Quest Journals
Journal of Research in Business and Management
Volume 11 ~ Issue 10 (2023) pp: 86-93
ISSN(Online):2347-3002
www.questjournals.org

Research Paper

Semantic Layers Reimagined: Building Knowledge Graphs for Next-Generation Business Intelligence

Ramesh Betha
Independent Researcher
Holly Springs NC, US.
ramesh.betha@gmail.com

Abstract—*The evolution of business intelligence (BI) has reached a pivotal juncture, where traditional data warehousing and analytics frameworks are increasingly inadequate for addressing the complexity and scale of modern data ecosystems. Semantic layers, which act as a bridge between raw data and business insights, are being reimagined through the lens of knowledge graphs. This paper explores the transformative potential of knowledge graphs in redefining semantic layers for next-generation BI. By enabling contextual, interconnected, and dynamic data representations, knowledge graphs offer a paradigm shift in how organizations derive value from their data. This paper delves into the architectural principles, implementation challenges, and real-world applications of knowledge graph-driven semantic layers, while also providing a forward-looking perspective on their role in shaping the future of BI.*
Keywords—*Semantic Layers, Knowledge Graphs, Business Intelligence, Data Integration, Contextual Analytics, Data Fabric, Ontology.*

I. INTRODUCTION

The digital transformation of enterprises has led to an exponential growth in data volume, variety, and velocity. Traditional business intelligence (BI) systems, built on relational databases and static semantic layers, struggle to keep pace with the demands of modern analytics. Semantic layers, which abstract the underlying data structures and provide a business-friendly interface for querying and analysis, have long been a cornerstone of BI. However, their conventional implementations often fall short in addressing the need for contextual, real-time, and interconnected insights [1].

Enter knowledge graphs—a powerful paradigm that reimagines semantic layers as dynamic, graph-based structures capable of capturing relationships, context, and meaning. Knowledge graphs are not merely a technological innovation; they represent a fundamental shift in how data is modeled, integrated, and consumed [2]. By embedding semantics directly into the data fabric, knowledge graphs enable organizations to move beyond tabular representations and embrace a more intuitive, graph-oriented approach to BI.

This paper explores how the integration of knowledge graph technologies with semantic layer concepts creates a new paradigm for business intelligence—one that preserves the accessibility benefits of traditional semantic layers while dramatically expanding their capabilities. We examine how knowledge graph-powered semantic layers enable more intuitive data exploration, contextual analytics, and AI-augmented decision support. Furthermore, we investigate how this approach addresses critical enterprise challenges including data governance, cross-domain integration, and the democratization of data access.

As organizations increasingly recognize data as a strategic asset, the need for more sophisticated, flexible, and intelligent approaches to data management and analysis has never been more acute. The reimagined semantic layer concept presented in this paper represents a significant step toward fulfilling the promise of truly knowledge-driven enterprises.

II. THE EVOLUTION OF SEMANTIC LAYERS

A. Traditional Semantic Layer

Semantic layers have been a critical component of BI systems since the early days of data warehousing. They serve as an abstraction layer that translates complex database schemas into business-friendly terms, enabling non-technical users to interact with data without needing to understand its underlying structure. Traditional semantic layers are typically built on relational models, where data is organized into tables, and relationships are defined through foreign keys [3].

Traditional semantic layers typically included several key components:

- Business entities: Representations of real-world business concepts like customers, products, and transactions
- Attributes: Properties of these entities that business users would recognize
- Metrics: Pre-defined calculations and aggregations relevant to business analysis
- Hierarchies: Relationships between entities, often representing organizational or classification structures
- Security rules: Access controls defining what data different user groups could see

While effective for structured data, this approach has limitations. Relational models struggle to represent complex, hierarchical, or interconnected data. Moreover, they are often rigid and inflexible, making it difficult to adapt to changing business requirements [4]. As organizations increasingly deal with unstructured and semi-structured data—such as text, images, and sensor data—the limitations of traditional semantic layers become even more pronounced.

The rise of big data and the proliferation of data sources have further exacerbated these challenges. Organizations now need to integrate data from diverse systems, including cloud platforms, IoT devices, and external APIs. Traditional semantic layers, designed for centralized and homogeneous environments, are ill-equipped to handle this complexity [5].

