Towards Semantically Enriched Knowledge Graph Embedding through Ontology Axiom Integration

Ivan Diliso[1,*]

[1]University of Bari Aldo Moro, Bari, Italy

Abstract

Knowledge Graph Embedding models are widely used to support predictive and reasoning tasks involving structured data. However, they often fail to capture the full semantic richness of formal ontologies, particularly those expressed in Description Logics. This PhD research project investigates how ontological knowledge can be systematically incorporated into Knowledge Graph Embedding models to improve their semantic awareness, generalization, and interpretability. The work is application-driven, with a specific focus on the railway domain, which presents complex, hierarchical, and mission-critical data scenarios. A modular ontology-integration framework is proposed, incorporating ontological knowledge into negative sampling, loss constraints, and embedding alignment via geometric representations of Description Logic axioms. Additionally, the project addresses an area that has been largely overlooked in the literature: evaluating the semantic validity of Knowledge Graph Embedding models. The ultimate objective is to develop a reusable, domain-agnostic methodology for constructing ontology-aware machine learning methods that bridge the gap between symbolic knowledge and statistical learning, with relevance in both academia and industry.

Keywords

Knowledge Graph Embedding, Ontology, Ontology Injection, Description Logics, Railway Engineering

1. Introduction and Research Questions

The rapid growth of both structured and unstructured data has intensified the need for advanced techniques for representing, integrating and reasoning about complex knowledge. Knowledge graphs (KGs) have emerged as a robust solution for representing entities and their relationships in graph form. They facilitate the integration of diverse data sources, providing a unified and interconnected perspective on knowledge [1]. KGs are pivotal in enhancing data interoperability and are instrumental in tasks such as semantic search, link prediction, and fault detection. They offer a unified, contextualized view of heterogeneous data sources.

To further enhance the expressive power and reasoning capabilities over KGs, ontologies provide formal, logic-based vocabularies that define domain concepts, relationships, hierarchies, properties, and constraints. Ontologies promote semantic interoperability across systems and enable structured knowledge to be injected into machine learning models [2]. Integrating such ontological knowledge into knowledge graph embedding (KGE) models has shown promise in enhancing the performance of predictive tasks [3].

Yet current approaches often underutilize the full expressiveness of Description Logics (DL), which underpin most formal ontologies [4]. Current literature on KGE models primarily focuses on surface-level graph structures or limited aspects of ontologies, such as class hierarchies and domain and range constraints [5, 6], often ignoring richer semantic features such as functional and role properties, disjointness, and the explicit negation of concepts. Consequently, these models struggle with deep reasoning and generalization [7, 8], particularly in domains that require high semantic fidelity.

RuleML+RR'25: Companion Proceedings of the 9th International Joint Conference on Rules and Reasoning, September 22–24, 2025, Istanbul, Turkiye

*Supervised by Nicola Fanizzi and Claudia d'Amato
i.diliso1@phd.uniba.it (I. Diliso)
https://ivandiliso.github.io/ (I. Diliso)
0009-0007-2942-202X (I. Diliso)

This research aims to systematically explore the injection of ontological knowledge into KGE models, addressing both the theoretical and the practical challenges. Specifically, this research has a focus on the railway engineering domain, where real-world challenges around safety, complexity, and semantic consistency make it an ideal application for ontology-enhanced KGE methods.

The modern railway system has evolved with digital technologies and automation, shifting from manual operations to intelligent, autonomous control. Its complex, interconnected environment poses diverse and unpredictable risks, increasing the need for efficient formalization of domain knowledge, automatic maintenance and diagnosis, and system compliance with implant specifications [9]. Systems must also comply with the "mission critical" domain, which is defined as "one where a hazard can degrade or prevent the successful completion of an intended operation" [10].

This highlights the need for applications able to systematically formalize the domain knowledge (formalization), reason and produce insights and results on the produced knowledge (exploitation), and be able to support the human user in their operations, offering interpretation of its results (explainability). The field of railway engineering is a natural application area for this work. Knowledge graphs are a powerful way of representing this rich structure, and ontologies can formalize system-level rules and constraints. Nevertheless, integrating ontological knowledge into machine learning (ML) systems that support railway operations remains largely unexplored.

The aim of this research is to bridge the gap between symbolic reasoning and statistical learning by incorporating ontological semantics at various stages of the KGE pipeline, such as negative sampling, loss functions, and embedding geometry. My overarching goal is to enhance the semantic awareness, reasoning accuracy, and interpretability of ML systems operating in real-world, safety-critical contexts. To address these challenges and guide the development of a semantically enriched KGE framework, this research is structured around the following key questions:

- **RQ1**: *How can formal ontological knowledge be systematically injected into KGE models without compromising scalability or performance?*
- **RQ2**: *What impact do ontological constraints have on the semantic consistency and predictive accuracy of KGE models?*
- **RQ3**: *How can we design evaluation methodologies that go beyond ranking metrics to measure the semantic validity of predictions made by ontology-aware KGE models?*
- **RQ4**: *How can learned representations of ontologies and knowledge graphs support domain experts in complex, mission-critical applications such as railway system analysis and planning?*

2. State Of The Art

The core of this PhD research covers a wide range of areas of interest, ranging from: learning knowledge graph representations; formal ontologies embedding in their different degrees of expressiveness; methodologies able to integrate both structural and ontological knowledge in low dimensional embedding for automatic reasoning; as well as the in-depth evaluation of predictive models. This section will review the relevant literature on these research areas, from knowledge graph embeddings to their domain-specific application in railway engineering.

2.1. Knowledge Graph Embedding

Knowledge graph embedding (KGE), also referred to as multi-relational learning, is a type of representational learning that seeks to embed KG components (entities and relations) into continuous vector spaces. This facilitates the use and exploitation of knowledge in downstream applications [1]. KGE models fall into one of several categories: distance-based, where relations are seen as a vector transformation function that maps head entities closer to tail entities; semantic matching, which scores triplets based on a semantic similarity score; and graph neural network (GNN) approaches, which gather knowledge from the neighborhood of a target node and aggregate multi-hop structural information, placing entities with similar neighbors together in the embedding space [11].

While classical distance-based approaches are simple and interpretable, they typically assume fixed geometric transformations that cannot model intricate dependencies and hierarchical structures in KG [12]. More sophisticated methods with geometric or probabilistic interpretations have therefore emerged to address these limitations, using advanced representations that demonstrate the model's ability to embed relational properties and logical rules [13, 14]. Using neural network-based approaches in KGE offers significant advantages, particularly in inductive settings, due to their ability to generalize beyond the initial graph on which they were trained [15]. These graph convolution-based approaches generalize the message-passing algorithm in the multi-relational setting, enabling their application to KGs [16].

