Do you catch my drift?

On the usage of embedding methods to measure concept shift in knowledge graphs

Stella Verkijk
s.verkijk@vu.nl
Vrije Universiteit Amsterdam
Huygens Institute
Amsterdam, Netherlands

Ritten Roothaert
h.m.roothaert@vu.nl
Vrije Universiteit Amsterdam
Amsterdam, Netherlands

Romana Pernisch
r.pernisch@vu.nl
Vrije Universiteit Amsterdam
Discovery Lab, Elsevier
Amsterdam, Netherlands

Stefan Schlobach
k.s.schlobach@vu.nl
Vrije Universiteit Amsterdam
Amsterdam, Netherlands

ABSTRACT

Automatically detecting and measuring differences between evolving Knowledge Graphs (KGs) has been a topic of investigation for years. With the rising popularity of embedding methods, we investigate the possibility of using embeddings to detect Concept Shift in evolving KGs. Specifically, we go deeper into the usage of nearest neighbour set comparison as the basis for a similarity measure, and show why this approach is conceptually problematic. As an alternative, we explore the possibility of using clustering methods. This paper serves to (i) inform the community about the challenges that arise when using KG embeddings for the comparison of different versions of a KG specifically, (ii) investigate how this is supported by theories on knowledge representation and semantic representation in NLP and (iii) take the first steps into the direction of valuable representation of semantics within KGs for comparison.

CCS CONCEPTS

• **Computing methodologies** → *Description logics*; *Ontology engineering*; Lexical semantics.

KEYWORDS

Concept Shift, Knowledge Graph Embeddings, NLP, Semantics

ACM Reference Format:

Stella Verkijk, Ritten Roothaert, Romana Pernisch, and Stefan Schlobach. 2023. Do you catch my drift?: On the usage of embedding methods to measure concept shift in knowledge graphs. In *Knowledge Capture Conference 2023 (K-CAP '23), December 5–7, 2023, Pensacola, FL, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3587259.3627555

1 INTRODUCTION

Knowledge Graphs (KGs) are useful tools for representing factual information within a domain in a structured manner and enable knowledge inference. As a domain grows or changes, so should the KG in question. Updates in a KG can drastically change the inferences that can be made [14]. In the case of very big KGs (like DBpedia [1]), it becomes impossible for users to estimate the consequences of updates [15]. Consequently, researchers investigate ways of automatically estimating the amount of difference and the type of differences between two versions of a KG [3, 11, 16, 21]. In this paper we will refer to *concept drift* as the general notion of a change of meaning of a concept in a new version of a knowledge graph. To make things more concrete, we will showcase a specific notion of *Concept Shift* in Section 2.

Some of the methods that exist for concept drift detection are rule-based, such as SemaDrift [21] and OntoDrift [3]. While these are very transparent, they capture drift if there are changes detected on the explicit formal representation. Also, they often consider only parts of the KG, potentially missing out on important changes.

The learning of embeddings, a vector representation of KGs, has become a prominent way of using KGs where machine learning is concerned [9]. The intuition behind using KG embeddings is twofold: (i) their success with link prediction and data-mining tasks has sparked an increase in academic interest [18] and (ii) embedding methods have shown useful for *semantic shift* detection in language between corpora using Natural Language Processing (NLP). In this paper, we discuss the following two research questions:

- ***RQ1*** *Is there a fundamental difference in semantic change (and its detection) in natural language (through NLP) and KGs (specifically looking at Description Logic based ontologies?* (Section 3)
- ***RQ2*** *How transferable are NLP approaches using embeddings to determining concept change in KGs?* (Section 4)

There is a wide and confusing range of terminology to refer to semantic and concept change in evolving KGs [21]. Without restricting generalisability we use a specific formal definition (in this case of *Concept Shift* (Sec 2), before we discuss the above research questions. We found that there is a fundamental difference between semantic shift in language and Concept Shift in KGs, both from a philosophical and technological perspective. Therefore, transferring methods between the two embedding types is more challenging than initially expected. To evaluate this hypothesis, we consider a use-case based on a nearest neighbours comparison to show this.