B. Limitations of Traditional Approaches

The constraints of traditional semantic layers have become more pronounced in today's complex data environments, manifesting in several key areas:

1) Structural Rigidity: Conventional semantic layers typically employ fixed, hierarchical data models that cannot easily accommodate the organic, interconnected nature of modern business contexts. This rigidity requires significant technical intervention to incorporate new data sources or adapt to changing business requirements, creating bottlenecks in the analytical process

2) Limited Context Awareness: Traditional models struggle to represent the rich contextual relationships that give data its full meaning. They excel at answering predetermined questions but falter when faced with exploratory, cross-domain inquiries that require understanding how entities relate across different business contexts.

3) Semantic Fragmentation: As organizations adopt specialized analytical tools for different departments and use cases, semantic definitions often become siloed and inconsistent across the enterprise. This fragmentation undermines the promise of a "single version of the truth" that semantic layers originally aimed to provide

4) Integration Challenges: The increasing diversity of data sources—from structured databases to semi-structured documents, streaming data, and external datasets—has strained traditional semantic layers, which were primarily designed for structured, internal data. Integrating these heterogeneous sources while maintaining coherent business meaning has proven exceptionally difficult.

5) Scalability Concerns: As data volumes grow exponentially, conventional semantic layer architectures face performance challenges, particularly when supporting complex, ad-hoc queries across large datasets. This limitation often forces organizations to make trade-offs between semantic richness and analytical performance.

The cumulative impact of these limitations has led many organizations to question the viability of traditional semantic layers in meeting their evolving analytical needs. While these approaches continue to serve important functions, they increasingly represent a constraint rather than an enabler for advanced analytics initiatives.

III. KNOWLEDGE GRAPHS: A PARADIGM SHIFT

Knowledge graphs offer a compelling alternative to traditional semantic layers. At their core, knowledge graphs are graph-based data structures that represent entities (nodes) and their relationships (edges). Unlike relational models, which rely on predefined schemas, knowledge graphs are inherently flexible and extensible. They can capture complex, multi-dimensional relationships and provide a rich, contextual representation of data [6].

A knowledge graph fundamentally differs from traditional data models in several important ways:

1) Entity-Relationship Centricity: Knowledge graphs place primary emphasis on entities (nodes) and the relationships (edges) between them, rather than on tables and joins. This approach allows for more natural representation of real-world concepts and their interconnections.

2) *Flexible Schema Evolution:* Unlike rigid relational schemas, knowledge graphs can accommodate new entity types and relationships without requiring comprehensive restructuring, enabling more agile adaptation to changing information needs.

3) *Semantic Expressiveness:* By incorporating ontologies and taxonomies, knowledge graphs can encode rich semantic meaning about entities and relationships, supporting more sophisticated reasoning and inference capabilities.

4) *Contextual Integration:* Knowledge graphs excel at integrating diverse information sources around common entities and concepts, preserving the contextual relationships that give data its full meaning and value.

The adoption of knowledge graph technologies in enterprise settings has accelerated in recent years, driven by maturing graph database platforms, the standardization of semantic web technologies (RDF, OWL, SPARQL), and the growing recognition of their value for addressing complex data integration challenges. Organizations across industries have implemented knowledge graphs for applications ranging from master data management to recommendation engines, fraud detection, and regulatory compliance.

One of the key advantages of knowledge graphs is their ability to incorporate semantics directly into the data model. By leveraging ontologies—formal representations of knowledge—knowledge graphs enable machines to understand the meaning of data. This semantic richness allows for more sophisticated querying, reasoning, and inference, paving the way for advanced analytics and AI-driven insights [7].

For example, consider a retail organization that wants to analyze customer behavior. A traditional semantic layer might represent customers, products, and transactions as separate tables, with relationships defined through foreign keys. In contrast, a knowledge graph could model customers, products, and transactions as interconnected entities, capturing not only transactional data but also contextual information such as customer preferences, product attributes, and market trends. This holistic representation enables deeper, more nuanced insights [8].

The convergence of these capabilities with the needs of modern business intelligence has set the stage for a fundamental reimagining of the semantic layer concept—one that leverages the flexibility and expressiveness of knowledge graphs while preserving the accessibility and governance benefits that made traditional semantic layers valuable.

IV. ARCHITECTURAL PRINCIPLES FOR KNOWLEDGE GRAPH DRIVEN SEMANTIC LAYERS

A. Conceptual Model

The integration of knowledge graph technologies with semantic layer concepts creates a new architectural paradigm that addresses many limitations of traditional approaches while introducing powerful new capabilities. A knowledge graph-powered semantic layer fundamentally reconceptualizes how business meaning is represented, accessed, and evolved within the enterprise data ecosystem.

