

Neuro-Symbolic AI for conflict-aware learning over Knowledge Graphs

Laura Balbi

Faculdade de Ciências da Universidade de Lisboa, Portugal
lbalbi@lasige.di.fc.ul.pt

Abstract. Knowledge Graphs (KGs) are fundamental to digital ecosystems, facilitating data integration, annotation, and interoperability across diverse domains. However, their construction and update depend on the merging of heterogeneous data sources that can introduce ambiguities and contradictions, particularly when introduced conflicting perspectives. Existing contradiction detection methods primarily resolve conflicts by discarding one of the contradictory statements, often overlooking the possibility of a nuanced, composite truth.

When contradictions explicitly violate KG semantics, they can be identified through logical reasoning. However, in the absence of domain-specific semantics, external contextual information becomes crucial for accurate detection. While advancements in KG representation learning (KGRL) methods, such as negation-aware KG embeddings, have improved semantic expressiveness, they are still limited in their ability to capture and reason about contradictions. Among neural-based KGRL are Language Models (LM), that still struggle with reasoning over conflicting information. This work explores how a neuro-symbolic approach — integrating symbolic and LM neural representations — can enhance contradiction detection and improve KGRL robustness, reliability, and predictive performance, paving the way for more trustworthy machine learning applications over KGs.

Keywords: Knowledge Graphs · Neuro-Symbolic AI · Language Models · Contradictions

1 Introduction

Knowledge Graphs (KG) have become an integral yet often unnoticed component of digital ecosystems that significantly enhances how individuals, researchers, and companies interact with data. A KG is a structured approach to storing and representing knowledge on real-world entities and concepts, through a graph-based data model where entities are represented as nodes and their relationships as edges, in a queriable format that is useful for applications such as question-answering and generating explanations to support machine learning decision-making [1]. KG representation learning (KGRL) is the process of mapping KG entities and their semantic relations into a low-dimensional sub-symbolic vector space [2] for ML usage. KGRL methods function under the closed-world assumption (CWA) — that a KG captures all factuality in its domain and that which is not represented as true in the KG is assumed to be false —, whilst in reality most KGs are incomplete and contain misinformation or conflicts that could compromise the accuracy and trustworthiness of their applications.

Conflicts, or contradictions, are commonly introduced during KG construction or update and can either be logical, if they pose a direct logical violation, or implicit if not detectable without introduction of external domain context and semantics. Platforms such as Wikidata [3] have actively addressed such issues through community validation efforts and by introducing provenance and reliability measures. This is a relevant avenue for the platform as about 23.71% of Wikidata’s [4] statements have multiple values for the same pair of item and property – the subject and predicate of a statement. A fraction of these are contradictions and do not necessarily correspond to data annotation errors but instead reflect polyvocal narratives regarding a subject, stemming from different human perspectives worthy of being represented in data; hence the need for context and provenance information, which is only available for 16.3% of Wikidata. The process of identifying facts that cannot be true simultaneously is referred to as contradiction detection. Although recent detection techniques over KGs include error-aware KGRL (e.g. [5]), current KGRL is still ill-equipped to handle and model contradiction. Subsequent efforts aimed to increase the semantic expressiveness of both shallow KG embedding (KGE) approaches [6], and deep KGRL-based methods such as Graph Neural Networks (GNNs) [7]. [8] designed a KGE approach that considers negation and disjunction – relevant axioms for contradiction modelling, but how they translate to learning under contradiction is unclear - and GNNs mostly ignore the semantic implications of KG edges [9]. Recent works on KGRL have proposed the usage of Language Models (LMs) to generate semantic representations for downstream KG Completion tasks [10]. However, current LMs exhibit reasoning deficiencies in capturing conflicting information [11]. Nevertheless, a refined consideration of the semantic value and ontological implications of negation is key in improving the representation of entities by KGE approaches, with a clear impact in reducing Type I errors (misclassification of False Positives) [12] in biomedical ML applications, paving the way for further research on KGRL with contradictory facts.

