

CAFE: Fact Checking in Knowledge Graphs using Neighborhood-Aware Features

Agustín Borrego^{a,*}, Daniel Ayala^a, Inma Hernández^a, Carlos R. Rivero^b and David Ruiz^a

^a *University of Seville, Av. Reina Mercedes s/n, Seville, Spain*

E-mails: borrego@us.es, dayala1@us.es, inmahernandez@us.es, druiz@us.es

^b *Rochester Institute of Technology, 92 Lomb Memorial Drive, Rochester NY, USA*

E-mail: crr@cs.rit.edu

Abstract. Knowledge Graphs (KGs) currently contain a vast amount of structured information in the form of entities and relations. Because KGs are often constructed automatically by means of information extraction processes, they may miss information that was either not present in the original source or not successfully extracted. As a result, KGs might potentially lack useful and valuable information. Current approaches that aim to complete missing information in KGs either have a dependence on embedded representations, which hinders their scalability and applicability to different KGs; or are based on long random paths that may not cover relevant information by mere chance, since exhaustively analyzing all possible paths of a large length between entities is very time-consuming. In this paper, we present an approach to completing KGs based on evaluating candidate triples using a novel set of features, which exploits the highly relational nature of KGs by analyzing the entities and relations surrounding any given pair of entities. Our results show that our proposal is able to identify correct triples with a higher effectiveness than other state-of-the-art approaches (up to 60% higher precision or 20% higher recall in some datasets).

Keywords: Knowledge Graph Completion, Fact Checking, Machine Learning

1. Introduction

Knowledge Graphs (KGs) are vast repositories of structured information whose growth over the past years can be related to that of the Web of Data [1]. Large volumes of information about different domains can be found in some well-known KGs such as DBpedia [2], NELL [3], Freebase [4] or the Google Knowledge Vault [5], which can then be used for tasks such as question answering [6].

KGs are most commonly built by extracting non-structured [3, 5] or semi-structured [2, 7] information from web sources, though some smaller KGs can be manually curated by domain experts [8]. After information extraction systems have been applied to extract knowledge from online sources, said information is then semantized [9, 10] and stored in a KG as entities and relations between those entities. Regardless of the specific process by which a KG is constructed, said KG

usually lacks a certain amount of information, either because said information was not originally present in the original information source, or because it was not extracted or semantized correctly [11]. Because of this inherent incompleteness, KGs operate under the Open World Assumption, i.e., a piece of information that is not contained in a KG is not considered to be incorrect, but rather just unknown [12]. Therefore, it is mandatory to refine KGs after their creation in order to expand the knowledge they contain, so that they contain more accurate information [13]. Such is the example of the Google Knowledge Vault [5], a KG that is commonly used for question answering and that is under constant expansion and refinement.

As a motivational example, Figure 1 presents a KG that contains information about fictional works, actors, writers and characters. This KG is incomplete, as it is missing certain information that is true in the real world, such as the fact that *Daniel Radcliffe* played the character *Harry Potter*.

*Corresponding author. E-mail: borrego@us.es.