K-CAP '23, December 5–7, 2023, Pensacola, FL, USA

ACM ISBN 979-8-4007-0141-2/23/12.
https://doi.org/10.1145/3587259.3627555

2 PRELIMINARIES

Concept Shift. We will use **Concept Shift** as an example formalisation for how to measure change of the meaning of concepts across versions of KGs. Concept shift between two versions O_o and O_n[1] exists if there is a new version B_n of a concept B that is more similar to the new version A_n of another concept A than to its original version A_o. Formally, this means $shift(A_o, A_n) = 1$ iff $\exists B_n\ sim(A_o, A_n) < sim(A_o, B_n) \land A_n \neq B_n$ The definition still leaves some core notions undefined, such as the similarity between concepts $sim(\cdot,\cdot)$, and the nature of what a concept is in the first place. Again, without loss of generality, we chose a specific formalism (Description Logic based Ontologies) as an illustrative example.

Knowledge Graphs and Description Logic Ontologies. In KGs in general the notion of concept drift is often under-defined, as there is no consensus as to what can be considered a concept in the first place, and often not even a commitment to formal semantics. We will here commit to a Description Logic perspective, as this is semantically close to some of the philosophical ideas about concept evolution and drift in the literature. Specifically, KGs are interpreted as axioms describing objects, sets and relations, in other words ontologies[2]. Following [2] we separate ontologies into terminology (Tbox) and assertions (Abox). An ontology is then defined a tuple of concepts (C), relations ($\mathcal{R}$), individuals ($\mathcal{I}$), literals ($\mathcal{L}$), as well as a set of axioms ($\mathcal{A}$). The axioms in the Tbox ($\mathcal{A}^{Tbox}$) define concepts, their hierarchies and relations between concepts, the Abox ($\mathcal{A}^{Abox}$) provides information about individuals, their type and their relations to other individuals.

It has been argued that the meaning of a concept A can be represented by three aspects, its label, intension ($int(A)$) and extension ($ext(A)$) [17, 26]. The label of a concept can be the name or URI used to refer to A. The intension consists of its properties, and traditionally, the Tbox would be where the intension of concepts is defined, while the extension of a concept represents its usage so is directly related to the Abox. The similarity relations needed to formally define Concept Shift are often based on similarities between labels (how similar are the names), extension (how similar are the objects described by a concept) and intensions (how similar are the logical properties of the description of concepts).

Calculating Concept Shift in Ontologies. Two approaches that are typically used to measure semantic change between two KGs are morphing-chain and identity-based approaches [26]. The morphing chain approach compares concepts in one version of the ontology to concepts in the other version. The identity-based approach assumes a pairing between a concept in one version to its new version in the new (possibly updated) ontology.

There are two main recent frameworks that provide a measuring-suite for assessing semantic drift between KGs: SemaDrift [21], a morphing-chain approach, and an extension of SemaDrift, OntoDrift [3], a hybrid approach. While the first measures the label aspect, the intension and the extension, the second adds the metrics URI, Subclasses, Superclasses and Equivalent classes. The advantage of these methods is their explainability. Unfortunately, the rules need to be explicitly defined and thus automatically limit the change you are measuring to what the rules can capture. The definitions also depend on the completeness of the definition of concepts in the KG.

Knowledge Graph Embeddings. As mentioned before, more often than not knowledge graphs are in practise semantically underspecified, which makes formal, symbolic, approaches to define meaning (let alone change of meaning) of concepts all but impossible. A way to deal with this problem is to represent KGs as points in vector space: KG embedding. This has proven useful in tasks such as link prediction [9, 24]. However, research is surfacing that is questioning what these embeddings really learn to represent [18, 19].