1.1 Problem Statement and Hypothesis

In the context of a KG, a contradiction can be formally defined as a logical violation resulting from two or more statements asserting mutually exclusive or logically incompatible information about a shared entity, relation or event, under the same semantic framework [13]. Contradictions can either be explicit, direct conflicts between statements that are easily identified if there is a logical violation: e.g. $\langle \textit{Paris}, \textit{capital of}, \textit{France} \rangle$ and $\langle \textit{Versailles}, \textit{capital of}, \textit{France} \rangle$, through the violation of an existing cardinality constraint of 1 for the relation “*capital of*”; or implicit, if not detectable by logical reasoning due to lack of assertions but identifiable when introducing logical context, e.g. $\langle \textit{Republic of Crimea}, \textit{instance of}, \textit{Republic of Russia} \rangle$ and $\langle \textit{Republic of Crimea}, \textit{instance of}, \textit{disputed territory} \rangle$, retrieved from Wikidata, are not considered a direct violation as there is no cardinality constraint for “instance of”, but are ambiguous enough that if given the proper contextual semantics may become obvious contradictions. On the other hand, contradictions may represent different facets of a complex, multifaceted truth that, if combined, contribute to a more comprehensive and accurate representation of entities. For example, contradictions that can introduce knowledge if enriched with temporal, spatial or modal agents that contextualize their ambiguity. However, current KGRL largely overlooks this perspective and how modeling contradictions as complementary components of

a broader contextual reality could enhance the accuracy and applicability of KG-based ML approaches.

The underlying hypothesis of this work is that a neuro-symbolic approach, by virtue of integrating symbolic and sub-symbolic representations, can bridge KGs and LLMs to answer the following research questions:

RQ1: To what extent does the existence of contradictions in a KG, both implicit and explicit, impact the performance, robustness, reliability and transparency of state-of-the-art KGRL?

RQ2: Does explicitly modeling contradictions and exploring this information in KG entity representations improve ML performance with contradictory facts?

RQ3: Can external sources of knowledge, such as LLMs, be explored to detect implicit contradictions?

RQ4: Do confidence scores and provenance information improve the handling of contradictory facts by KGRL as reflected by predictive performance and usefulness?

The main objective of this work is to investigate how contradictory facts can be captured by KGRL algorithms to improve the robustness, reliability, transparency and predictive performance of ML methods. This study aims to contribute to the field of KGRL by: i) Providing a formal classification of contradictions in KGs and definition of implicit contradictions; ii) The first effort to model contradictions into entity representations; iii) Researching how different contradictions can be captured by KGRL algorithms to improve the robustness, reliability, transparency and predictive performance of ML over KGs; iv) Systematic evaluation of the developed approaches over several benchmark datasets and on different domain tasks; v) Addressing two major themes for future Artificial Intelligence – supporting polyvocal perspectives and mitigating misconceptions.

2 State-of-the-art

2.1 On Knowledge Graph Representation Learning

As KGs grow in size and complexity, efficient computation increases in importance. KGRL can embed entities and relations into a continuous low-dimensional vector space that preserves graph positional and connectivity information while bridging graph-structured data and traditional ML methods.

The state-of-the-art in KGRL can be broadly categorized into three groups: (i) translation-based, (ii) semantic matching-based, (iii) random-walk based and (iv) neural-based methods.

Translation-based methods (e.g. TransE [14]) interpret relations as translational operations between the head and tail entities that make up a triple in the KG. Semantic matching-based KG embedding methods (e.g. RESCAL [15], distMult [16]) work under the assumption that entities with similar semantics have similar embeddings and therefore use similarity-based scoring functions to capture the latent semantics of KG components in their vector space representations. Random walk-based methods (e.g. RDF2Vec [17]) transform the KG structure into node sequences that are randomly sampled with “walks” and given as input to a LM that learns latent entity representations. While the latter methods are shallow – use simple linear transformations -, neural network-based methods generate deeper representations based on the learning of graph structure to solve a specific task. Graph-based networks have been popular in KGRL due to accounting for both structure and node features in the representation

learning process by pooling representations of node local neighborhoods in the KG [7], under the assumption that connectivity and proximity in the KG reflects similarity.

However, most GNN approaches fail to capture the various semantic aspects of entities and relations, including those with reverse or opposing meanings such as negation and contradiction, which may reduce the quality of its representations.