The popularization of the attention mechanism in the field of natural language processing [17], led to the development of several approaches that adapt this technique to graph data by assigning different importance to nodes within the neighborhood [18]. All the attention-based approaches differ in the level of the graph on which the attention mechanism is applied: whether it is on triplet level [19], query [20], relations and entities [21], meta-paths [22] or classes. Specifically in [23] the authors incorporate the self-attention mechanism on domain and range properties (approximated using the local closed word assumption on node neighbors), providing a first example of how ontological knowledge can be used to enrich knowledge representation.

2.2. Ontology Embedding

Ontologies are fundamental to the Semantic Web [24], enabling the representation, reasoning, and sharing of knowledge across domains. An ontology is a formal explicit specification of a shared conceptualization of a domain [4], written in formal languages such as Web Ontology Languages (OWL) and rooted in Description Logics (DL), the formal language that provides the mathematical underpinning for ontology design and reasoning [25]. The aim of ontology embedding is to project axioms into a low-dimensional embedding space while preserving the logical properties of the knowledge base [26]. This generates an approximate model of the knowledge base in the sense of an interpretation of the DL, which is often formalized in the literature as either $\mathcal{ALC}$ or $\mathcal{EL}^{++}$. One solution is to treat the ontology as a knowledge graph using classical KGE algorithms. However, this approach results in a significant loss of information, discarding the richness of language [27]. In [7] the authors provide an in-depth analysis of the projection capabilities of ontology embedding models. Another solution is to use geometric models that project axioms into a topological space. These models are able to preserve the semantics of concept and role descriptions. Examples include spheres, boxes [28], lattices [26] or cones [29].

2.3. Joint Ontology and Knowledge Graph Embedding

The core of this PhD research lies in jointly learning the structural-relational information of KGs and the rich semantics of DL. Although the usefulness of the SOTA KGE models in link prediction tasks has been demonstrated across various datasets, these methods often fail to capture even simple DL rules [30]. Typically, these models ignore DL semantics, focusing their learning on surface-level patterns and primarily representing the ABox. They disregard the richer schema-level information contained in the TBox [31], often focusing only on hierarchical type information [32]. Some approaches focus specifically on concept subsumption (e.g., node types and hierarchies) and view the KG as having two levels: an upper level that describes concept subsumption relationships and a lower level that describes relationships between entities and their associated concepts [5]. Others inject and leverage ontological information during the negative sampling phase. For example, soft-type constraints are used to guide negative sample generation [33] (triples assumed to be false statements). Alternatively, in [34], the authors propose an iterative negative sampling approach that uses a reasoner to gather inconsistent predictions that populate the negative sample set for the next training phase. Another approach involves using relation properties to constrain the score function in KGE [3], accounting for inverse and equivalent relations on both classes and properties, as well as their hierarchical structure.

In [6] the proposed method further injects knowledge by using a reasoner to produce negative samples that exploit ontology axioms, following a similar formulation to the methods described earlier. A common oversight in literature is the consideration of axioms and hierarchies over relationships, with few works incorporating both concept and relation hierarchies, as discussed in [2].

2.4. Evaluation of Ontology-Enriched Embeddings

An underexplored area in the literature is the in-depth evaluation of KGE models that incorporate ontological knowledge (i.e., constraints derived from axioms such as domain, range, or class hierarchies). Most existing evaluation practices, particularly in the context of link prediction tasks (e.g., predicting the most probable tail entity t' given a triple $\langle h, r, ?\rangle$, or head prediction in the reverse form), rely on metrics borrowed from information retrieval. These include ranking-based metrics such as Mean Reciprocal Rank (MRR) or Hits@K, which assess how well the model ranks the correct entity among all possible candidates based on its internal scoring function. Although these metrics accurately measure the model's ability to retrieve the correct entity, they do not capture the quality of the learned semantics. In other words, they fail to evaluate whether the predicted entities are semantically valid in relation to the ontology.

For example, [35] introduces a semantic-aware evaluation metric that measures the proportion of semantically valid triples among the top-K predictions. This metric uses compatibility constraints derived from domain and range axioms, as well as class hierarchies. For instance, predicting '*dog*' as the object for the '*hasOccupation*' relation would be invalid if the relation's range expects instances from '*Profession*'. Similarly, [34] proposes a metric that counts the number of semantically inconsistent predictions among the top-K candidates.

Despite these contributions, there has been little large-scale, systematic evaluation of KGE models that integrate ontological axioms using metrics specifically designed to assess semantic validity. This presents a critical gap: in scenarios where standard ranking metrics yield comparable results, ontology-enriched models should, in principle, outperform traditional models when evaluated using semantic-aware metrics. Addressing this gap is a key objective of our work, as it is essential for a full understanding of, and demonstration of the benefits of, ontology-aware KGE.

2.5. Applications of Ontologies and KG to Railway Engineering

This research project focuses on incorporating ontological knowledge into machine learning models for knowledge graphs, particularly in the field of railway engineering. Knowledge graphs and formal ontologies provide a robust basis for handling the complexity and diversity of information in this field, improving data organization, facilitating semantic integration, and supporting sophisticated analytical capabilities via learning algorithms [36].

Most existing research in this area concentrates on extracting schema-level and instance-level information from textual specification documents, since many system rules are not yet formalized in machine-readable languages. Despite this, the explicit use of formal ontological knowledge in the railway domain remains largely underexplored [36], even though ontologies have already demonstrated significant value in Industry 4.0 contexts, where there is high demand for formal, semantically rich, and interoperable data infrastructure [37, 38].

Examples of the potential impact are beginning to emerge. For instance, [39] introduces a two-view knowledge graph to support engineers in path alignment tasks, showing how integrating knowledge graphs enables systems to adapt dynamically to new knowledge. More broadly, railway applications such as traffic management, safety rule compliance, predictive maintenance, and infrastructure interoperability could benefit from ontological models that make implicit domain knowledge explicit and machine-actionable. However, systematic approaches that combine such ontological representations with learning-based methods remain scarce, highlighting a research gap that this project aims to address.

3. Methodology

3.1. Central Research Problem

Despite considerable progress in both KGE and ontology embedding, existing methods remain limited in their ability to jointly capture the structural patterns of KGs and the formal semantics provided by ontologies. KGE models have achieved strong performance in link prediction tasks, but they largely ignore schema-level information such as class hierarchies, relation properties, domain and range constraints, or explicit negation, focusing instead on surface-level statistical patterns. Ontology embedding approaches, by contrast, are designed to preserve logical semantics but are typically developed and evaluated in isolation, without integration into predictive models. This disconnect has resulted in a lack of systematic frameworks that enrich KGE with ontological knowledge while maintaining compatibility with existing architectures and evaluation practices.

