*. These records and relationships are manually specified as methodological assumptions and do not represent actual museum collections or externally retrieved knowledge graphs.

The proposed graph construction procedure would begin with the extraction of relevant information from catalog metadata and textual descriptions. Entity mentions would be linked to

canonical identifiers using verified contextual information. Relationship assertions would be checked against a predefined ontology to ensure that subject and object types satisfy the expected semantic constraints. Duplicate entities would be consolidated only when their identities could be established reliably. The resulting graph would support query interpretation and document association while preserving the provenance of each asserted relationship.

4.3 Search Configuration and Comparative Evaluation

The proposed query dataset contains 240 information-seeking requests divided equally among exact-entity, alias-based, relationship-oriented, and multi-hop categories. Half of the queries would support development, one quarter validation, and one quarter independent testing. Relevance judgments would use a three-level scale distinguishing irrelevant, partially relevant, and highly relevant catalog records. Independent assessors with appropriate domain knowledge would apply predefined criteria, with disagreements resolved through documented adjudication procedures.

Retrieval performance would be evaluated using $nDCG@10$, $Recall@10$, and $MRR@10$. Graded relevance judgments would inform $nDCG$, while $Recall@10$ would measure the proportion of judged-relevant records recovered within the first ten results. $MRR@10$ would emphasize the rank of the first highly relevant document. The proposed assessment would also examine entity-linking accuracy, erroneous relation expansion, and ranking latency. Repeated evaluation and uncertainty estimation would be necessary before establishing meaningful performance differences. No retrieval experiments were executed in preparing this manuscript.

5. Analysis

5.1 Illustrative Distribution of Semantic Information-Seeking Queries

The characteristics of information-seeking queries can influence the comparative performance of lexical and knowledge-enhanced retrieval systems. Exact-entity queries may contain names that directly match indexed catalog descriptions, while alias-based requests require recognition of alternative expressions. Relationship-oriented searches involve connections among entities, and multi-hop queries may require the interpretation of several successive semantic relationships. Table 1 presents the hypothetical distribution of these query categories across the proposed evaluation partitions.

Table 1: Hypothetical Semantic Search Query Distribution

Query Category	Development Queries	Validation Queries	Test Queries	Total
Exact Entity Queries	30	15	15	60
Entity Alias Queries	30	15	15	60
Relationship-Oriented Queries	30	15	15	60
Multi-Hop Semantic Queries	30	15	15	60
Total	120	60	60	240

The fictional dataset contains 240 queries, with 60 allocated to independent testing. The structure enables query-category comparisons but does not establish actual retrieval performance or user information-seeking behavior.

5.2 Illustrative Evaluation of Knowledge-Enhanced Retrieval Configurations

Knowledge-enhanced retrieval systems can be compared by examining how effectively they identify and rank records satisfying a user's information need. $nDCG@10$ evaluates ranking quality while accounting for graded relevance, $Recall@10$ describes coverage of relevant records within the first ten results, and $MRR@10$ emphasizes the position of the first highly relevant result. Table 2 presents hypothetical performance values for four retrieval configurations. The numerical values were manually constructed and do not represent experimentally generated rankings or measured model outputs.

The illustrative comparison assigns the highest ranking scores to hybrid semantic retrieval with graph-informed reranking. Entity-alias expansion and relationship-based expansion receive progressively higher hypothetical values than conventional BM25. This pattern demonstrates one possible scenario in which structured semantic information contributes to retrieval. However, the assigned differences cannot establish actual improvement. A completed investigation would require implemented retrieval systems, independent relevance assessments, reproducible graph processing, and suitable statistical evaluation before identifying a preferred retrieval configuration.