As opposed to DL ontologies, KG embedding, e.g. according to [9], consider KGs to be simply graphs (without any further commitment to specific interpretations of the nodes and relations. $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{F}\}$ with $(h, r, t) \in \mathcal{F}, h \in \mathcal{E}, t \in \mathcal{E}, r \in \mathcal{R}$. Often therefore, no additional semantics are taken as input for embedding methods like TransE [25]. While there is no proper definition of concepts in KGs, $\mathcal{G}$ is most closely related to Aboxes, where the head h and tail t entities are interpreted to be individuals in $\mathcal{I}$ (which makes sense, as in most applications of KGs Tbox information, even if available, is ignored). Embeddings are sets of vectors for the head ($\mathbf{h}$), relationship ($\mathbf{r}$) and tail ($\mathbf{t}$), where $h, r, t \in \mathbb{R}^n$, n denoting the number of dimensions in the embedding space. [3] The embedding is learned by randomly sampling facts and initiating their entities in the embedding space. The position of these entities is then optimised often in batches using a specific loss function, dependent on the embedding method. We refrain from also defining other embedding spaces and their loss functions, and refer the reader to one of the recent surveys such as the one of Ji et al. [9].

While these embedding methods capture different aspects of the KG when fitting it in latent space, there are some similarities [18]. Entities or concepts that can appear in similar contexts within the KG (i.e. the in- and outgoing relationships are similar), have a relatively small distance (both cosine and euclidean) between them in latent space. This is due to the methods used to train/fit these embeddings. Here we make a distinction between two different methods: translational methods such as TransE [25], and word2Vec based methods such as RDF2Vec [20]. Note that both of these approaches are generally used to embed the Abox of a KG.

Translational methods optimise towards the condition $\mathbf{h} + \mathbf{r} \approx \mathbf{t}$. This forces entities with similar relationships into closer proximity, clustering them. Often these entities share the same class, even if this information is not included in the learning process. Word2Vec [8] based methods achieve a similar result using different techniques. These methods are based on random walks through the KG, which are interpreted as 'sentences' and fed into Word2Vec. Each node and edge in the walk is considered a word in the sentence. There are two methods to train such a model, the first using the context of the concept to predict the concept (continuous bag of words), the second uses the concept to predict the context of the concept (skip-gram). For both methods concepts that appear in similar contexts should be difficult to distinguish from each other. Hence, the embedding of these concepts are similar.

[1] where n and o refer to the **n**ew and **o**ld version respectively

[2] It is natural to publish DL axioms as triples (s, r, o) and thus as (knowledge) graphs. We will therefore interchangeably use the terms KG and ontology

[3] This definition is specific for point-wise space like TransE, where also the translational property $\mathbf{h} + \mathbf{r} \approx \mathbf{t}$ holds.

Embedding-based Concept Drift. Since link prediction and data mining tasks with KG embeddings work well it seems like a natural step to use embedding methods on KGs to assess change between evolving KGs. Even more so since assessing **semantic shift** between corpora with word embedding models has also shown successful in Natural Language Processing (NLP). In the shared task on Linguistic Semantic Shift Detection (2022), the best performing system fine-tuned a multilingual language model (XLM-R [4]) to obtain contextualized embeddings for words, after which the degree of change for each word was calculated as the Average Pairwise Distance between the embeddings for word usages in the old and new corpus [28]. There are also other works dealing with semantic shift (e.g. [10, 27]) showing the feasibility of using word embeddings to detect it.

3 CONCEPT DRIFT IN NLP VS KGS

An important difference between semantic shift detection in NLP and concept change, such as shift, in KGs is that the first is about measuring changes in word meaning over time, i.e. assessing a shift in language use, whereas the second is about concept modelling. For example, there is an important difference between studying whether the word *car* refers to a specific type of vehicle in a specific language and studying whether the concept *car* in a KG refers to only gas-fueled vehicles in one version of a graph and in the next also includes electronically-fueled vehicles. While the first studies the interplay between sense and reference, the second focuses more on sense, following Frege [7]. To explain this further: the labels of concepts in a KG can in principle be arbitrary (they could be replaced by, for example, sequences of numbers). Note, however, that this question is related to the decenia old discussion about semantics on the web, between formal and social semantics, etc.