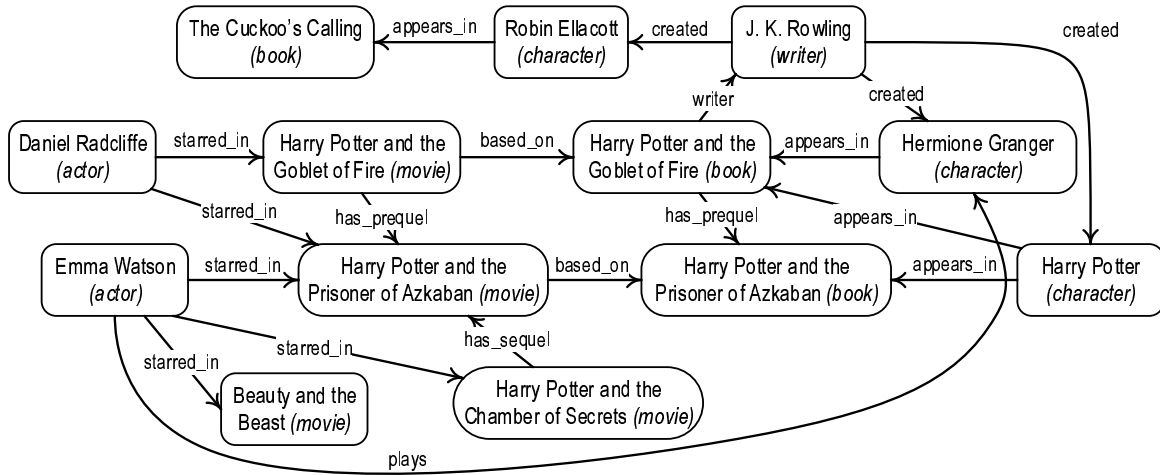


Fig. 1. Example KG describing works, actors, writers and characters

Deriving additional knowledge from an existing KG to augment it is a task known as Knowledge Graph completion [13]. KG completion can be approached with two complementary goals in mind: completing information about entity types [14], or completing the relations between entities [15–18]. Furthermore, these techniques can be based either on external features, i.e., features that depend on external information sources, or on purely internal features, i.e., features that are computed using only the KG itself [19, 20]. In this work, we focus in KG completion using only internal features, since this line of work has shown promising results without dependencies to external sources of information.

When completing a KG, one has to produce a set of facts that are not contained in said KG but that have a reasonable chance of being true. This is usually done through a prediction model that has been trained beforehand. The task of creating, training and applying such a model is known as fact checking [21, 22]. Thus, it is desirable for these models to achieve a high precision to ensure that the facts that are considered to be correct are actually true in the real world [23].

There are several fact checking proposals in the literature that use solely internal features [12, 15, 18, 20, 24, 25], which we discuss in more detail in Section 2. Unfortunately, these approaches can suffer from a number of drawbacks. Some of them rely on embedded representations of the entities and/or relations, and both these representations and the associated neural models need to be completely recomputed whenever a few new triples, which is a relatively frequent event [3, 5]. Additionally, there exists a number of different,

random walk-based proposals that may miss relevant information due to their non-deterministic nature.

In this paper we present CAFE, our approach for Fact Checking using Neighborhood-Aware Features, which uses a novel feature set that extracts information from the neighborhoods of the entities, i.e., their nearby entities and relations. This feature set is used to transform triples in the KG into feature vectors, which are fed to neural prediction models that discern between correct triples that should be added to the KG, and incorrect ones that should be disregarded. Different than embedding-based proposals, the features CAFE is based on can be directly applied to any KG without the need of computing embedded representations, and do not require to be recomputed as the KG grows. They are also not based on random paths, which ensures that all contextual information is always taken into account. Furthermore, the experimental results show that CAFE is able to predict which relations should be added to a KG with a higher precision than other state-of-the-art approaches.

The rest of this paper is organized as follows: Section 2 analyzes the related work in KG completion proposals, Section 3 describes our terminology and the set of features that CAFE uses, Section 4 presents our proposal for training predictive models for triples in a KG, Section 5 reports our experimental results; finally, Section 6 presents our conclusions.

2. Related Work

In the field of Fact Checking, there are rule-based, embedding-based and path-based proposals:

1 Rule-based proposals rely on discovering logical
 2 rules that determine when a given triple is correct.
 3 In this regard, some authors propose using Induc-
 4 tive Logic Programming to mine Horn rules that ex-
 5 press the relations between entities in a KG, and then
 6 applying said rules to produce new explicit knowl-
 7 edge [16, 26], under the assumption that all triples that
 8 match the rules are correct. The performance of the ob-
 9 tained rules is usually measured using metrics such as
 10 rule support or confidence, or a variant of these metrics
 11 [12].