Table 2: Hypothetical Comparative Retrieval Performance With Knowledge Graph Integration

Retrieval Configuration	$nDCG@10$	$Recall@10$	$MRR@10$
Conventional BM25 Retrieval	0.48	0.60	0.45
BM25 with Entity-Alias Expansion	0.56	0.68	0.52
Relationship-Aware Knowledge Graph Expansion	0.63	0.74	0.58
Hybrid Semantic Retrieval with Graph Reranking	0.70	0.82	0.65

The illustrative pattern demonstrates how graph relationships and semantic representations might contribute additional information beyond lexical matching. Nevertheless, the numerical differences are not verified findings and do not demonstrate that graph-enhanced search consistently outperforms conventional retrieval.

6. Discussion

6.1 Knowledge Graphs and Semantic Relationship Interpretation

Knowledge graphs provide a method for representing information through explicitly defined entities and relationships. Within a cultural heritage repository, these relationships can connect objects to creators, movements, collections, and historical periods. Such connections may support queries whose meaning depends on associations not stated directly within a document. However, the usefulness of these relationships depends on the accuracy of entity identification and the validity of the underlying semantic assertions. A graph containing incorrect connections can introduce misleading evidence into the retrieval process.

The proposed framework illustrates how knowledge graphs could provide additional context for information retrieval. Entity aliases may connect different expressions referring to the same creator, while relational paths may identify objects satisfying particular conceptual constraints. Nevertheless, semantic closeness does not automatically establish that every connected entity is relevant to a query. A relationship may be valid in the graph but unrelated to the user's intended information need. Consequently, graph expansion requires appropriate predicate restrictions and relevance-aware ranking rather than unrestricted traversal.

6.2 Entity Linking and Contextual Search Accuracy

Entity linking plays an important role in connecting unstructured query text with structured knowledge graph representations. A search query may refer to an artist through an abbreviated name, a translated expression, or a title shared by several entities. An entity-linking procedure must therefore consider contextual information before selecting the relevant graph identifier. Incorrect entity selection can direct retrieval toward unrelated documents even when the subsequent graph traversal is computationally correct. Reliable semantic search depends on the quality of this initial disambiguation stage.

Research by Hoffart et al. (2011) and Shen et al. (2015) demonstrates the methodological importance of contextual entity identification. Their contributions indicate that entity linking involves more than recognizing matching strings. Candidate entities must be evaluated according to context and their relationships with other recognized concepts. The proposed framework incorporates these principles through controlled linking and relationship validation. However, the hypothetical analysis does not measure actual entity-linking accuracy, and any claimed retrieval improvements would require separate experimental verification.

6.3 Hybrid Ranking and Intelligent Information Retrieval

Hybrid ranking provides a methodological approach for combining conventional lexical relevance with structured semantic information. BM25 can identify documents containing important query terms, while knowledge graph features can provide additional evidence concerning entity relationships. Combining these signals may support retrieval when documents satisfy relational conditions without repeating the exact wording of the query. However, the weights assigned to different ranking components can influence the final results and may introduce unnecessary complexity when graph information is incomplete.

The hypothetical comparison presents a possible pattern in which graph-informed hybrid retrieval receives favorable ranking scores. Nevertheless, these values do not establish empirical superiority. A real investigation should assess whether the improvement arises from entity linking, relationship expansion, or the final reranking procedure. Component-level ablation experiments would help distinguish these contributions. Retrieval latency and the risk of incorrect semantic expansion would also require examination. Consequently, the preferred configuration should be determined through reproducible evaluation rather than aggregate ranking scores alone.

7. Challenges

7.1 Entity Ambiguity and Incomplete Knowledge Representation

Knowledge graph integration can be affected by ambiguity when several entities share similar names or descriptions. Cultural heritage repositories may contain artworks with identical titles, creators with

related names, and historical references that require contextual interpretation. An entity-linking system may select an incorrect graph identifier when the available textual evidence is incomplete. Such mistakes can influence subsequent retrieval because graph traversal and ranking procedures depend on the selected entity. Consequently, semantic search requires accurate disambiguation and careful handling of unresolved references.