While traditional KGRL struggles to represent these statements in a consistent vector space, methods capable of modeling edges of opposing values (i.e., positive and negative edges) can naturally model contradictions in KGs. [12] proposed TrueWalks, a random walk -based embedding method that integrates negative annotations into the KGRL process, differentiates between positive and negative annotations and takes into account the semantic implications of negation in ontology-rich KGs.

Current neural -based KGRL takes two approaches to exploring relations of opposing values, with a focus on handling negation: the signed GNNs for modeling relations of different signals onto node embeddings, and GNNs enhanced with losses for training node embeddings that model positive edges as proximal and negative edges as more distant. From this approach the Signed Graph Convolutional Network was developed, commonly applied to social networks under the “Balance Theory” assumption on triangle relationships between nodes, modeling positive and negative edges into node representations with two separate aggregation mechanisms. Similarly, the Signed Graph Attention Network [18] uses dual attention mechanisms for positive and negative edges. However, while the assumption of similarity of traditional message passing in GNNs works well with positively connected neighbors, the same does not hold for negatively connected neighbors, that should be considered more dissimilar and therefore more distant in a graph instead of belonging to the same local neighborhood.

On the other hand, the second approach to modeling opposing relations for node representation was introduced in [19] with the triplet loss, which trains the model to learn and embed the similarity (or dissimilarity) between node representations, so that related nodes are projected closer and disparate nodes are projected further apart. The loss ensures that the positive triple is at least α distance closer than the negative triple. If this condition is met, the loss is zero, meaning the network is correctly separating data samples. Since then, a triplet network [20] architecture has emerged, composed of three identical neural branches that share weights and produce three distinct node embeddings that are then compared using the triplet Loss. The N-pair Loss [21] extends the triplet Loss to multiple negative "neighbors" per anchor-positive pair. Its goal is to maximize the similarity between the anchor and positive node while minimizing similarity between the anchor and multiple negatives. A key limitation of embedding learning with such loss functions is the process of sampling negatives, that while under a more flexible and relaxed CWA, may still introduce assumed negative samples whose actual truth value is unknown. This could lead to misclassifications that compromise the quality of the trained node representations.

Furthermore, the methods above are tailored to handling binary signals and cannot adapt to the richer spectrum of conflicting relational semantics found in complex KGs. The ability to handle contradicting relations beyond simple negation remains limited.

2.2 Knowledge Graphs with contradictions

It is very common for large KGs, and in particular those crowd-sourced or constructed through semi-automated web-scraping, to contain contradictions due to the merging of information from different sources or reflecting different world-views [4]. Despite their prevalence, contradictions are often non-acknowledged or are subjected to fact-checking methods for selecting one of the conflicting sides. This is due to the "majority-view" nature of most data collection and ML strategies, that solve contradictions based on popularity-based metrics. However, a contradicting fact's popularity may not reflect its contribution nor its sufficiency towards a realistic and comprehensive description of the entity involved. Real-world entities may have multiple "real" facets that combine into a complex description.

3 Methodology

As KG construction becomes more automated, contradicting statements, both explicit and implicit, can be expected to become even more prevalent. In this context, Neuro-Symbolic (NeSy) AI emerges as a promising approach for implicit contradiction detection, as it combines neural-based ML methods with symbolic techniques. Current NeSy methods mostly work with smaller rulesets but can be extended to the much larger size and complex semantics of a KG. This allows for their integration with other approaches, such as LLMs, to detect and model contradictions to learn entity representations off of.

3.1 Data

The developed approaches will be tested over three distinct domain KGs: on factual knowledge - Wiki KG -, commonsense knowledge - Concept KG -, and a biomedical KG - PPI KG. The KGs result of the combining of positive and negative statements. This combination will potentially introduce mutually negating and conflicting relations, amplifying the contradictions already present within the original knowledge sources. The following sections detail their statistics and data sources.