Addressing this gap is particularly important in safety-critical and data-intensive domains such as railway engineering. Railway systems involve heterogeneous data sources, strict safety and compliance rules, and evolving operational constraints, all of which require models that are not only predictive but also semantically consistent with domain knowledge.

Developing approaches that combine structural learning with ontological semantics is therefore essential to improve both the accuracy and trustworthiness of predictive models while enabling their practical application in domains where semantic validity is as critical as raw performance.

3.2. Proposed Approach

In this work, the term "*ontology injection*" refers to the systematic process of aligning a KGE model with the semantics of an ontology. By ontology injection we mean not only enriching the model with additional features but also ensuring that the learned representations explicitly respect ontological rules and axioms, such as class hierarchies, relation properties, domain and range constraints, or explicit negations. In practice, injection entails incorporating ontological knowledge directly into the learning process so that the resulting embeddings are both topologically informed by the graph structure and semantically constrained by the ontology.

The first step is to exploit the key properties of ontologies, either by injecting them when they are explicitly available or by inferring them from the graph structure when schema-level information is missing. This enables the discovery of new, exploitable knowledge while grounding the embeddings in a formally defined semantic framework. Unlike purely structural models, ontology-injected embeddings are expected to reuse existing domain knowledge rather than relearning patterns already encoded at the schema level, leading to improvements in both accuracy and reliability.

A central novelty of this research as opposed to current literature, lies in its focus on developing a generic ontology injection module, rather than a single ontology-enriched KGE model. The goal is not to design yet another specialized architecture, but to provide a flexible component that can be integrated with any KGE method to enhance semantic validity. This modular perspective is what motivates the use of the term "*ontology injection*": ontological semantics are injected into existing models in a way that is architecture-agnostic, extensible, and reproducible.

In line with these objectives, this research investigates three complementary strategies: injection during negative sampling, where ontological axioms are used to guide the generation of corrupted triples, avoiding semantically invalid samples and thus producing more informative training pairs (positive triples and corrupted triples); injection into embedding spaces, where axioms are translated into geometric constraints; and injection via loss constraints, where the model's optimization objective is augmented with penalties or regularizers derived from logical axioms, explicitly discouraging inconsistent predictions.

Together, these strategies form a structured framework for integrating ontological knowledge into KGE models. The overarching goal is to move towards semantically grounded representations that are consistent, interpretable, and applicable in safety-critical domains such as railway engineering. A conceptual overview of this framework is presented in Figure 1.

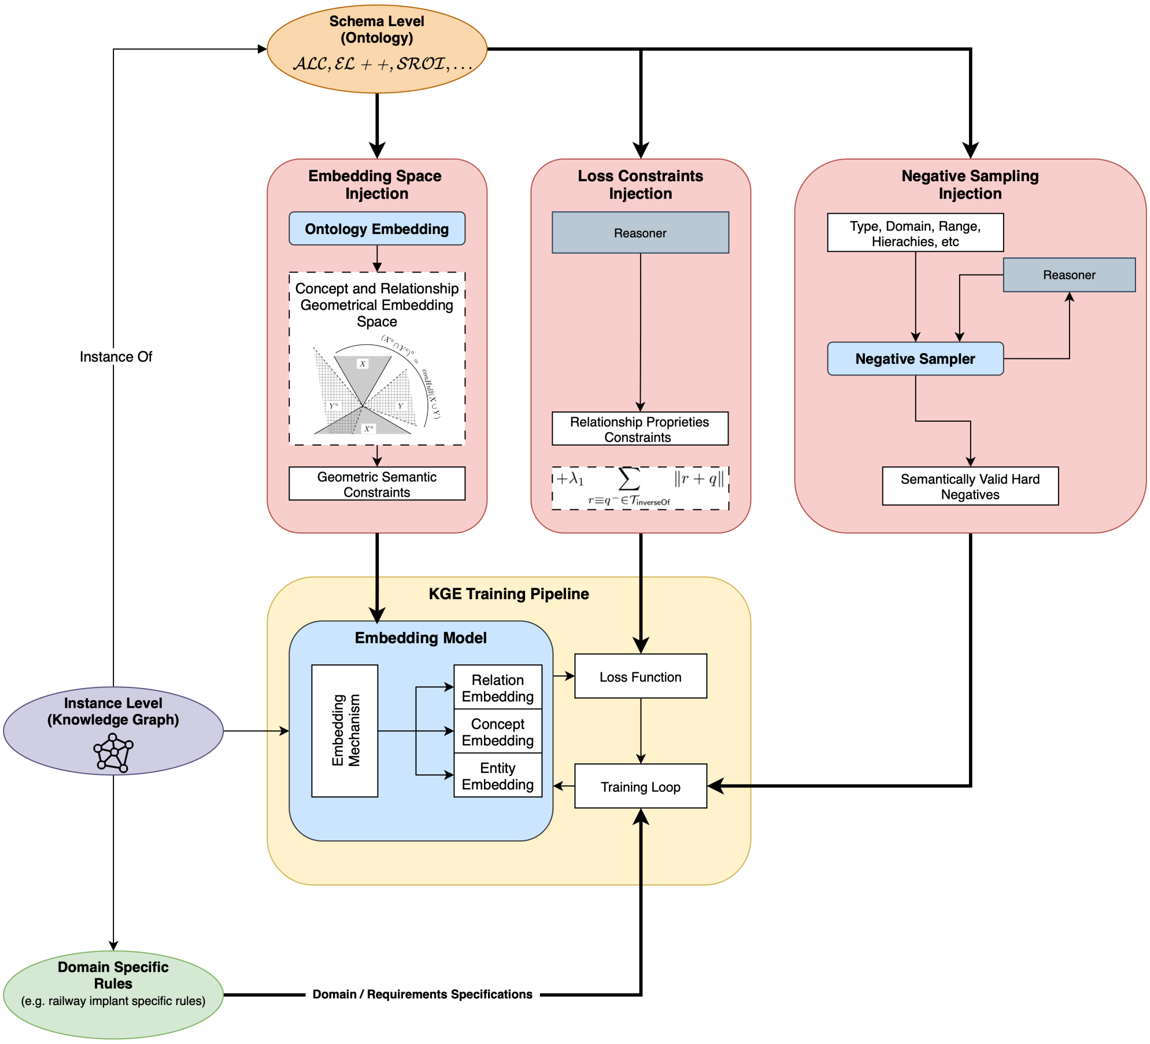

Figure 1: General conceptual framework for KGE ontology injection, encompassing the areas of embedding space injection, loss constraints, and negative sampling (from left to right). The bigger arrows indicate the source of additional schema/rule-based knowledge in the KGE model.