At its core, this architecture consists of several interconnected components:

1) Ontological Fooundation

Ontologies provide the foundation for knowledge graphs by defining the concepts, relationships, and rules that govern a domain. An ontology-driven design ensures that the semantic layer is aligned with the business context and can evolve as new requirements emerge. For example, a healthcare organization might use an ontology to define concepts such as patients, diagnoses, and treatments, along with their relationships [9].

2) *Entity Resolution Layer:* This component handles the critical task of identifying and reconciling the same real-world entities across disparate data sources, creating unified entity representations that integrate information from multiple systems.

3) Graph Storage and Processing

Knowledge graphs excel at integrating data from diverse sources. By representing data as a graph, organizations can seamlessly combine structured and unstructured data, as well as internal and external data sources. This approach eliminates the need for complex ETL (extract, transform, load) processes and enables real-time data integration [10].

4) *Semantic Query Interface:* This layer translates business user queries into graph traversal operations, shielding users from the complexity of underlying graph query languages while preserving the rich relationship context that makes knowledge graphs valuable.

5) *Inference Engine:* Leveraging the semantic richness of the knowledge graph, this component enables automated reasoning over the data, uncovering implicit relationships and generating new insights that aren't explicitly encoded.

6) *Metadata Management:* Comprehensive metadata capture and management ensures that the provenance, quality, and usage patterns of data are tracked throughout the system, supporting governance and continuous improvement.

7) *API Ecosystem:* A robust API layer enables applications, analytics tools, and AI systems to interact with the knowledge graph-powered semantic layer in standardized ways, fostering integration across the enterprise technology landscape

This architecture represents a significant evolution from traditional semantic layers, which typically operated as static translation layers between databases and BI tools. Instead, it establishes a dynamic, continuously evolving knowledge fabric that adapts to new information and contexts while maintaining semantic coherence.

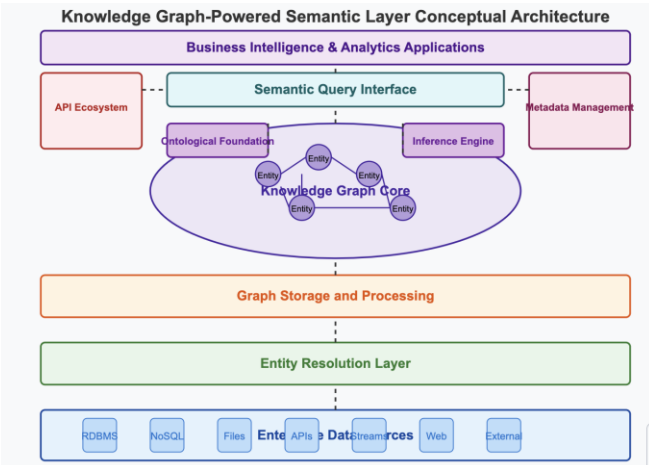

Fig.1 Conceptual Architecture of a knowledge graph powered Semantic Layer

B. Integration with Existing BI Ecosystems

For organizations with significant investments in existing business intelligence infrastructure, the transition to knowledge graph-powered semantic layers need not be disruptive. Instead, these new capabilities can be integrated with existing BI ecosystems through several complementary approaches:

1) *Federated Query Model:* Knowledge graph semantic layers can be implemented alongside traditional semantic layers, with federation mechanisms routing queries to the appropriate system based on the nature of the analytical need. This approach allows organizations to leverage knowledge graph capabilities for complex, relationship-centric analyses while maintaining existing systems for more straightforward reporting use cases.

2) *Progressive Enhancement:* Organizations can gradually enhance traditional semantic models with knowledge graph capabilities, starting with specific domains or use cases where the limitations of current approaches are most acute. This progressive approach minimizes disruption while demonstrating value.

3) *Semantic Virtualization:* Rather than physically migrating all data to graph structures, virtualization techniques can create knowledge graph views over existing data sources, preserving investments in data warehouses and lakes while adding semantic richness and relationship context.

4) *Bidirectional Synchronization:* Mechanisms for synchronizing semantic definitions between traditional BI semantic layers and knowledge graphs ensure consistency across platforms and prevent the creation of new semantic silos.

5) *Unified Governance Framework:* Extending data governance frameworks to encompass both traditional and knowledge graph semantic layers ensures that regardless of the technical implementation, organizational data assets remain subject to appropriate policies and controls.

This integration approach recognizes that the transition to knowledge graph-powered semantic layers represents an evolution rather than a replacement of existing BI investments. By thoughtfully integrating these capabilities, organizations can address the limitations of traditional approaches while leveraging their strengths and preserving the value they continue to deliver.