Wiki KG Wikidata is a free and open KG that supports Wikimedia projects such as Wikipedia, and often serves as a reliable source of factual information. It's current version describes a total of 114,682,266 entities and contains 15,019,727,427 facts. Several works that have performed quality analyses over Wikidata [4] have reported logical inconsistencies, including contradicting and/or ambiguous claims. Furthermore, the Wikinegata project [26] released a set of 681,000,000 relevant negative facts on 600,000 of Wikidata's entities, inferred through a peer-based ranking method. The inclusion of the Wikinegata's set into the Wikidata KG will generate a KG with an increased amount of logical contradictions, the Wiki KG.

Concept KG The ConceptNet [23] is a commonsense KG that connects words and phrases of natural language with labeled edges. It contains 34,000,000 facts between over 8,000,000 ConceptNet entities. Uncommonsense [24] is a set of 13,600,000 real negative facts for 8,000 of the entities covered by ConceptNet. By combining both, there is a potential to either detect existing contradictions or generate new ones to construct a conflict-enriched commonsense Concept KG.

PPI KG The PPI KG is composed of a Protein-Protein Interaction (PPI) network with 11,953 proteins of the human interactome and their annotations to the Gene Ontology (GO) [22]. The PPI data contains 479,665 experimentally verified positive interactions. The combining of PPI network with the GO graph through positive and negative annotations (292,377 and 67,300, respectively) totals 65,321 entities and 906,768 triples. In summary, in the PPI KG an edge between a protein (instance) and a GO class indicates that a particular protein performs a specific function described in the GO. Similarly, an edge between two classes represents a "subclass of" relation. An edge between two proteins, typed "positive PPI", will correspond to an experimentally verified physical association between the two.

3.2 Contradiction detection approach and KG enrichment

This stage of the project will focus on the development of different methods of contradiction detection. We propose the exploration of pre-trained LMs and LM prompt-based and in-context learning with external knowledge sources (i.e. domain KGs, ontologies) for implicit contradiction detection and investigation of methods that identify and mitigate self-contradiction in LLMs and can be translated into KG implicit contradiction detection by leveraging mapping functions between the latent KG space and the LLM space. These approaches will be compared with current baseline detection methods based on traditional techniques such as logic-based reasoners and rule-based inference engines. Afterwards, the KGs will be enriched with the semantic information that renders implicit contradictions logically detectable and introduces context into the ambiguity of contradicting statements to enhance downstream tasks. The enrichment process should produce three variants of each KG, reflecting the complexity of the detection process: a) a KG solely with direct logical contradictions; b) a KG with both direct and inferred logical contradictions; iii) a KG containing both logical and implicit contradictions.

3.3 KG Representation Learning with contradictions

This step of the approach will focus on the development of a theoretical model for the sub-symbolic representation of entities considering disjointness, property inheritance and negation – criteria that are essential for modeling contradiction – and capturing and taking into account available information on confidence and provenance. This theoretical model will be the foundation for designing novel neural-based KGRL approaches, whose modeling mechanism is rooted in the notion that closeness in the graph reflects similarity, which is not the case when contradiction comes into play. An important challenge to tackle is how, dissimilarly to nodes connected by positive edges, negative neighbor nodes connected by direct edges should not be modeled closer in the vector space nor be included in the same local neighborhood. This involves investigating how contradictions should shape the aggregate and update functions of GNNs.

3.4 Domain tasks

The developed approaches will be extensively evaluated in three use cases to assess the impact of contradictory information on the robustness, reliability, transparency and performance of KGRL applications.

The Wiki KG and the Concept KG will serve as testbeds for KGC tasks. KGC pertains to the task of predicting missing triples in a KG, either by identifying a missing relation or entity. The PPI KG provides a test bed in the biomedical domain where it is common for research findings to contradict each other [12]. This task focuses on testing the developed approaches on PPI prediction, with a binary classification of pairs of proteins on whether they are interacting or not. For the KGC tasks over the Wiki KG and Concept KG, since the KGs lack a hierarchical structure there are no concerns such as class inheritance or subsumption (a broader class including a more specific subclass). In contrast, for the PPI prediction task, since the PPI KG contains the taxonomy-based GO, which contains parent-child relationships between classes, there is a need to ensure that protein representations respect the hierarchical constraints, to align the protein embeddings with biological reality [25]. If a protein is known to have a specific function (i.e., has a positive annotation with a GO class), we can infer that it also has its broader functions (superclasses). Therefore positive annotations should be propagated upwards in the GO hierarchy. On the other hand, if a protein is known not to have a certain function (i.e., a negative annotation), it should also not have any of the more specific sub-functions (subclasses), and therefore negative annotations should be propagated downwards in the hierarchy.