12 Those that are based on embeddings aim to evaluate
 13 possible relations in a KG by learning embedded rep-
 14 resentations of its entities and relations, either by using
 15 them as feature vectors for a prediction model [15], or
 16 by performing different transformations in an embed-
 17 ding space [18, 20, 25, 27, 28]. The resulting embed-
 18 ding space (or spaces) is subsequently used to evalu-
 19 ate the likelihood of a candidate triple to be correct or
 20 incorrect, since entities that are supposed to be related
 21 by means of a certain relation are expected to be close
 22 to each other in the embedding space. A similar line of
 23 work includes representing a KG as a tensor, and then
 24 factorizing it to obtain latent, more compact represen-
 25 tations of the triples contained within [19, 29]. These
 26 proposals usually rank triples by decreasing order of
 27 likelihood, according to how close the two entities in
 28 the triples are in the generated space. Thus, they are
 29 commonly evaluated using metrics that are based on
 30 some ranking of likely correct candidate triples [30],
 31 however, a likelihood threshold can be set to obtain
 32 classical classification measures such as precision, re-
 33 call and F1 [31].

34 Finally, those that are path-based exploit the highly
 35 relational nature of KGs to learn how to predict new
 36 relations between entities. Our approach, CAFE, falls
 37 under this category. In this line of work, Lao and Co-
 38 hen [17] introduce the Path Ranking Algorithm (PRA),
 39 a two-step process to find which paths may be use-
 40 ful to predict a certain relation. An evolution of PRA
 41 named Subgraph Feature Extraction (SFE) is proposed
 42 in [24] by Gardner and Mitchell. SFE achieves a bet-
 43 ter performance than PRA and produced more expres-
 44 sive results while running faster than PRA. It also in-
 45 troduced a handmade “Alias” relation, which relates
 46 entities in the same KG that have different labels but
 47 refer to the same element in the real world. In [32],
 48 Mazumder and Liu propose a random walk-based ap-
 49 proach using neighborhood-guided path finding, where
 50 semantic similarities between entities are computed by
 51 applying a Word2vec-based embedding model on the

1 names of the entities. Reinforcement learning has also
 2 been used to find valuable paths that can help to suc-
 3 cessfully complete a KG [33]. The evaluation metrics
 4 used by these proposals are similar to those used by
 5 the proposals based on embeddings.

6 Nowadays, there exist a number of KG completion
 7 proposals that combine different approaches. For ex-
 8 ample, in [34] the authors propose obtaining embed-
 9 dings of the entities and the relations, and then com-
 10 bining these embeddings in the forms of paths, which
 11 can be helpful for determining whether a certain triple
 12 is correct. A different way of obtaining embedded rep-
 13 resentations of the elements in a KG is presented in
 14 [30], where they use graph neural networks to capture
 15 information pertaining the structure of the graph. En-
 16 tity embeddings have also been used to aid in the pro-
 17 duction of KG completion rules [35].

3. Neighborhood-Aware Features

22 In this Section, we first introduce some preliminary
 23 concepts that are necessary to understand our proposal,
 24 and then we define our set of neighborhood-aware fea-
 25 tures.

3.1. Preliminaries

29 We now define some preliminary concepts and the
 30 notation that is used throughout the paper.

32 **Definition 1. Triple:** Let \mathcal{E} be a set of entities, and
 33 let \mathcal{R} be a set of relations. We define a triple as a 3-
 34 tuple that represents the existence of a relation $r \in \mathcal{R}$
 35 between a source entity $s \in \mathcal{E}$ and a target entity $t \in \mathcal{E}$.
 36 We denote triples as (s, r, t) .

37 In the example KG depicted in Figure 1, a sample
 38 triple is $(Emma\ Watson, starred_in, Beauty\ and\ the\ Beast)$.

41 **Definition 2. Knowledge Graph:** Let \mathcal{E} be a set of
 42 entities, let \mathcal{R} be a set of relations, and let \mathcal{T} be a set
 43 of triples of the form $\{(s, r, t) \mid s, t \in \mathcal{E}, r \in \mathcal{R}\}$. We
 44 define a Knowledge Graph as a 3-tuple $\mathcal{KG} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$.