Algorithms like word2vec [12] capture the meaning of a word by modelling its context, e.g. since the word *cat* and *dog* both occur often within the context of *feeding* and *petting*, the algorithm sees them as similar. However, it also captures some relatedness, i.e. *Angela Merkel* and *Berlin* are often mentioned in similar contexts even though they are very different entities (a human and a city share little attributes). For the detection of semantic shift in language, this is less problematic, as shift detection in NLP is about the detection of a change in word meaning. The distributional hypothesis [6] widely adopted in NLP states that word meaning is represented by a word's context, i.e. changes in context represent a change in meaning. However, this inclusion of relatedness might be problematic for measuring shift in KGs.

Portisch et al. [18] show that RDF2Vec assigns close vectors to similar entities as well as related entities, giving the example of Merkel and Berlin being placed close together in the vector space by RDF2Vec since they both share some relationship to Germany. Portisch et al. show that node2vec, DeepWalk and KGlove follow a similar pattern and how this can be problematic for some downstream tasks: separating cities and countries in two clusters becomes difficult when a city is related to the country it is located in. However, they also mention that the problem of relatedness and similarity being mixed is not problematic when working with entities that are all of the same kind.

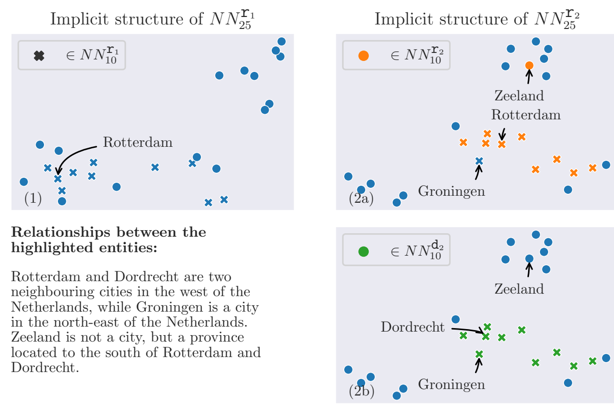

Figure 1: t-SNE visualisation of the implicit structure the nearest neighbour set NN of Rotterdam (r): $NN_{25}^{\mathbf{r}_1}$ (1) and $NN_{25}^{\mathbf{r}_2}$ (2a, 2b). In all figures, the crosses indicate the entities that are also elements of $NN_{10}^{\mathbf{r}_1}$.

In order to answer Research Question **RQ1** we introduced semantic shift in both NLP and DL-based ontologies, and provided a preliminary analysis of the two related, though different, challenges. While the first primarily studies change in the usage of words, Concept Shift in KGs studies the change of the underlying concept. The question of using algorithms based on NLP definitions of semantics depends on which embedding method you apply which in turn depends on the problem and the way knowledge is represented in the graph. It is also important to note that studying word meaning change is, naturally, a lexical task. Multiple words can refer to the same concept, however, in a KG, we capture one concept as one node.

4 DRIFT DETECTION WITH NNS

Having discussed methods to define semantic change, and characteristics of language and KGs, let us analyse some of the differences and commonalities in a more specific scenario. This section studies nearest neighbour (NN) approaches to find semantic change. Sets of NNs are often used in classification tasks to narrow down the possible candidates for better performance, instead of using all possible individuals in the embedding [5, 22]. Zhou et al. [29] successfully adopt a NN walk network embedding for link prediction. Pernisch et al. [13] introduced a measure called Embedding Resemblance Indicator (ERI) to compare/quantify how different two embeddings learned on consecutive versions of the input graph are. [13] is based on a different idea from [27] to detect shifts using word embeddings. As we believe, though, that semantic shift in NLP and KGs are fundamentally different, those methods should not be adopted without more detailed review.