4 Baseline Evaluation on PPI KGs

The project is currently at the stage of conducting baseline experiments over the original PPI KG without the enrichment step in 3.2, and by exploring the following conditions: i) **PPI w/o LC**: both positive and negative annotations for the same protein, while ensuring there are no explicit logical contradictions (LC); ii) **PPI Full**: both positive and negative annotations for the same protein, regardless of their causing contradictions in the KG; iii) **PPI w/ only Pos Annotations**: only positive annotations; iv) **PPI w/ only Neg Annotations**: only negative annotations. The PPI w/o LC KG was obtained by fully removing all conflicting sides of 6,575 logical contradictions detected by running the Hermit [27] reasoner over the PPI KG as a baseline detection method. It is noted that the contradictions detected were either of direct negations between positive and negative annotations or of indirect negations between negative annotations to superclasses of positive annotations for the same protein. Futurely, the detection process in 3.2 will allow for the detection of other types of contradictions, be it explicit or implicit.

Baselines i) and ii) should contribute to assess the impact of the inexistence or improper modeling of contradictions, while iii) and iv) contribute to understanding the impact of modeling opposite relationships, and the role that each of them has, into entity representations. Furthermore, the KGRL approaches employed will process the different KGs (baselines and enriched) over different degrees of graph heterogeneity and edge types to reflect the approaches' abilities to distinguish between the semantic values of the different edges. Table 1 describes the heterogeneity settings explored and the models trained over them, along with the types of edges considered for each experiment. Current baseline experiences were tested with GCN [28] and GAT [29] architectures extended to support a flexible number of edge types [30]. Traditional GNNs were trained over a single edge type, the "BiGNNs" incorporated two edge types to distinguish between positive and negative semantic relationships and heterogeneous GNNs (HGNNs) were trained considering a different edge type per relation in

the KG – i.e. PPI, positive or negative annotations, "subclass of" links –. The models were trained with random initialization of features for the PPI prediction task on a 10-fold cross-validation setting for hyperparameter tuning and subsequent testing of the best model over a test set of 15% of the dataset’s proteins and with 479,665 randomly sampled negative PPIs for test proteins, repeated for 5 runs.

Table 1. Description of experiments with different edge types for each KG and their composition. Edge Types: E = "Edge"; PE = "Positive Edge" and NE = "Negative Edge". GO = GO "subclass of" links; Pos = Positive Annotations; and Neg = Negative Annotations.

Models	KG	Edge Types	Type Composition
GCN & GAT	w/o LC	1 type – E	E: GO, Pos, Neg & PPI
	Full		E: GO, Pos, Neg & PPI
	only Neg Annotations		E: GO, Neg & PPI
	only Pos Annotations		E: GO, Pos & PPI
BiGCN & BiGAT	w/o LC	2 types – PE & NE	PE: GO, Pos & PPI; NE: Neg
	Full		PE: GO, Pos & PPI; NE: Neg
	only Neg Annotations	1 type – PE	PE: GO & PPI; NE: Neg
	only Pos Annotations		PE: GO, Pos & PPI
HGCN & HGAT	w/o LC	4 types	GO, Pos, Neg & PPI
	Full	4 types	GO, Pos, Neg & PPI
	only Neg Annotations	3 types	GO, Neg & PPI
	only Pos Annotations	3 types	GO, Pos & PPI

5 Preliminary Results

This section discusses the preliminary results obtained from the baseline experiments across four variants of the PPI KG— PPI w/o LC, PPI Full, PPI Pos, and PPI Neg in a first effort to understand the role of contradictions and opposite semantic relations in shaping protein representations and their downstream impact on PPI prediction. Table 2 shows the results for baseline experiments for a PPI prediction task over the PPI KGs when contradictions are not explicitly detected and modeled into KG representations.