46 Figure 1 graphically represents a KG, with 13 enti-
 47 ties and 7 distinct relations.

49 **Definition 3. Path between entities:** Let $\mathcal{KG} =$
 50 $(\mathcal{E}, \mathcal{R}, \mathcal{T})$ be a Knowledge Graph, and let $s, t \in \mathcal{E}$
 51 be two entities in \mathcal{KG} . We define a path p between

1 s and t as a sequence of triples of the form $p =$
 2 $\langle (e_i, r_i, e_{i+1}) \rangle$ for $i = 1..n$, where $e_1 = s$, $e_{n+1} = t$
 3 and $(e_i, r_i, e_{i+1}) \in \mathcal{T}$ for $i = 1..n$. We define the length
 4 of a path as the number of triples it contains, i.e., $|p|$.
 5 We denote a path p of length n between s and t as
 6 $path(s, t, r_1, r_2, \dots, r_n)$, or $path_n(s, t)$ for short. We de-
 7 note the set of all possible distinct paths of the form
 8 $path(s, t, r_1, r_2, \dots, r_n)$ as $\mathcal{P}(s, t, r_1, r_2, \dots, r_n)$.

9 In the KG depicted in Figure 1, an example path
 10 of length 2 between the entities *J.K. Rowling* and *The*
 11 *Cuckoo's Calling* is $\langle (J.K. Rowling, created, Robin
 12 *Ellacott), (Robin Ellacott, appears_in, The Cuckoo's*
 13 *Calling) \rangle*$.

14 **Definition 4. Reachability:** Let $\mathcal{KG} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ be a
 15 Knowledge Graph, let $s, t \in \mathcal{E}$ be two entities in \mathcal{KG} ,
 16 let $r \in \mathcal{R}$ be a relation in \mathcal{KG} , and let $n \geq 1$ be a
 17 natural number. We define reachability as a predicate
 18 that determines whether there exists a path of length
 19 n between s and t in \mathcal{KG} such that the relation r
 20 appears in the last triple of the path, i.e., $Reach(\mathcal{KG}, s, t,$
 21 $r, n) \iff \exists path_n(s, t) \wedge \exists a \in \mathcal{E} \mid last(path_n(s, t)) =$
 22 (a, r, t) . We define the set of entities that can be reached
 23 from s through a relation r at distance n as the set of
 24 entities that match the predicate *Reach* under such cir-
 25 cumstances, i.e., $\{t \in \mathcal{E} \mid Reach(\mathcal{KG}, s, t, r, n)\}$. We
 26 denote the previously defined set as $Reachable(s, r, n)$.

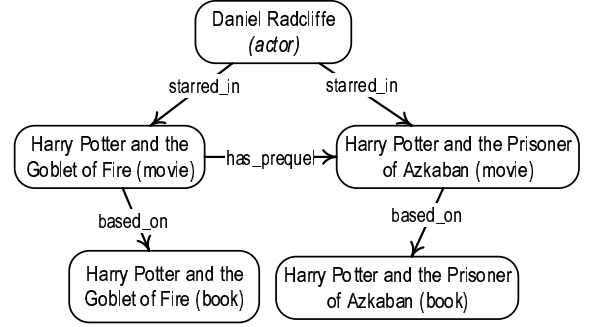
27 In the example KG depicted in Figure 1,
 28 $Reachable(Hermione Granger, writer, 2) =$
 29 $\{Hermione Granger, J.K. Rowling\}$.