3.2.1. Injection in Negative Sampling

A first approach to leverage ontologies during the negative sampling phase is to generate more semantically accurate negative samples for training KGE models. Although the use of type constraints and reasoning in negative sampling has been explored in previous studies and proven effective [33, 34], existing work still lacks comprehensive approaches that exploit the full expressive power of ontologies. This highlights an open research direction focused on developing novel sampling strategies capable of capturing the complete range of available background knowledge. This strategy aligns closely with adversarial negative sampling approaches, where an auxiliary model is used to predict harder negatives. Rather than directly injecting ontological constraints, an adversarial negative sampler can leverage a KGE or ontology embedding model that has been specifically trained to capture ontological semantics. Assuming this model accurately reflects the semantic structure of the domain, it can guide the sampling process to produce more informative and challenging negative samples for training. The conceptual schema of this proposed implementation is provided in Figure 2. Adversarial negative samplers based on ontology-injected models are still lacking in the literature and present a promising research direction. However, this approach relies on the availability of ontology-injected models, which are developed through the other research directions explored in this PhD study.

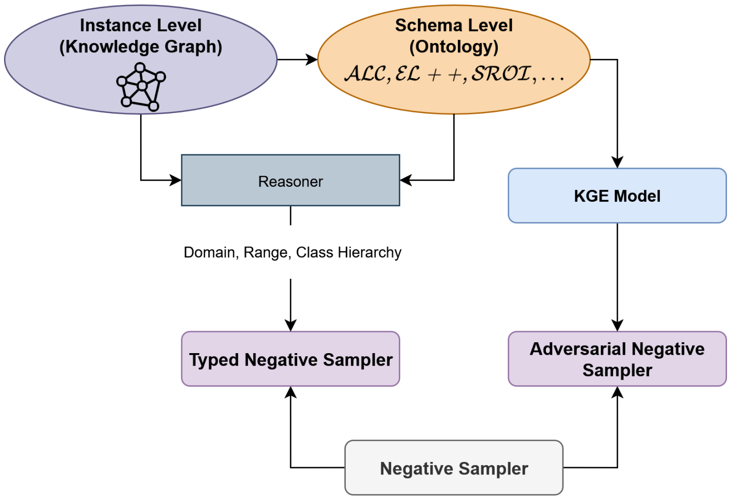

Figure 2: Conceptual schema of ontology injection in negative sampling, leveraging ontological properties or ontology-injected adversarial embedding models.

3.2.2. Injection in Embedding Spaces

The geometric interpretation of DL in vector spaces has been proven to be an efficient and sound method of representing symbolic knowledge in numerical form [29]. These approaches make it possible to embed both defined concepts and the axioms that describe their relationships (including explicit negation, often ignored in SOTA KGE models) into a structured vector space.

In the proposed approach, this process is organized into two phases. First, an "*ontology embedding space*" is constructed in which logical axioms, class hierarchies, role properties, and constraints are represented in a machine-readable vector form. This space serves as a semantic reference model that encodes the formal structure of the ontology. Second, this ontology-informed space is used to constrain and guide the training of KGE models: entity and relation embeddings learned from the ABox are projected into or regularized against the ontology space, ensuring that predictions remain consistent with schema-level semantics. In this way, the ontology embedding acts as a supervisory signal that aligns structural learning with domain rules.

To further refine this process, when multiple ontological constraints apply simultaneously to the same embedding (e.g., domain and range restrictions in addition to hierarchical constraints), an attention mechanism (called in this work "*ontological attention*") can be introduced to dynamically weight their relative influence, aiming to minimize prediction inconsistencies. This allows the model to prioritize the most relevant constraints in context, rather than enforcing them uniformly. The resulting integration balances symbolic fidelity with flexibility, reducing semantic violations without over-constraining the embeddings. Overall, this coupling of ontology embedding spaces with loss-based constraints provides a principled method of ontology injection, transforming abstract DL semantics into operational guidance for KGE models and enabling consistent, semantically grounded representations across both ABox and TBox levels.

3.2.3. Injection in Loss Constraints

Loss constraints have been shown to enrich the learned model semantics by leveraging relational properties [3]. While these approaches have proven useful, a general framework for broadly applicable ontology injection is still missing, which hinders large-scale evaluation and comparison. Moreover, the application of loss constraints to neural-based models remains relatively understudied. I intend to explore this topic, particularly in the context of GNNs, which are well-suited for generalizing over KG structure and can operate in inductive settings (making predictions even on previously unseen relationships between entities). A promising research direction is the development of a model that can adapt to new knowledge and rules, ensuring that the learned embeddings align with underlying semantic constraints and relationships. This alignment is critical for enabling generalization and inductive reasoning, thus supporting broader applicability across heterogeneous domains. Furthermore, my research is closely tied to a specific application domain: the railway ecosystem. This domain is characterized by schema-level knowledge (capturing general semantics), instance-level knowledge (e.g., implants), and implant-specific rules, known as "*domain specifications*" or "*domain requirements*". While a general model is necessary to capture the core semantics of the domain, it is equally important to optimize it for context-specific rules. Loss or structural constraints in KGE models offer a direct mechanism to fine-tune and adapt a general model to these domain-specific requirements. Unlike prior work, which often focuses on a narrow subset of ontological features, this approach also aims to leverage a broader range of axioms and semantics, incorporating the ontology embedding spaces developed in earlier phases.

3.3. Goals, Motivations, and Challenges

A key obstacle in advancing ontology-aware KGE methods is the lack of standardized, reproducible resources. Current datasets, ontologies, and implementations are fragmented across ad hoc efforts, forcing researchers to manually formalize inputs and build custom pipelines. This fragmentation slows progress, raises the entry barrier, and prevents systematic evaluation and comparison of methods. A further challenge lies in bridging symbolic and neural paradigms. Ontology reasoners are typically JVM-based, while KGE models rely on GPU-accelerated tensor frameworks. Efficiently integrating these ecosystems requires careful design to preserve both scalability and semantic fidelity.

The goal of this research is to develop a general, extensible framework for incorporating ontological knowledge into KGE models. The framework is designed to be domain-agnostic yet adaptable, supporting both generic use cases and domain-specific applications. Railway engineering serves as a key case study, where heterogeneous data sources (e.g., interlocking systems, track infrastructure, sensor networks) and strict safety rules highlight the need for models that are not only predictive but also semantically consistent. Ultimately, the ambition is to deliver a flexible and reusable solution that unifies symbolic reasoning with neural learning. Such a framework would lower technical barriers, enable reproducible research, and support the development of ontology-enriched KGE pipelines across both academic and industrial contexts.