When translating the Concept Shift definition into embeddings, the idea is to use a 'footprint', e.g., use NNs as a footprint [13]. However, this approach does not work out of the box for KGs because of (1) the stochasticity of embedding learning, (2) the embedding of individuals instead of concepts and (3) a questionable reduction of meaning of an entity to (only) its closest neighbours.

To study this in more detail, we performed an experiment where we use NN sets to detect Concept Shift where we know there is none[4]. We took two embeddings calculated in [13] in experiments with TransE. These two embeddings are learned on the same version of FB15k-237. We will refer to them as E_1 and E_2, where $E_1 = \mathtt{TransE}(\mathcal{G})$, and $E_2 = \mathtt{TransE}(\mathcal{G})$. Figure 1 visualises the 25 NNs of the concept *Rotterdam* (r) in the two embeddings respectively: r_1 and r_2, where $\mathsf{r}_1 \in E_1$ (shown in Figure 1.1), and $\mathsf{r}_2 \in E_2$ (shown in Figure 1.2a and 1.2b). We determined the NNs by euclidean distance and reduced the number of dimensions for visualisation purposes using t-distributed stochastic neighbourhood embeddings (t-SNE) to maintain the implicit structure of the data [23]. The set of 10 NNs of the concept Rotterdam in the first embedding ($NN_{10}^{\mathsf{r}_1}$), has an overlap of nine out of ten (9/10) entities with the concept of Rotterdam in the second embedding ($NN_{10}^{\mathsf{r}_2}$). However, there is a different concept, namely Dordrecht, in the second embedding ($\mathsf{d}_2 \in E_2$) of which the ten NNs have a complete overlap of those of Rotterdam in the first embedding, i.e., $NN_{10}^{\mathsf{d}_2}$ has a complete overlap with $NN_{10}^{\mathsf{r}_1}$, which means that $sim(\mathsf{r}_1, \mathsf{d}_2) > sim(\mathsf{r}_1, \mathsf{r}_2)$. Hence, conform our definition of Concept Shift in Section 2 this means that we have detected Concept Shift in these two embeddings E_1 and E_2 while they are learned on the same version of the same KG, and we thus know there is no shift between E_1 and E_2.

This simple example shows that a NN set is vulnerable to small changes, which is problematic given the stochastic nature of embedding methods. Further experiments showed that close to 25% of all concepts in E_2 would be falsely classified as having experienced Concept Shift in the current experimental set-up.

Another big issue of using NN sets is the usage of individuals rather than concepts in popular embedding methods like TransE [25] and RDF2vec [20]. When calculating a NN set, we take an individual as the center of the set. However, per definition of Concept Shift we are interested in capturing the change in meaning of concepts, not its individuals one by one. We therefore look to clustering methods. The idea of clusters stems from the fact that embedding methods automatically group individuals with the same relations, due to the loss functions employed that optimise representations capturing the similarity and relatedness of individuals. In the embedding space, we can thus potentially detect small clusters of individuals using this approach to calculate extension(al shift) of a concept. Additionally, should there be multiple clusters detected of the same class, one can investigate if subclasses could help distinguish entities across clusters. This can also signal concept drift, as a class could potentially have evolved into multiple subclasses. Or even, the modelling of that class has not been granular enough.

To answer **RQ2**, we investigated the prominent approach of NNs from word embeddings and its potential to be used for Concept Shift detection in KG. Unfortunately, there are multiple drawbacks and therefore, NNs are problematic to detect Concept Shift without serious consideration on the semantics captured in the KG embedding as well as the representation used in the KG itself. There is potential in investigating clustering approaches instead, as they are closely related to the extensional aspect of a concept.