30 **Definition 5. Neighborhood subgraph:** Let $\mathcal{KG} =$
 31 $(\mathcal{E}, \mathcal{R}, \mathcal{T})$ be a Knowledge Graph, let $e \in \mathcal{E}$ be an en-
 32 tity in \mathcal{KG} , and let $n \geq 1$ be a natural number. We
 33 define the neighborhood subgraph of e of size n as a
 34 Knowledge Graph $\mathcal{KG}_e^n = (\mathcal{E}_e^n, \mathcal{R}, \mathcal{T}_e^n)$ that contains
 35 the triples whose target entities can be reached from e
 36 at a distance of at most n through any relation, and the
 37 entity set that can be derived from such triples, where
 38 $\mathcal{T}_e^n = \{(s', r', t') \in \mathcal{T} \mid Reach(\mathcal{KG}, e, t', r', i), i =$
 39 $1..n\}$ and $\mathcal{E}_e^n = \bigcup \{s, t\} \subseteq \mathcal{E} \mid (s, r, t) \in \mathcal{T}_e^n\}$.

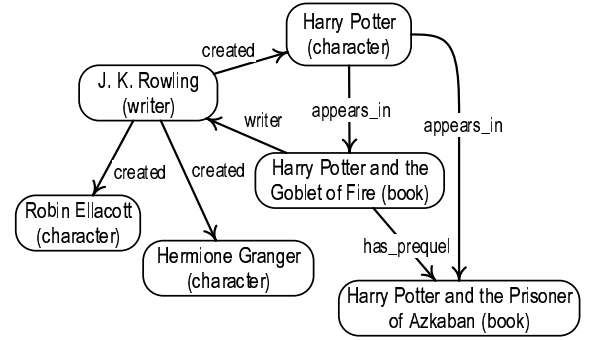
40 Figure 2 illustrates this definition, with two possible
 41 neighborhood subgraphs for the KG shown in Figure
 42 1.

43 3.2. Feature set

44 We devised a neighborhood-aware set of features
 45 which takes neighborhood subgraphs, reachable enti-
 46 ties and paths into account. Due to the large number



(a) Neighborhood subgraph of size 2 for the entity *Daniel Radcliffe*



(b) Neighborhood subgraph of size 3 for the entity *Harry Potter*

Fig. 2. Two neighborhood subgraphs for the example KG depicted in Figure 1

of possible variations of each feature, we present our
 feature set in terms of groups of features, where each
 group can be parameterized to obtain a specific feature,
 which we call an instance of the feature group.

Definition 6. Feature: Let $\mathcal{KG} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ be a
 Knowledge Graph. We define a feature f as a function
 $f : \mathcal{T} \rightarrow \mathbb{R}$ that assigns a real number to a triple.

For example, a feature f may convert a triple into the
 number of entities in the neighborhood subgraph
 of size 2 of the source entity, i.e., $f : (s, r, t) \mapsto |\mathcal{E}_s^2|$.

Definition 7. Feature group: Let $\mathcal{KG} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ be
 a Knowledge Graph. We define a feature group f_n as
 a function $f_n : \mathcal{X} \rightarrow (\mathcal{T} \rightarrow \mathbb{R})$ that receives a set of
 parameters \mathcal{X} and returns a feature.

For example, a feature group f_0 may return a feature
 that converts a triple into the number of entities in the
 neighborhood subgraph of size n of the source entity,
 i.e., $f_0(n) = f : (s, r, t) \mapsto |\mathcal{E}_s^n|$, where n is a parameter
 of the feature group. Thus, $f_0(2) : (s, r, t) \mapsto |\mathcal{E}_s^2|$,
 which is the feature shown in the previous example.

Consequently, feature groups allow us to represent a set of very similar features in a more compact way, where the only distinction between said features is a given set of parameters.

In the following, we present the feature groups that CAFE uses. For the sake of clarity, we illustrate a possible instance of every feature group and its value using the example triple $example = (Daniel\ Radcliffe, plays, Harry\ Potter)$, the KG shown in Figure 1 and the neighborhood subgraphs described in Figure 2.

Feature group f_1 : Number of entities in the neighborhood subgraph of size n of the source entity in the triple.