3.4. Evaluation Plan and Success Criteria

Datasets. Evaluation will be carried out on a mix of standard benchmark datasets and ontology-enriched datasets. Classical benchmarks such as WN18-RR, FB15K-237, NELL, YAGO3-10, DBPedia50K, and DBPedia100K will be used to ensure comparability with prior work in the KGE literature. However, these datasets primarily consist of raw triples (ABox assertions) and generally lack schema-level information (TBox axioms), limiting their ability to test semantic validity. To address this, evaluation will also include datasets that combine triples with subsets of ontologies, such as YAGO4-20, YAGO39K, DBPedia39K, and DBPediaYAGO, which trade off computational feasibility with rich schema information. In addition, ad hoc datasets in OWL format will be developed to support controlled experiments on specific ontological constructs. Finally, a domain-specific case study based on railway engineering data will be conducted. This will provide a setting where industry-standard rules and compliance requirements are encoded in ontologies, enabling an assessment of the practical utility of the proposed methods in a safety-critical environment.

Metrics. Since the primary tasks are knowledge graph completion (link prediction) and triple classification, evaluation will adopt the standard metrics used in the field. For link prediction, these include Mean Reciprocal Rank (MRR) and Hits@K, while for triple classification, Precision, Recall, Accuracy, F1-score, and Area Under the Curve (AUC) will be reported. While these metrics capture general predictive performance, they do not account for semantic validity, i.e., whether predictions respect ontological constraints. To address this, semantic-aware metrics will also be applied. For link prediction, *Sem@K* [35] will measure the proportion of semantically consistent triples in the top-K predictions, while *Inc@K* [34] will count inconsistent triples among the top-K. For triple classification, particular emphasis will be placed on detecting cases where semantically inconsistent triples are predicted as true, as these are especially problematic in domains such as railway engineering. A further contribution of this research will be the design and evaluation of additional semantic-aware metrics, extending current approaches to more comprehensively capture the consistency of predictions with respect to ontologies.

Success Criteria and Expected Results. The primary measure of success is the ability of ontology-enriched KGE models to consistently produce semantically valid predictions, thereby reducing the frequency of inconsistent outputs that violate ontological constraints. While improvements in classical predictive accuracy remain relevant, they are considered secondary to gains in semantic reliability. Concretely, models will be evaluated on three dimensions: maintaining or surpassing baseline performance on established benchmarks to ensure competitiveness, achieving measurable improvements on semantic-aware metrics such as higher Sem@K and lower Inc@K, which directly quantify the reduction of inconsistent predictions, and demonstrating interpretable alignment between predictions and domain rules, enhancing the trustworthiness of the model.

In the railway case study, success is specifically tied to the model's ability to respect safety-critical constraints and compliance rules while preserving predictive utility. A model that avoids semantically inconsistent predictions (such as proposing impossible states or invalid relations) will be considered more successful than one achieving marginally higher accuracy at the cost of invalid outputs. The overarching expectation is to deliver models that combine strong predictive capability with semantic integrity, prioritizing reliability and consistency as the key enablers of adoption in real-world, safety-critical systems.

4. Preliminary Results

4.1. Negative Sampling Extended Framework

Building on the challenges identified in integrating ontological knowledge into KGE workflows (particularly the lack of standardized, reproducible tools) the initial research direction focused on addressing this gap through the injection of ontological properties into the negative sampling phase of KGE training.

The first concrete contribution was to extend a widely adopted knowledge graph embedding framework with a negative sampling module that can incorporate additional semantics from ontologies. This module also supports a generalized interface for adversarial sampling strategies. This work marks a foundational step towards enriching existing KGE methods with ontology-aware capabilities. Integrating support for ontological axioms directly into the negative sampling process enables the framework to be used seamlessly across a broad range of datasets and knowledge graphs, regardless of domain or complexity. Crucially, this extension also provides a modular and standardized interface for building and experimenting with new negative samplers. This significantly accelerates the development, evaluation, and ablation of novel techniques by eliminating the need for fragmented or ad hoc implementations. It paves the way for the subsequent research directions described earlier (such as ontology injection via loss constraints or geometric DL embeddings) by establishing a reusable and extensible foundation. Furthermore, by aligning with a widely used benchmarking framework, it ensures compatibility with existing literature and promotes reproducibility, thereby facilitating meaningful comparisons and broader community adoption.

Specifically, I extended the PyKEEN framework[1] by implementing six additional negative samplers [40]. These samplers are based on a generalized standardization process that simplifies the integration of new sampling strategies. Four of these samplers are directly derived from the research questions and directions outlined in the previous sections. The `TypedNegativeSampler` leverages

[1]https://pykeen.readthedocs.io/en/stable/

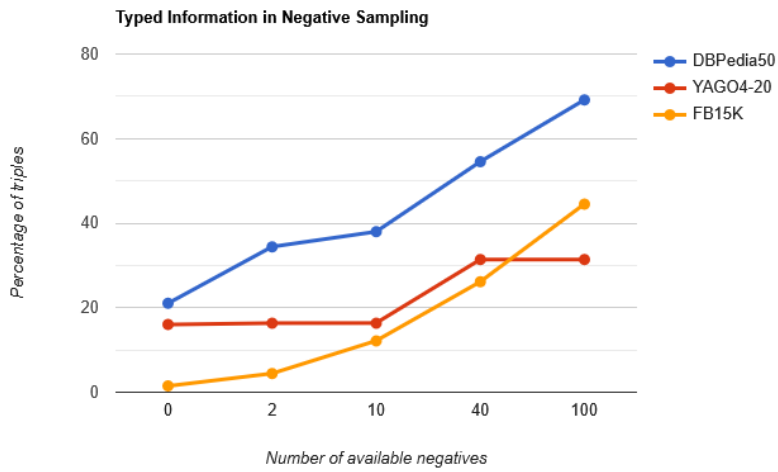

Figure 3: Negative pool availability analysis. Each point represents the percentage of triples with at least X distinct negative samples available.

domain and range constraints obtained through reasoning over the schema (TBox) to restrict the negative sampling pool at the triple level. A variant of this method uses instance-level (ABox) information to approximate domain and range classes. The `AdversarialNegativeSampler` provides an abstract implementation of adversarial sampling, where candidate negatives are selected based on a distance measure between the true entity and the entity predicted by an auxiliary adversarial model. To the best of my knowledge, this is the only implementation currently available that supports both adversarial and ontology-based sampling within the same feature-rich framework, using any user-defined model. Additionally, we provided a cleaned version of the datasets with pre-processed ontological information ready to use for experimentation. All the materials and documentation are available in the repository. [2] [3]

4.2. Type Information Availability Analysis

Following the implementation of the new negative sampling strategy, it became possible to further analyze sampling behavior when integrating additional semantic information from ontologies. A key challenge in working with knowledge graphs is their inherent incompleteness (partly due to the open-world assumption) which means not all entities and relations are connected to schema-level ontological information. This introduces a significant problem: if ontology-based sampling strategies cannot be applied consistently across the entire dataset, it becomes difficult and potentially misleading to evaluate and compare different methods.