C. Key Differentiating Capabilities

Knowledge graph-powered semantic layers introduce several transformative capabilities that fundamentally change how organizations derive meaning and value from their data assets:

1) Dynamic Contextual Relevance: Unlike static semantic models, knowledge graphs can dynamically adapt the presentation and interpretation of data based on the user's context, query patterns, and the evolving network of relationships within the graph. This context-awareness enables more personalized and relevant analytical experiences.

*2) **Cross**-Domain Knowledge Integration:* By representing information as interconnected entities rather than domain-specific tables, knowledge graph semantic layers naturally bridge organizational silos, enabling analyses that span traditionally separate domains like finance, marketing, operations, and HR. This integration reveals insights at the intersections of these domains that would remain hidden in conventional approaches.

3) Semantic Reasoning and Inference: The formal ontological foundations of knowledge graphs enable automated reasoning over data, allowing systems to infer new facts, identify inconsistencies, and generate insights that aren't explicitly encoded. This capability amplifies the value of existing data by uncovering hidden patterns and relationships.

4) Natural Language Interaction: The entity-relationship structure of knowledge graphs aligns naturally with how humans think and communicate about business concepts, enabling more intuitive natural language interfaces for data exploration and analysis. Users can express complex analytical needs in business language rather than having to understand technical query constructs.

5) Evolutionary Schema Adaptation: Knowledge graphs can organically evolve as new entities, relationships, and properties emerge in the business environment. This evolutionary capability eliminates the rigid schema migrations that plague traditional systems, enabling more agile responses to changing analytical needs.

6) Knowledge Discovery Acceleration: By preserving the rich relationship context around data entities, knowledge graph semantic layers facilitate serendipitous discovery of relevant information that traditional query-based approaches would miss. This capability is particularly valuable for exploratory analyses where users may not know exactly what they're looking for.

These capabilities represent a fundamental shift in how semantic layers function within the enterprise—from passive translation layers to active knowledge systems that continuously evolve and adapt to changing business contexts while preserving semantic coherence across the organization.

V. IMPLEMENTATION STRATEGIES AND CHALLENGES

A. Ontology Development and Management

The foundation of any effective knowledge graph-powered semantic layer is a well-designed ontology that captures the concepts, relationships, and rules relevant to the organization's business domains. Developing and managing these ontologies presents both significant opportunities and challenges:

1) Balancing Specificity and Flexibility: Organizations must strike a careful balance in ontology design, creating models specific enough to provide meaningful semantic structure while remaining flexible enough to accommodate evolving business concepts and use cases. Overly rigid ontologies become maintenance burdens, while excessively loose ones fail to provide adequate semantic guidance [14].

2) Collaborative Development Approaches: Effective ontologies require input from both domain experts who understand the business concepts and knowledge engineers who can formalize these concepts in machine-readable formats. Establishing collaborative workflows that bridge these perspectives is essential for developing ontologies that are both technically sound and business-relevant.

3) Ontology Lifecycle Management: As business concepts and relationships evolve, ontologies must adapt accordingly. Implementing governance processes for ontology versioning, change management, and impact assessment ensures that semantic models remain accurate and relevant over time without disrupting dependent systems.

4) Standards and Interoperability: Leveraging established semantic web standards (RDF, OWL, SKOS) and industry-specific ontologies where available promotes interoperability and reduces development effort.

However, organizations must carefully adapt these standards to their specific contexts while maintaining compliance with core principles.

5) *Ontology Validation and Testing:* Formal validation techniques, including logical consistency checking and test-case validation, help ensure that ontologies accurately represent the intended business semantics. Regular validation against real-world data and use cases identifies gaps and inconsistencies before they impact downstream analytics.

Organizations that successfully navigate these challenges create ontologies that serve as living semantic foundations, continuously evolving while maintaining the semantic coherence necessary for reliable analytics and decision support.

B. Data Integration and Entity Resolution

Bringing diverse data sources into a unified knowledge graph while ensuring that entities are properly identified and linked across sources represents one of the most significant technical challenges in implementing knowledge graph semantic layers:

1) *Source Data Diversity:* Organizations typically need to integrate structured data from operational systems, semi-structured data from documents and logs, and increasingly, unstructured text and multimedia content. Each source type requires specialized extraction and transformation approaches to identify entities and relationships.