Features in this group can be computed as $f_1(n) : (s, r, t) \mapsto |\mathcal{E}_s^n|$.

In the example shown in Fig. 2(a), $f_1(2)$ applied to the example triple is $|\{Daniel\ Radcliffe, Harry\ Potter\ and\ the\ Goblet\ of\ Fire\ (movie), Harry\ Potter\ and\ the\ Prisoner\ of\ Azkaban\ (movie), Harry\ Potter\ and\ the\ Goblet\ of\ Fire\ (book), Harry\ Potter\ and\ the\ Prisoner\ of\ Azkaban\ (book)\}| = 5$.

Feature group f_2 : Number of entities in the neighborhood subgraph of size n of the target entity in the triple.

Features in this group can be computed as $f_2(n) : (s, r, t) \mapsto |\mathcal{E}_t^n|$.

In the example shown in Fig. 2(b), $f_2(3)$ applied to the example triple is $|\{Harry\ Potter, J.K.\ Rowling, Harry\ Potter\ and\ the\ Goblet\ of\ Fire\ (book), Harry\ Potter\ and\ the\ Prisoner\ of\ Azkaban\ (book), Robin\ Ellacott, Hermione\ Granger\}| = 6$.

Feature group f_3 : Number of common entities between the neighborhood subgraph of size n of the source entity and the neighborhood subgraph of size m of the target entity in the triple.

Features in this group can be computed as $f_3(n, m) : (s, r, t) \mapsto |\mathcal{E}_s^n \cap \mathcal{E}_t^m|$.

In the example shown in Fig. 2, $f_3(2, 3)$ applied to the example triple is $|\{Harry\ Potter\ and\ the\ Goblet\ of\ Fire\ (book), Harry\ Potter\ and\ the\ Prisoner\ of\ Azkaban\ (book)\}| = 2$.

Feature group f_4 : Jaccard index of similarity between the entity sets for the neighborhood subgraph of size n of the source entity and the neighborhood subgraph of size m of the target entity in the triple.

Features in this group can be computed as $f_4(n, m) : (s, r, t) \mapsto jaccard(\mathcal{E}_s^n, \mathcal{E}_t^m)$.

In the example shown in Fig. 2, $f_4(2, 3)$ applied to

the example triple is $|\{Harry\ Potter\ and\ the\ Goblet\ of\ Fire\ (book), Harry\ Potter\ and\ the\ Prisoner\ of\ Azkaban\ (book)\}| / |\{Daniel\ Radcliffe, Harry\ Potter\ and\ the\ Goblet\ of\ Fire\ (movie), Harry\ Potter\ and\ the\ Prisoner\ of\ Azkaban\ (movie), Harry\ Potter\ and\ the\ Goblet\ of\ Fire\ (book), Harry\ Potter\ and\ the\ Prisoner\ of\ Azkaban\ (book), Harry\ Potter, J.K.\ Rowling, Robin\ Ellacott, Hermione\ Granger\}| = 2 / 9 = 0.22$.

Feature group f_5 : Number of reachable entities through the relation r at distance n from the source entity in the triple.

Features in this group can be computed as $f_5(r, n) : (s, r, t) \mapsto |Reachable(s, r, n)|$.

In the example shown in Fig. 2(a), $f_5(2, hasPrequel)$ applied to the example triple is $|\{Daniel\ Radcliffe, Harry\ Potter\ and\ the\ Prisoner\ of\ Azkaban\ (movie)\}| = 2$.

Feature group f_6 : Number of reachable entities through the relation r at distance n from the target entity in the triple.

Features in this group can be computed as $f_6(r, n) : (s, r, t) \mapsto |Reachable(t, r, n)|$.

In the example shown in Fig. 2(b), $f_6(2, created)$ applied to the example triple is $|\{Harry\ Potter, Hermione\ Granger, Robin\ Ellacott\}| = 3$.