Motivated by these considerations, an analysis was conducted on the negative pool statistics produced by the newly implemented samplers [40]. Specifically, the evaluation focused on the number of distinct negative entities that can be generated per triple under each sampling strategy. Figure 3 presents these statistics for three datasets, all enriched with ontology data, such as domain and range constraints and class hierarchies, except for FB15K, where the SoftType variant was applied for compatibility with the `TypedNegativeSampler`.

Results show that missing ontological annotations (despite reasoning to complete deducible missing information) lead to significantly small negative pools (reminding that the most used negative samplers, namely the random one, can use the whole set of entities as a candidate negative pool). For example, in the DBpedia50 dataset, over 30% of triples have a negative pool of fewer than two entities. This implies that during training, the negative sampler repeatedly selects from a minimal set, increasing the risk of overfitting and limiting learning effectiveness. Additionally, around 20% of triples cannot be used at all

[2]Repository: https://github.com/ivandiliso/refactor-negative-sampler
[3]Documentation: https://ivandiliso.github.io/refactor-negative-sampler/

for ontology-based sampling due to missing schema-level information (value 0 in the Figure 3). These issues persist across datasets, with nearly 70% of DBpedia50 triples having a negative pool of 100 or fewer entities, and over 30% in the other two datasets showing similar limitations.

This analysis provides an important insight into the practical applicability of ontology injection and, more generally, the use of background knowledge in KGE workflows. It highlights the need for thorough statistical examination of available data and underscores the importance of more complete and semantically annotated knowledge bases. These findings also reinforce the need for adaptable methods capable of handling incomplete ontology coverage while still leveraging available semantics.

5. Research Plan

The research activities outlined in this project are organized into distinct phases, each of which targets a specific aspect of the overall framework. While these phases provide a logical breakdown of the work, they are not intended to be carried out strictly sequentially. Many of them can be carried out in parallel and developed iteratively or in overlap, depending on the needs of the project and the results that emerge.

5.1. Phases of the Research

Phase 1: Study of Literature Comprehensive study and review of literature related to the PhD research themes, including knowledge graph embedding methods, machine and deep learning on multi-relational data, Description Logics, numerical and geometric interpretations of DL, embedding spaces for both KGs and ontologies, and techniques for injecting background knowledge into embedding models. The analysis also covers evaluation methodologies for link prediction and semantic consistency, as well as the practical application of these topics within the railway engineering domain, including implant rule specification and industrial-grade verification and validation.
Expected Result: A literature review summarizing the analyzed works and identifying the key approaches, strengths, and weaknesses, as well as potential areas for innovation.

Phase 2: Data Gathering and Analysis Collection and analysis of SOTA datasets in the fields of knowledge graphs and ontologies, with a focus on those that provide both instance-level (ABox) and schema-level (TBox) views, formalized in OWL. This includes gathering, cleaning, and possibly generating new datasets; extracting large-scale ontologies for experimentation (e.g., subsets of DBpedia or YAGO); and conducting statistical analysis relevant to ontology injection, such as completeness and structural properties. The phase also includes analysis of industrial formalizations used in railway implants, including preprocessing and domain-specific knowledge extraction.
Expected Result: A curated collection of datasets and associated resources for subsequent experimentation, along with detailed statistical reports to guide experimental design in later phases.

Phase 3: Ontology Injection Framework Definition and Development Formalization and development of a structured framework for injecting ontological knowledge into KGE models. The framework is divided into three main components: (1) Injection during negative sampling, (2) Injection via loss constraints, and (3) Injection through the embedding space using DL-based representations. Each component will be studied and implemented individually before being integrated into a coherent system.
Expected Result: A robust and extensible ontology-injection framework ready for large-scale evaluation in the next phase.

Phase 4: Evaluation Framework Definition and Development Design and development of a comprehensive evaluation suite capable of assessing both the predictive performance of the model and the quality of the semantic knowledge learned. This includes defining or reusing evaluation metrics

that align with common practices in the literature and domain-specific requirements.
Expected Results: A standardized and reusable evaluation framework for assessing the effectiveness of ontology-aware KGE models in both general and domain-specific scenarios

Phase 5: Experiment Execution and Discussion Execution of extensive experiments to evaluate the proposed models using the defined evaluation suite. The analysis will assess performance, robustness, reliability, and applicability to mission-critical domains. Results will be compared with state-of-the-art models, and ablation studies will be conducted to measure the specific impact of ontology injection. This phase also includes evaluation of the model's adaptability to different domains and datasets.
Expected Results: A comprehensive report detailing the experimental results and critical analysis, offering insights for further improvements, refinement of the models, and potential directions for future research.

Phase 6: Domain Specific Application in Railway Domain Adaptation and fine-tuning of the general framework to meet industry-specific needs within the railway domain. This includes translating general methods into specialized models that handle railway-specific data, semantics, and operational constraints. The phase also involves on-site collaborations (e.g., internships) with local companies to analyze and apply the proposed methods in real-world scenarios.
Expected Results: Deployment-ready, domain-tuned models and processes suitable for industrial application, along with validation through field-specific use cases.

5.2. Research Timeline and Initial Contributions

PhD Year 1 (Current) The first year of this PhD project has primarily focused on Phases 1, 2, and 3. Following a thorough literature review, the data gathering and analysis phase commenced, leading to the collection and study of several state-of-the-art datasets. Preliminary results have been presented in earlier sections [40]. In collaboration with AI2, the company co-funding this PhD, an analysis was also carried out on industry-relevant technologies for formalizing domain knowledge, including RailML [4] (both the markup language and its associated schema ontology [5]), the ERA Ontology [6], and the ERA Knowledge Graph [7]. Current efforts are focused on anonymizing and transforming these resources into Semantic Web standards to facilitate broader reuse. Regarding Phase 3, initial work concentrated on ontology injection during the negative sampling process, with implementation and evaluation outlined previously. In parallel, Phase 4 was initiated through the design of an evaluation framework, including the development of new evaluation metrics and early implementation using PyKEEN, PyTorch, and Torch Geometric. The incremental development of both the ontology injection framework and the evaluation suite will continue through the remainder of the first year.