Incomplete knowledge representation introduces additional difficulties because the absence of a relationship within a graph does not necessarily establish that the relationship is false. A catalog record may lack information concerning an object's historical period or association with an artistic movement. Strict graph-based filtering can therefore exclude relevant documents when the graph is incomplete. Conversely, unrestricted inference can introduce unsupported connections. The proposed framework emphasizes provenance-aware relationship management and cautious interpretation of missing graph information.

7.2 Graph Consistency and Semantic Integration Quality

Integrating information from different catalog records and knowledge resources can introduce inconsistencies in entity descriptions, relationship types, and classification structures. One source may describe an object through a broad artistic category, while another uses a more specialized classification. These representations may be compatible but require appropriate ontology alignment. In other cases, conflicting metadata may indicate unresolved factual differences. Consequently, graph construction should include consistency assessment and transparent rules for integrating heterogeneous information.

Relationship quality also depends on the validity of the ontology used to organize graph assertions. Predicates should describe meaningful semantic connections, while domain and range constraints should prevent incompatible assertions where appropriate. However, formal consistency alone does not establish factual correctness. A graph may satisfy its schema while containing an incorrect relationship between two entities. The proposed methodology therefore emphasizes source verification, provenance documentation, and targeted quality review. These safeguards are necessary before graph-derived relationships are used to influence search rankings.

7.3 Computational Complexity and Retrieval Evaluation Bias

Knowledge graph-enhanced retrieval may require additional computation for entity identification, relationship traversal, candidate expansion, and reranking. Large graphs can contain numerous connections, making unrestricted traversal inefficient and potentially overwhelming the retrieval process with irrelevant candidates. Multi-hop queries may introduce further processing requirements because multiple relationship paths must be examined. Consequently, graph expansion depth and candidate selection should be controlled according to the intended search application and available computational resources.

Evaluation bias presents another challenge because relevance assessments can vary among reviewers and may depend on the interpretation of complex semantic queries. Documents satisfying only part of a multi-hop condition may be considered partially relevant by one reviewer and irrelevant by another. Incomplete judgments can also influence reported recall and ranking scores. The proposed framework therefore emphasizes clear assessment criteria, independent adjudication, and transparent reporting of unjudged documents. A genuine experiment should also evaluate performance separately across query categories rather than relying exclusively on overall averages.

8. Conclusion

Knowledge graph integration provides a methodological foundation for enhancing semantic search through structured entity representation, contextual disambiguation, and explicit relationship modeling. The present paper developed a hypothetical comparative framework involving conventional BM25 retrieval, entity-aware query expansion, relationship-based graph expansion, and hybrid semantic retrieval with graph-informed reranking. Its principal focus was the relationship between structured knowledge and the identification of relevant documents. The proposed design demonstrates how semantic information could be incorporated into an established information retrieval workflow.

The illustrative tables demonstrate possible procedures for organizing semantic query categories and comparing retrieval performance. However, all numerical values were manually constructed and do not represent actual catalog records, constructed knowledge graph experiments, or independently verified search outcomes. Establishing meaningful performance improvements would require authentic repositories, validated entity relationships, implemented retrieval systems, and independent relevance judgments. Consequently, the hypothetical results cannot establish that knowledge graph integration provides superior retrieval performance under real operating conditions.

Overall, the study identifies entity ambiguity, graph completeness, ontology consistency, relationship accuracy, computational complexity, and evaluation reliability as important considerations in intelligent information retrieval. Future empirical investigations should examine these factors through reproducible experiments, domain-specific datasets, and comparative evaluation across different query categories. Knowledge graphs may strengthen information discovery when their structured relationships provide trustworthy context beyond lexical similarity. Nevertheless, reliable semantic search requires carefully validated knowledge representations and independently supported retrieval performance rather than assumptions concerning the benefits of graph integration.