[4]The code and data used for this experiment can be found at https://zenodo.org/records/10026567

5 CONCLUSIONS

It highly depends on the downstream task and the way knowledge is represented in a graph which embedding method is most useful [18]. When using KG embeddings for any downstream task, one should be aware of what is represented in a KG and how. We find that method transfer between NLP and KGs requires serious considerations about the differences between sense and reference and similarity and relatedness. While semantic shift detection in NLP studies a change in (the interplay between sense and) reference, Concept Shift detection in evolving KGs mostly studies a change in sense. Studying Firth's distributional hypothesis experimentally in the context of KGs will be interesting future research. When it comes to Concept Shift detection between KGs, we argue that using NN sets is problematic. We will move towards clustering approaches instead. Studying embedding methods besides TransE is also necessary. Our paper provides interesting initial insights into the relation between embedding methods and formal semantics, for the task of detecting Concept Shift specifically, but also in general. Our future research will start from a use-case, by choosing specific KGs and embedding methods to further investigate the feasibility of using clustering or other approaches to detect semantic shift.

REFERENCES

[1] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. In *The Semantic Web*. Springer Berlin Heidelberg, Berlin, Heidelberg, 722–735.
[2] Franz Baader, Diego Calvanese, Deborah L McGuinness, Daniele Nardi, and Peter F (Eds.). Patel-Schneider. 2007. *The Description Logic Handbook: Theory, Implementation and Applications* (2 ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511711787
[3] Giuseppe Capobianco, Danilo Cavaliere, and Sabrina Senatore. 2020. OntoDrift: a Semantic Drift Gauge for Ontology Evolution Monitoring *(CEUR Workshop Proceedings, Vol. 2821)*. CEUR-WS.org, 1–10. https://ceur-ws.org/Vol-2821/xpreface.pdf
[4] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 8440–8451. https://doi.org/10.18653/v1/2020.acl-main.747
[5] Sébastien Ferré. 2019. Link Prediction in Knowledge Graphs with Concepts of Nearest Neighbours. In *The Semantic Web*. Springer International Publishing, Cham, 84–100.
[6] John R. Firth. 1957. A synopsis of linguistic theory 1930-55. 1952-59 (1957), 1–32.
[7] Gottlob Frege. 1892. Über sinn und bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* 100 (1892), 25–50.
[8] Yoav Goldberg and Omer Levy. 2014. word2vec Explained: deriving Mikolov et al.'s negative-sampling word-embedding method. arXiv:1402.3722 http://arxiv.org/abs/1402.3722
[9] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. 2022. A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. *IEEE Transactions on Neural Networks and Learning Systems* 33, 2 (2022), 494–514.
[10] Andrey Kutuzov, Lilja Øvrelid, Terrence Szymanski, and Erik Velldal. 2018. Diachronic word embeddings and semantic shifts: a survey. In *Proceedings of the 27th International Conference on Computational Linguistics*. Association for Computational Linguistics, Santa Fe, New Mexico, USA, 1384–1397. https://aclanthology.org/C18-1117
[11] Albert Meroño Peñuela, Romana Pernisch, Christophe Guéret, and Stefan Schlobach. 2021. Multi-Domain and Explainable Prediction of Changes in Web Vocabularies. In *Proceedings of the 11th on Knowledge Capture Conference* (Virtual Event, USA) *(K-CAP '21)*. 193–200.
[12] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2* (Lake Tahoe, Nevada) *(NIPS'13)*. Curran Associates Inc., Red Hook, NY, USA, 3111–3119.
[13] Romana Pernisch, Daniele Dell'Aglio, and Abraham Bernstein. 2021. Toward Measuring the Resemblance of Embedding Models for Evolving Ontologies. In *Proceedings of the 11th on Knowledge Capture Conference*. ACM, Virtual Event

USA, 177–184. https://doi.org/10.1145/3460210.3493540

[14] Romana Pernisch, Daniele Dell'Aglio, and Abraham Bernstein. 2021. Beware of the hierarchy — An analysis of ontology evolution and the materialisation impact for biomedical ontologies. *Journal of Web Semantics* 70 (2021), 100658. https://doi.org/10.1016/j.websem.2021.100658

[15] Romana Pernisch, Daniele Dell'Aglio, Mirko Serbak, Rafael S. Gonçalves, and Abraham Bernstein. 2022. Visualising the effects of ontology changes and studying their understanding with ChImp. *Journal of Web Semantics* 74 (2022), 100715.