With exception of the experiments for BiGCN and HGCN, the KG with only negative annotations consistently outperforms others, which underscores the potential value of including negative knowledge in KGRL and yet also suggests that in PPI prediction it is more relevant to include information on functions that proteins do not perform than on the functions they do perform— an insight that may seem counterintuitive. Such results could hint at a need for a more nuanced and smart approach to the simultaneous usage of positive and negative annotations to enhance model performance rather than altogether or by simply removing contradictions in a naive fashion as done in the baselines.

The consensual increase in performance from the GCN to the HGCN also highlights a potential advantage in actively modelling opposing relations and contradictions in separate. Interestingly, the traditional GAT often performed better than the BiGAT and HGAT, especially accuracy and recall -wise. This could be due to the GAT’s ability to weigh edges according to their importance being applied across edge types whilst the Bi- and HGAT weigh importance only within each edge type and not across types.

Table 2. Preliminary results for baseline experiments over the 4 KG variants PPI w/o LC, PPI Full, PPI w/ only Pos Annotations and PPI w/ only Neg Annotations. Best results in bold for each ML model. Results pertain to mean of 5 runs with a maximum standard deviation of 0.076. Best results for each GNN in bold.

ML	KG	Accuracy	F-measure	Precision	Recall	ROC-AUC
GCN	Full	0.6502	0.6071	0.9440	0.3193	0.9023
	w/o LC	0.6484	0.6041	0.9444	0.3153	0.8998
	only Neg Annotations	0.6567	0.6175	0.9346	0.3369	0.8954
	only Pos Annotations	0.6438	0.5970	0.9477	0.3045	0.9007
GAT	Full	0.5721	0.4795	0.9375	0.1546	0.6913
	w/o LC	0.5690	0.4760	0.9310	0.1491	0.6778
	only Neg Annotations	0.6002	0.5287	0.9269	0.2175	0.6777
	only Pos Annotations	0.5672	0.4693	0.9487	0.1419	0.6888
BiGCN	Full	0.6107	0.5451	0.9574	0.2316	0.9002
	w/o LC	0.6275	0.5719	0.9546	0.2675	0.8976
	only Neg Annotations	0.6342	0.5840	0.9391	0.2869	0.8909
	only Pos Annotations	0.6540	0.6121	0.9475	0.3259	0.9010
BiGAT	Full	0.5673	0.4688	0.9452	0.1438	0.6642
	w/o LC	0.5498	0.4420	0.8853	0.1140	0.6255
	only Neg Annotations	0.5872	0.5113	0.9105	0.1945	0.6990
	only Pos Annotations	0.5705	0.4743	0.9470	0.1504	0.6636
HGCN	Full	0.6656	0.6294	0.9402	0.3536	0.8916
	w/o LC	0.6808	0.6501	0.9441	0.3844	0.8888
	only Neg Annotations	0.6704	0.6358	0.9398	0.3640	0.8904
	only Pos Annotations	0.6732	0.6396	0.9429	0.3687	0.8912
HGAT	Full	0.5662	0.4702	0.9386	0.1414	0.7030
	w/o LC	0.5661	0.4702	0.9391	0.1415	0.6985
	only Neg Annotations	0.5766	0.4891	0.9393	0.1639	0.7021
	only Pos Annotations	0.5506	0.4386	0.9599	0.1057	0.6775

6 Conclusions

The first baseline results reinforce the idea that the existence of opposing relations – and subsequently, contradictions – within a KG may provide a more accurate representation of entities. While baseline models still struggle to model opposite semantic values and contradictions effectively, these insights point to a promising shift in KGRL where the active modeling of contradictions could significantly enhance ML performance and the richness of knowledge representation. Future work will focus on developing KGRL approaches capable of modeling contradictions into entity representations to assess the impact of contradictions in KGRL and graph-based ML expressiveness, accuracy, and applicability across diverse domains and tasks.

Acknowledgments. This work was supported by FCT through fellowship 2024.01208.BD (LB), and the LASIGE Research Unit, ref. UID/00408/2025), partially supported by the KATY project which has received funding from the EU Horizon 2020 research and innovation program under grant agreement No 101017453, and also partially supported by project 41, HfPT: Health from Portugal, funded by the Portuguese Plano de Recuperação e Resiliência.