2) *Identity Resolution at Scale:* Accurately determining when references across different sources point to the same real-world entities becomes exponentially more complex as the number of sources and entities grows. Probabilistic matching algorithms, machine learning approaches, and human-in-the-loop verification workflows all play important roles in addressing this challenge [15].

3) *Temporal Dimension Management:* Business entities and their relationships evolve over time, requiring semantic layers to track temporal dimensions efficiently. Capturing both transaction time (when data was recorded) and valid time (when facts were true in the real world) enables accurate point-in-time and longitudinal analyses.

4) *Incremental Graph Building:* Rather than attempting comprehensive knowledge graph creation in a single project, successful implementations typically adopt incremental approaches, starting with high-value domains and progressively expanding coverage. This approach delivers business value earlier while refining integration patterns based on experience.

5) *Feedback Loops for Continuous Improvement:* Establishing mechanisms for users to identify and correct entity resolution errors creates virtuous cycles that progressively improve data quality while engaging business stakeholders in the knowledge-building process.

Organizations that effectively address these data integration challenges create unified knowledge assets that preserve the rich context around business entities while harmonizing information from disparate sources—a foundational capability for advanced analytics and AI applications.

C. Governance and Security Implications

The flexible, interconnected nature of knowledge graphs introduces both new opportunities and challenges for data governance and security:

1) *Granular Access Control:* Knowledge graph structures enable more nuanced access control models that can restrict access based on entity types, relationship patterns, and attribute values rather than simply on tables or columns. This granularity better aligns security with business concepts but requires careful implementation to avoid performance penalties.

2) *Provenance Tracking:* Capturing and preserving the lineage of entities and relationships throughout the knowledge graph enables transparency about how information was derived and transformed. This provenance information supports auditability while helping users assess the reliability of insights.

3) *Privacy-Preserving Knowledge Representation:* Organizations must implement privacy controls that prevent unintended exposure of sensitive information through relationship analysis and inference. Techniques such as differential privacy and relationship filtering help mitigate these risks while preserving analytical utility.

4) *Semantic Policy Definition:* Expressing governance policies in terms of the same ontologies that structure the knowledge graph creates alignment between technical implementation and governance intent. This semantic approach to policy definition enables more consistent enforcement across heterogeneous data sources.

5) *Active Compliance Monitoring:* The inferential capabilities of knowledge graphs can be leveraged to actively monitor compliance with data governance policies, automatically detecting potential violations and policy conflicts as the graph evolves.

Organizations that thoughtfully address these governance considerations create trustworthy semantic foundations that enable broader data democratization while ensuring appropriate protection of sensitive information and compliance with regulatory requirements.

VI. FUTURE DIRECTIONS AND RESEARCH OPPORTUNITIES

As knowledge graph-powered semantic layers continue to evolve, several promising directions for future development and research emerge:

1) *Automated Ontology Learning:* Techniques for automatically extracting and evolving ontologies from existing data assets and user interactions represent a critical frontier for scaling semantic layer implementations. Machine learning approaches that can identify emerging concepts and relationship patterns with minimal human supervision will significantly accelerate knowledge graph adoption

2) *Multimodal Knowledge Representation:* Extending knowledge graph models to seamlessly incorporate multimodal information—including text, images, video, and sensor data—will enable more comprehensive semantic representations of business realities. These multimodal capabilities will be particularly valuable for industries where visual and spatial information plays critical roles.

3) *Cognitive Architecture Integration:* Deeper integration between knowledge graph semantic layers and emerging cognitive architectures, including large language models and neural-symbolic systems, presents opportunities for creating analytics systems with more human-like understanding of business contexts and natural communication capabilities.

4) *Federated Knowledge Networks:* As organizations increasingly collaborate within business ecosystems, techniques for federating knowledge graphs across organizational boundaries while preserving security and sovereignty will become essential. These federated approaches will enable richer analytics across supply chains, partner networks, and industry consortia.

5) *Real-Time Knowledge Evolution:* Methods for continuously evolving knowledge graphs in response to streaming data and events represent an important frontier for organizations operating in dynamic environments. These capabilities will be particularly valuable for use cases requiring immediate responses to changing conditions, from financial trading to supply chain management.

These future directions suggest that knowledge graph-powered semantic layers will continue to advance rapidly, further transforming how organizations understand, analyze, and derive value from their information assets. Organizations that establish the foundational capabilities described in this paper will be well-positioned to adopt these emerging innovations as they mature.