PhD Year 2 Building on the foundational datasets and preliminary framework developed during the first year, the second year of the project will focus on further developing ontology injection mechanisms directly within KGE models. In Q1–Q2, Phase 3 will continue with an emphasis on ontology injection through geometric embedding spaces grounded in Description Logics. This phase will involve the formalization of new techniques, the design of model architectures, and the definition of appropriate training strategies. The finalization of the evaluation suite (Phase 4) is also planned during this period. In Q3–Q4, Phase 5 will be initiated with large-scale experimental evaluations of the proposed models. These evaluations will be carried out using the developed framework and guided by the finalized semantic and performance metrics. Based on the evaluation results, an iterative refinement process will be undertaken to improve model robustness, enhance semantic alignment, and address identified limitations—ensuring the scalability and effectiveness of the proposed approach.

[4] https://www.railml.org/en/
[5] https://ontology.railml.org
[6] https://data-interop.era.europa.eu/era-vocabulary/
[7] https://www.era.europa.eu/domains/registers/era-knowlege-graph_en

PhD Year 3 Assuming successful completion of the previous phases, the third year of the project will focus on Phase 6: adapting and applying the proposed framework to domain-specific requirements within the railway sector. This will require iterative refinement of both Phase 3 (model definition) and Phase 5 (experimental evaluation) to align the ontology injection methodology with the specific constraints, formalisms, and standards of the target industrial domain. This phase will also include an internship with an industry partner to acquire in-depth domain knowledge and practical experience. Close collaboration with domain experts will be essential to ensure the correctness, formal compliance, and applicability of the proposed models in a mission-critical environment. In parallel, previously conducted experiments will be re-evaluated and extended, with the goal of delivering a mature, validated, and industry-ready solution.

6. Conclusion

This PhD project takes a structured and comprehensive approach to improving knowledge graph embedding models by integrating ontological knowledge. In response to the current limitations of semantic representation, the research will focus on developing a general, domain-agnostic framework that supports ontology injection at various stages of the KGE pipeline, including negative sampling, loss constraints, and embedding space alignments. The proposed work aims to enrich existing models with ontological semantics and emphasizes practical usability through standardized implementations, reusable components, and robust evaluation tools, thereby aligning with software engineering best practices.

During the first year, the project laid the foundation through an in-depth literature review and dataset analysis, as well as the initial development of the ontology injection framework through negative sampling. These contributions have already enabled a deeper analysis, exposing critical challenges such as data incompleteness and the absence of standardized tools in the literature.

Looking ahead, the next phases will extend the framework to support ontology injection through description logic embedding spaces and semantic loss constraints, followed by large-scale evaluation and domain-specific adaptations. The ultimate goal is to deliver a unified, flexible methodology that can support both academic research and industrial applications. By bridging the gap between symbolic reasoning and statistical learning, this research will advance the field of knowledge representation and enable more semantically grounded machine learning systems.

Acknowledgments

This work has been carried out within the program "*Patti territoriali dell'alta formazione per le imprese*" (CUP F61B23000370006), with the support of the company *AI2 – Applicazioni di Ingegneria ed Informatica Srl*. The research was conducted as part of the project "*SymbioticRAIL2 – Symbiotic AI techniques supporting design and maintenance of railway systems*". This work was done in collaboration with my PhD supervisor and co-supervisor, Nicola Fanizzi and Claudia d'Amato, I thank them for the continued support throughout this research.

Declaration on Generative AI

During the preparation of this work, the author used DeepL Write, Grammarly in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the author reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, et al., Knowledge graphs, ACM Computing Surveys (Csur) 54 (2021) 1–37. doi:10.1145/3447772.

[2] Z. Li, X. Liu, X. Wang, P. Liu, Y. Shen, Transo: a knowledge-driven representation learning method with ontology information constraints, World Wide Web 26 (2023) 297–319. doi:10.1007/s11280-022-01016-3.

[3] C. d'Amato, N. F. Quatraro, N. Fanizzi, Injecting background knowledge into embedding models for predictive tasks on knowledge graphs, in: The Semantic Web: 18th International Conference, ESWC 2021, Virtual Event, June 6–10, 2021, Proceedings 18, Springer, 2021, pp. 441–457. doi:10.1007/978-3-030-77385-4_26.

[4] T. R. Gruber, Toward principles for the design of ontologies used for knowledge sharing?, International journal of human-computer studies 43 (1995) 907–928. doi:10.1006/ijhc.1995.1081.

[5] J. Hao, M. Chen, W. Yu, Y. Sun, W. Wang, Universal representation learning of knowledge bases by jointly embedding instances and ontological concepts, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019, pp. 1709–1719. doi:10.1145/3292500.3330838.

[6] C. Wang, T. Gan, X. Li, L. Zhang, X. Wang, An ontology-enhanced knowledge graph embedding method, in: Proceedings of the 2023 12th International Conference on Computing and Pattern Recognition, 2023, pp. 51–57. doi:10.1145/3633637.3633645.

[7] F. Zhapa-Camacho, R. Hoehndorf, From axioms over graphs to vectors, and back again: Evaluating the properties of graph-based ontology embeddings, in: 17th International Workshop on Neural-Symbolic Learning and Reasoning, NeSy 2023, CEUR-WS, 2023, pp. 85–102. doi:10754/690862.

[8] S. M. Kazemi, D. Poole, Simple embedding for link prediction in knowledge graphs, in: Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18, Curran Associates Inc., Red Hook, NY, USA, 2018, p. 4289–4300.

[9] Q. Li, S. Cao, L. Zhang, A domain knowledge-based railway equipment fault diagnosis framework, in: International Symposium on Knowledge and Systems Sciences, Springer, 2024, pp. 330–344. doi:10.1007/978-981-96-0178-3_23.

[10] K. Fowler, Mission-critical and safety-critical development, IEEE Instrumentation & Measurement Magazine 7 (2004) 52–59. doi:10.1016/b978-0-7506-8567-2.00001-9.

[11] Z. Ye, Y. J. Kumar, G. O. Sing, F. Song, J. Wang, A comprehensive survey of graph neural networks for knowledge graphs, IEEE Access 10 (2022) 75729–75741. doi:10.1109/ACCESS.2022.3191784.

[12] Z. Zhang, J. Cai, Y. Zhang, J. Wang, Learning hierarchy-aware knowledge graph embeddings for link prediction, in: Proceedings of the AAAI conference on artificial intelligence, volume 34, 2020, pp. 3065–3072.

[13] H. Ren, W. Hu, J. Leskovec, Query2box: Reasoning over knowledge graphs in vector space using box embeddings, in: International Conference on Learning Representations (ICLR), 2020. doi:10.1117/12.2683290.