[16] Romana Pernischova, Daniele Dell'Aglio, Matthiew Horridge, Matthias Baumgartner, and Abraham Bernstein. 2019. Toward Predicting Impact of Changes in Evolving Knowledge Graphs. In *ISWC 2019 Posters & Demonstrations*. ISWC, CEUR-WS.org. https://doi.org/10.5167/uzh-174975

[17] Catia Pesquita, Daniel Faria, André O. Falcão, Phillip Lord, and Francisco M. Couto. 2009. Semantic Similarity in Biomedical Ontologies. *PLOS Computational Biology* 5, 7 (07 2009), 1–12. https://doi.org/10.1371/journal.pcbi.1000443

[18] Jan Portisch, Nicolas Heist, and Heiko Paulheim. 2022. Knowledge graph embedding for data mining vs. knowledge graph embedding for link prediction–two sides of the same coin? *Semantic Web* 13, 3 (2022), 399–422.

[19] Jan Portisch and Heiko Paulheim. 2022. The DLCC Node Classification Benchmark for Analyzing Knowledge Graph Embeddings. In *The Semantic Web – ISWC 2022*. Springer International Publishing, Cham, 592–609.

[20] Petar Ristoski and Heiko Paulheim. 2016. RDF2Vec: RDF Graph Embeddings for Data Mining. In *The Semantic Web – ISWC 2016*. Springer International Publishing, Cham, 498–514.

[21] Thanos G. Stavropoulos, Stelios Andreadis, Efstratios Kontopoulos, and Ioannis Kompatsiaris. 2019. SemaDrift: A hybrid method and visual tools to measure semantic drift in ontologies. *Journal of Web Semantics* 54 (2019), 87–106.

[22] Bram Steenwinckel, Gilles Vandewiele, Michael Weyns, Terencio Agozzino, Filip De Turck, and Femke Ongenae. 2022. INK: knowledge graph embeddings for node classification. *Data Mining and Knowledge Discovery* 36, 2 (March 2022), 620–667. https://doi.org/10.1007/s10618-021-00806-z

[23] Laurens J.P. van der Maaten and Geoffrey E. Hinton. 2008. Visualizing High-Dimensional Data Using t-SNE. *Journal of Machine Learning Research* 9, nov (2008), 2579–2605.

[24] Meihong Wang, Linling Qiu, and Xiaoli Wang. 2021. A survey on knowledge graph embeddings for link prediction. *Symmetry* 13, 3 (2021), 485.

[25] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering* 29, 12 (2017), 2724–2743.

[26] Shenghui Wang, Stefan Schlobach, and Michel C. A. Klein. 2011. Concept drift and how to identify it. *J. Web Semant.* 9, 3 (2011), 247–265.

[27] Anna Wegmann, Florian Lemmerich, and Markus Strohmaier. 2020. Detecting Different Forms of Semantic Shift in Word Embeddings via Paradigmatic and Syntagmatic Association Changes. In *The Semantic Web – ISWC 2020*. Springer International Publishing, Cham, 619–635.

[28] Frank D. Zamora-Reina, Felipe Bravo-Marquez, and Dominik Schlechtweg. 2022. LSCDiscovery: A shared task on semantic change discovery and detection in Spanish. In *Proceedings of the 3rd Workshop on Computational Approaches to Historical Language Change*. Association for Computational Linguistics, Dublin, Ireland, 149–164. https://doi.org/10.18653/v1/2022.lchange-1.16

[29] Mingqiang Zhou, Qizhi Han, Mengjiao Li, Kunpeng Li, and Zhiyuan Qian. 2023. Nearest neighbor walk network embedding for link prediction in complex networks. *Physica A: Statistical Mechanics and its Applications* 620 (2023), 128757.