VII. CONCLUSION

The integration of knowledge graph technologies with semantic layer concepts represents a fundamental reimagining of how businesses organize, access, and derive value from their data assets. This fusion creates a new paradigm for business intelligence—one that preserves the accessibility benefits of traditional semantic layers while dramatically expanding their capabilities to address the complexities of modern data environments.

Knowledge graph-powered semantic layers offer several transformative advantages over traditional approaches. They provide more natural representations of business contexts through entity-relationship models that mirror how humans understand the world. They enable more dynamic and adaptive analytics experiences that evolve with changing business needs rather than constraining them. They facilitate cross-domain integration that reveals insights at the intersections of traditionally siloed business functions. Perhaps most importantly, they establish the semantic foundation necessary for truly intelligent systems that can reason about business realities rather than simply reporting on them.

While implementing these capabilities presents significant technical and organizational challenges, the emerging patterns and practices described in this paper provide practical pathways for organizations to begin this transformation. By taking incremental approaches that deliver value at each stage while building toward a comprehensive knowledge architecture, organizations can manage risk while progressively enhancing their analytical capabilities.

As artificial intelligence continues to transform business operations and decision-making, the importance of rich semantic foundations will only increase. Knowledge graph-powered semantic layers provide the contextual awareness, relationship understanding, and reasoning capabilities that AI systems require to deliver truly intelligent and trustworthy business insights. Organizations that establish these foundations today will be better positioned to leverage the full potential of AI in the future.

The journey toward knowledge graph-powered semantic layers represents more than a technical evolution—it signals a fundamental shift in how organizations conceptualize and interact with their information assets. In this new paradigm, data transcends its role as a static resource to be queried and becomes a dynamic knowledge fabric that actively supports understanding, discovery, and decision-making across the enterprise.

REFERENCES

[1] F. van Harmelen, V. Lifschitz, and B. Porter, *Handbook of Knowledge Representation*. Elsevier, 2008. https://www.sciencedirect.com/bookseries/foundations-of-artificial-intelligence/vol/3/suppl/C

[2] D. Fensel et al., *Knowledge Graphs: Methodology, Tools and Selected Use Cases*. Springer, 2020. https://link.springer.com/book/10.1007/978-3-030-37439-6

[3] A. Hogan et al., "Knowledge Graphs," *ACM Computing Surveys*, vol. 54, no. 4, pp. 1–37, 2021. https://dl.acm.org/doi/10.1145/3447772

[4] P. Hitzler et al., *Semantic Web Technologies: Trends and Research in Ontology-Based Systems.* Wiley, 2012.

[5] M. Ehrlinger and W. Wöß, "Towards a Definition of Knowledge Graphs," *SEMANTiCS 2016*, 2016. https://ceur-ws.org/Vol-1695/paper4.pdf

[6] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, 2001. https://www.scientificamerican.com/article/the-semantic-web/

[7] G. Klyne and J. Carroll, *Resource Description Framework (RDF): Concepts and Abstract Syntax*. W3C, 2004. https://www.w3.org/TR/rdf-concepts/

[8] D. Brickley and R. Guha, *RDF Schema 1.1*. W3C, 2014. https://www.w3.org/TR/rdf-schema/

[9] S. Harris and A. Seaborne, *SPARQL 1.1 Query Language*. W3C, 2013. https://www.w3.org/TR/sparql11-query/

[10] M. Uschold and M. Gruninger, "Ontologies: Principles, Methods, and Applications," *Knowledge Engineering Review*, vol. 11, no. 2, pp. 93–136, 1996.

[11] J. Lehmann et al., "DBpedia: A Large-Scale, Multilingual Knowledge Base Extracted from Wikipedia," *Semantic Web Journal*, vol. 6, no. 2, pp. 167–195, 2015. https://content.iospress.com/articles/semantic-web/sw203

[12] A. Singhal, "Introducing the Knowledge Graph: Things, Not Strings," *Google Blog*, 2012. https://blog.google/products/search/introducing-knowledge-graph-things-not/

[13] M. Nickel, K. Murphy, V. Tresp, and E. Gabrilovich, "A Review of Relational Machine Learning for Knowledge Graphs," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 11–33, 2016. https://ieeexplore.ieee.org/document/7358050

[14] L. Ehrlinger and W. Wöß, "Towards a Definition of Knowledge Graphs," *SEMANTiCS 2016*, 2016. https://ceur-ws.org/Vol-1695/paper4.pdf

[15] D. Calvanese et al., "Ontology-Based Data Integration: A Survey," *IJCAI 2018*, 2018. https://www.ijcai.org/proceedings/2018/0777.pdf