[14] Z. Huang, M.-F. Chiang, W.-C. Lee, Line: Logical query reasoning over hierarchical knowledge graphs, in: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2022, pp. 615–625. doi:10.1145/3534678.3539338.

[15] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. Van Den Berg, I. Titov, M. Welling, Modeling relational data with graph convolutional networks, in: The semantic web: 15th international conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, proceedings 15, Springer, 2018, pp. 593–607. doi:10.1007/978-3-319-93417-4_38.

[16] S. Vashishth, S. Sanyal, V. Nitin, P. Talukdar, Composition-based multi-relational graph convolutional networks, in: International Conference on Learning Representations, 2019.

[17] D. Bahdanau, K. H. Cho, Y. Bengio, Neural machine translation by jointly learning to align and translate, in: 3rd International Conference on Learning Representations, ICLR 2015, 2015.

[18] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, Y. Bengio, Graph attention networks,

in: International Conference on Learning Representations, 2018.
[19] D. Nathani, J. Chauhan, C. Sharma, M. Kaul, Learning attention-based embeddings for relation prediction in knowledge graphs, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019, pp. 4710–4723. doi:10.18653/v1/p19-1466.
[20] T. Bansal, D.-C. Juan, S. Ravi, A. McCallum, A2n: Attending to neighbors for knowledge graph inference, in: Proceedings of the 57th annual meeting of the association for computational linguistics, 2019, pp. 4387–4392. doi:10.18653/v1/p19-1431.
[21] Z. Zhang, F. Zhuang, H. Zhu, Z. Shi, H. Xiong, Q. He, Relational graph neural network with hierarchical attention for knowledge graph completion, in: Proceedings of the AAAI conference on artificial intelligence, volume 34, 2020, pp. 9612–9619. doi:10.1609/aaai.v34i05.6508.
[22] X. Zhao, Y. Jia, A. Li, R. Jiang, K. Chen, Y. Wang, Target relational attention-oriented knowledge graph reasoning, Neurocomputing 461 (2021) 577–586. doi:10.1016/j.neucom.2021.03.135.
[23] Y. Wang, H. Wang, J. He, W. Lu, S. Gao, Tagat: Type-aware graph attention networks for reasoning over knowledge graphs, Knowledge-Based Systems 233 (2021) 107500. doi:10.1016/j.knosys.2021.107500.
[24] D. Allemang, J. Hendler, Semantic web for the working ontologist: effective modeling in RDFS and OWL, Elsevier, 2011. doi:10.1145/3382097.
[25] The Description Logic Handbook: Theory, Implementation and Applications, 2 ed., Cambridge University Press, 2007. doi:10.1017/CBO9780511711787.
[26] F. Zhapa-Camacho, R. Hoehndorf, Lattice-preserving alc ontology embeddings, in: International Conference on Neural-Symbolic Learning and Reasoning, Springer, 2024, pp. 355–369. doi:10.1007/978-3-031-71167-1_19.
[27] J. Chen, P. Hu, E. Jimenez-Ruiz, O. M. Holter, D. Antonyrajah, I. Horrocks, Owl2vec*: Embedding of owl ontologies, Machine Learning 110 (2021) 1813–1845. doi:10.1007/s10994-021-05997-6.
[28] M. Jackermeier, J. Chen, I. Horrocks, Dual box embeddings for the description logic el++, in: Proceedings of the ACM on Web Conference 2024, 2024, pp. 2250–2258. doi:10.1145/3589334.3645648.
[29] Ö. L. Özcep, M. Leemhuis, D. Wolter, Embedding ontologies in the description logic alc by axis-aligned cones, Journal of Artificial Intelligence Research 78 (2023) 217–267. doi:10.1613/jair.1.13939.
[30] V. Gutiérrez-Basulto, S. Schockaert, From knowledge graph embedding to ontology embedding? an analysis of the compatibility between vector space representations and rules, in: International Conference on Principles of Knowledge Representation and Reasoning, 2018. URL: https://api.semanticscholar.org/CorpusID:51935922.
[31] C. d'Amato, L. Mahon, P. Monnin, G. Stamou, et al., Machine learning and knowledge graphs: Existing gaps and future research challenges, Transactions on Graph Data and Knowledge 1 (2023) 1–35. doi:10.4230/TGDK.1.1.8.
[32] M. Alam, F. van Harmelen, M. Acosta, Towards semantically enriched embeddings for knowledge graph completion, Neurosymbolic Artificial Intelligence 1 (2025) NAI–240731. doi:10.3233/NAI-240731.
[33] R. Xie, Z. Liu, M. Sun, et al., Representation learning of knowledge graphs with hierarchical types., in: IJCAI, volume 2016, 2016, pp. 2965–2971.
[34] N. Jain, T.-K. Tran, M. H. Gad-Elrab, D. Stepanova, Improving knowledge graph embeddings with ontological reasoning, in: International Semantic Web Conference, Springer, 2021, pp. 410–426. doi:10.1007/978-3-030-88361-4_24.
[35] N. Hubert, P. Monnin, A. Brun, D. Monticolo, Sem@ k: Is my knowledge graph embedding model semantic-aware?, Semantic Web 14 (2023) 1273–1309. doi:10.3233/sw-233508.
[36] H. Hadj-Mabrouk, Literature review on applications of ontologies and knowledge graphs in railway transport safety, Railway Transport and Engineering-A Comprehensive Guide [Working Title] (2024) 1–29. doi:10.5772/intechopen.1006278.
[37] C. Yang, Y. Zheng, X. Tu, R. Ala-Laurinaho, J. Autiosalo, O. Seppänen, K. Tammi, Ontology-based knowledge representation of industrial production workflow, Advanced Engineering Informatics

58 (2023) 102185. doi:10.1016/j.aei.2023.102185.

[38] D. Zhou, B. Zhou, Z. Zheng, E. V. Kostylev, G. Cheng, E. Jimenez-Ruiz, A. Soylu, E. Kharlamov, Enhancing knowledge graph generation with ontology reshaping–bosch case, in: European Semantic Web Conference, Springer, 2022, pp. 299–302. doi:10.1007/978-3-031-11609-4_45.

[39] H. Pu, T. Hu, T. Song, P. Schonfeld, X. Wan, W. Li, L. Peng, Modeling and application of a customized knowledge graph for railway alignment optimization, Expert Systems with Applications 244 (2024) 122999. doi:10.1016/j.eswa.2023.122999.

[40] C. d'Amato, I. Diliso, N. Fanizzi, Z. Saeed, Enhancing pykeen with multiple negative sampling solutions for knowledge graph embedding models, 2025. URL: https://arxiv.org/abs/2508.05587. arXiv:2508.05587.