Feature group f_7 : Number of common reachable entities through the relation r from the source entity at distance n and from the target entity at distance m .

Features in this group can be computed as $f_7(r, n, m) : (s, r, t) \mapsto |Reachable(s, r, n) \cap Reachable(t, r, m)|$.

In the example shown in Fig. 2, $f_7(2, 3, created)$ applied to the example triple is $|\{Daniel\ Radcliffe\} \cap \{Harry\ Potter, Hermione\ Granger, Robin\ Ellacott\}| = 0$.

Feature group f_8 : Jaccard index of similarity between the reachable entities through the relation r from the source entity at distance n and those reachable through the relation r from the target entity at distance m .

Features in this group can be computed as $f_8(r, n, m) : (s, r, t) \mapsto jaccard(Reachable(s, r, n), Reachable(t, r, m))$.

In the example shown in Fig. 2, $f_8(2, 3, created)$ applied to the example triple is $|\emptyset| / |\{Harry\ Potter, Hermione\ Granger, Robin\ Ellacott\}| = 0 / 3 = 0$.

1 *Feature group f_9 : Number of distinct paths of length n*
 2 *between the source and the target entity in the triple,*
 3 *using relations r_1, \dots, r_n .*

4 Features in this group can be computed as $f_9(n, r_1,$
 5 $\dots, r_n) : (s, r, t) \mapsto |\mathcal{P}(s, t, r_1, \dots, r_n)|$.

6 In the example shown in Fig. 1, $f_9(4, \textit{starred_in},$
 7 $\textit{based_on}, \textit{writer}, \textit{created})$ applied to the *example*
 8 *triple is 1, as there is one path of length 4 between the*
 9 *entities *Daniel Radcliffe* and *Harry Potter* that matches*
 10 *the given relations.*

11
 12 The rationale behind this set of feature groups is
 13 mani: first, computing the common entities and indices
 14 of similarities of the neighborhood subgraphs for the
 15 source and target entities allows CAFE to find out the
 16 degree of overlap that exists in the neighborhoods of
 17 the two relations, along with detailed information on
 18 whether this overlap occurs for different relations. Sec-
 19 ond, studying all kinds of paths that exist between the
 20 entities in a triple allows CAFE to also benefit from
 21 path-based information in a deterministic manner. Fi-
 22 nally, they can be applied at any time as a KG grows
 23 with new entities and relations, without the need of a
 24 complete recomputation in opposition to embedding-
 25 based approaches.

26 4. Our Proposal

27
 28
 29 Our proposal, CAFE, receives a KG and a set of re-
 30 lations from that KG as input, and outputs a classifi-
 31 cation model for each of the selected relations. These
 32 models are able to determine if a given triple that rep-
 33 represents an instance of the relation is correct, i.e., if
 34 it should belong to the KG. Our workflow is visually de-
 35 picted in Figure 3 and, in the following subsections,
 36 we describe each one of its steps.

37 4.1. Loading the KG

38
 39
 40 We are interested only in the triples contained in a
 41 KG, as both the entities and the relations sets can be
 42 derived from them. CAFE internally stores the KG it
 43 is given in the form of (s, r, t) triples, using an efficient
 44 graph-like structure. The triples that contain relations
 45 for which a predictive model does not need to be gen-
 46 erated will still be taken into account when comput-
 47 ing features, since they may provide valuable predic-
 48 tive information.

42 4.2. Generating negative examples

43
 44 A KG contains only positive information, i.e., it
 45 contains examples of the occurrence of a relation r be-
 46 tween two entities, but it does not contain explicit in-
 47 formation about pairs of entities for which r does not
 48 hold. Since our proposal relies on a neural prediction
 49 model that requires negative examples during its train-
 50 ing, a number of them has to be produced. To do this,
 51 we follow the same approach as other authors and con-
 sider any triple not present in the KG to be a negative
 example [15, 25, 32]. To train our models, we gener-
 ate one negative example for each triple existing in the
 KG by replacing its target entity t with a different one,
 t' , such that the resulting triple (s, r, t') does not exist
 in the KG. To preserve the range of each relation, we
 randomly choose a t' such that there exists some other
 triple in the KG where t' appears as the target entity for
 the relation r . A positive or negative label is then added
 to every triple, denoting whether it represents positive
 or negative evidence.