References

1. Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American*, 284(5), 34–43.
2. Bizer, C., Heath, T., & Berners-Lee, T. (2009). Linked data—The story so far. *International Journal on Semantic Web and Information Systems*, 5(3), 1–22. <https://doi.org/10.4018/jswis.2009081901>
3. Chaudhri, V. K., Baru, C., Chittar, N., Dong, X. L., Genesereth, M., Hendler, J., Kalyanpur, A., Lenat, D., Sequeda, J., Taylor, M., & Wang, D. Z. (2022). Knowledge graphs: Introduction, history, and perspectives. *AI Magazine*, 43(1), 17–29. <https://doi.org/10.1002/aaai.12033>
4. Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutierrez, 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), Article 71. <https://doi.org/10.1145/3447772>
5. Ji, S., Pan, S., Cambria, E., Martinen, 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. <https://doi.org/10.1109/TNNLS.2021.3070843>
6. Koumenides, C., & Shadbolt, N. (2014). Ranking methods for entity-oriented semantic web search. *Journal of the Association for Information Science and Technology*, 65(6), 1099–1112. <https://doi.org/10.1002/asi.23018>

7. Liu, Z., Xiong, C., Sun, M., & Liu, Z. (2018). Entity-duet neural ranking: Understanding the role of knowledge graph semantics in neural information retrieval. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (pp. 2395–2405). Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-1223>
8. Manning, C. D., Raghavan, P., & Schütze, H. (2008). Introduction to information retrieval. Cambridge University Press. <https://nlp.stanford.edu/IR-book/>
9. Nickel, M., Murphy, K., Tresp, V., & Gabrilovich, 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>
10. Noy, N. F., Gao, Y., Jain, A., Narayanan, A., Patterson, A., & Taylor, J. (2019). Industry-scale knowledge graphs: Lessons and challenges. Communications of the ACM, 62(8), 36–43. <https://doi.org/10.1145/3331166>
11. Pérez, J., Arenas, M., & Gutierrez, C. (2009). Semantics and complexity of SPARQL. ACM Transactions on Database Systems, 34(3), Article 16. <https://doi.org/10.1145/1567274.1567278>
12. Reinanda, R., Meij, E., & de Rijke, M. (2020). Knowledge graphs: An information retrieval perspective. Foundations and Trends in Information Retrieval, 14(4), 289–444. <https://doi.org/10.1561/15000000063>
13. Salton, G., & McGill, M. J. (1983). Introduction to modern information retrieval. McGraw-Hill.
14. Suchanek, F. M., Kasneci, G., & Weikum, G. (2007). YAGO: A core of semantic knowledge. In Proceedings of the 16th International Conference on World Wide Web (pp. 697–706). Association for Computing Machinery. <https://doi.org/10.1145/1242572.1242667>
15. Vrandečić, D., & Krötzsch, M. (2014). Wikidata: A free collaborative knowledgebase. Communications of the ACM, 57(10), 78–85. <https://doi.org/10.1145/2629489>
16. Xiong, C., Power, R., & Callan, J. (2017). Explicit semantic ranking for academic search via knowledge graph embedding. In Proceedings of the 26th International Conference on World Wide Web (pp. 1271–1279). International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/3038912.3052558>
17. Xiong, C., Liu, Z., Sun, M., & Liu, Z. (2017). Word-entity duet representations for document ranking. In Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (pp. 763–772). Association for Computing Machinery. <https://doi.org/10.1145/3077136.3080788>
18. Yang, B., Yih, W.-t., He, X., Gao, J., & Deng, L. (2015). Embedding entities and relations for learning and inference in knowledge bases. In Proceedings of the 3rd International Conference on Learning Representations. <https://arxiv.org/abs/1412.6575>
19. Zeng, D., Liu, K., Lai, S., Zhou, G., & Zhao, J. (2014). Relation classification via convolutional deep neural network. In Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers (pp. 2335–2344). Dublin City University and Association for Computational Linguistics.
20. Zhang, Z., Zhuang, F., Zhu, H., Shi, Z., Xiong, H., & He, Q. (2020). Relational graph neural network with hierarchical attention for knowledge graph completion. In Proceedings of the AAAI Conference on Artificial Intelligence, 34(4), 6855–6862. <https://doi.org/10.1609/aaai.v34i04.6161>