Interests. The author and supervisor, Professor Catia Pesquita, have no competing interests to declare that are relevant to the content of this article.

References

1. Liartis,J. et al.(2024). Searching for explanations of black-box classifiers in the space of semantic queries. *Semantic Web* 15 - 4, p.1085–1126
2. Jakus,G. et al.(2013). Knowledge Representation, In *Concepts, Ontologies, and Knowledge Representation*, p.47–62
3. Vrandečić,D. et al.(2014). Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57 - 10, p.78–85
4. Santos,V. et al.(2024). Can you trust Wikidata?. *Semantic Web*, p.1–22
5. Zhang,Q. et al.(2023). Integrating entity attributes for error-aware knowledge graph embedding. *IEEE Transactions on Knowledge and Data Engineering*
6. Wang,Q et al.(2017). Knowledge graph embedding: A survey of approaches and applications. *IEEE transactions on knowledge and data engineering*, 29 - 12, p.2724–2743
7. Hamilton,W. et al.(2017). Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30
8. Özcep,Ö. et al.(2023). Embedding ontologies in the description logic ALC by axis-aligned cones. *Journal of Artificial Intelligence Research*, 78 p.217–267
9. Barceló,P. et al. The logical expressiveness of graph neural networks. In *ICLR 2020*.
10. Lv,X. et al.(2022). Do pre-trained models benefit knowledge graph completion? a reliable evaluation and a reasonable approach. *ACL '22*
11. Kazemi,M. et al.(2024). Boardgameqa: A dataset for natural language reasoning with contradictory information. *Advances in Neural Information Processing Systems*, 36
12. Sousa,R. et al.(2023). Biomedical knowledge graph embeddings with negative statements. *ISWC 2023*
13. Hogan,A. et al.(2021) Knowledge graphs. *ACM Computing Surveys*, 54 - 4, p.1–37
14. Bordes,A. et al.(2013). Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*,26
15. Nickel,M. et al.(2011). A three-way model for collective learning on multi-relational data. *ICML Vol.11*
16. Yang,B. et al.(2015). Embedding entities and relations for learning and inference in knowledge bases. *ICLR 2015*
17. Ristoski,P. et al.(2016). Rdf2vec: Rdf graph embeddings for data mining. *ISWC 2016*, p.498–514
18. Huang,J. et al.(2019). Signed graph attention networks. In *ICANN 2019: Workshop and Special Sessions*, September 17–19, p.566–577
19. Schroff,F. et al.(2015). Facenet: A unified embedding for face recognition and clustering. In *CVPR 2015, IEEE*. p.815–823
20. Hoffer,E. et al.(2015). Deep metric learning using triplet network. In *SIMBAD 2015 Proceedings*, 3 p.84–92
21. Sohn,K.(2016). Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29
22. Gene Ontology Consortium.(2004). The Gene Ontology (GO) database and informatics resource. *Nucleic acids research*, 32(suppl_1), D258-D261.
23. Speer,R. et al.(2017). Conceptnet 5.5: An open multilingual graph of general knowledge. *AAAI'17*, 31, 1
24. Arnaout,H. et al.(2022). Uncommonsense: Informative negative knowledge about everyday concepts. In *CIKM 2023*, pp.37-46
25. Sousa,R. et al.(2020). Evolving knowledge graph similarity for supervised learning in complex biomedical domains. *BMC bioinformatics*, 21 p.1–19
26. Arnaout,H. et al.(2021). Wikinegata: a knowledge base with interesting negative statements. *Proceedings of the VLDB Endowment* 14(12),2807-2810.
27. Glimm,B. et al (2014). HermiT: an OWL 2 reasoner. *Journal of automated reasoning*, 53, 245-269.
28. Kipf,T. et al.(2016). Semi-supervised classification with graph convolutional networks. *ICLR 2017*
29. Veličković,P. et al.(2017). Graph attention networks. *ICLR 2018*
30. Wang,X. et al.(2019). Heterogeneous Graph Attention Network. *The world wide web conference*