In the example depicted in Figure 1, a valid nega-
 tive example is $(\textit{Hermione Granger}, \textit{appears_in}, \textit{The}$
 $\textit{Cuckoo's Calling})$, since we know that *The Cuckoo's*
Calling is a valid target for the relation *appears_in*,
 while $(\textit{Hermione Granger}, \textit{appears_in}, \textit{Daniel Rad-}$
 $\textit{cliffe})$ would not be allowed as a negative example, be-
 cause *Daniel Radcliffe* never appears as the target of
 the relation *appears_in*.

It can be argued that generating negative evidence
 in this manner can produce false negatives by mere
 chance, i.e., statements that are deemed false because
 they do not exist in the KG, but that are true in the
 real world. While this is indeed plausible, we as other
 authors [15, 18, 20] consider that the chances of this
 happening are low enough for it to not have a notice-
 able outcome in our results. Furthermore, it may even
 prevent our prediction models from overfitting to the
 available evidence.

43 4.3. Converting triples into feature vectors

44
 45 Once negative examples have been generated, our
 46 feature set is instantiated and applied to all triples. For
 47 all feature groups, we obtain all possible feature in-
 48 stances by applying the Cartesian product of its possi-
 49 ble parameters but, for efficiency reasons, we limit pa-
 50 rameters that represent neighborhood sizes to a max-
 51 imum value of 3. Each feature instance then assigns
 each triple a real number, and thus applying several
 features to a triple results in a feature vector, in which

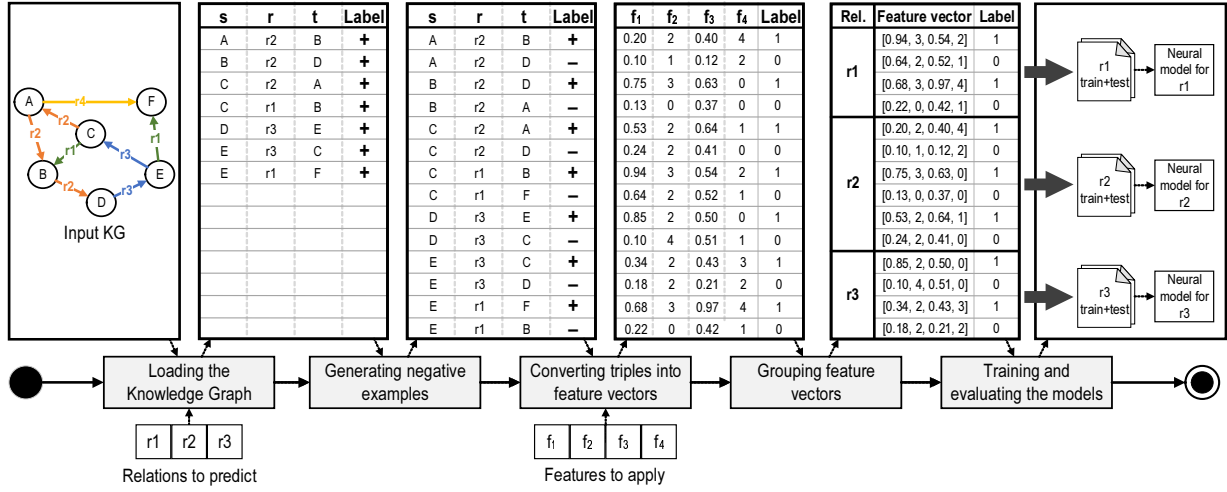


Fig. 3. Workflow of CAFE

each position represents the value that the corresponding feature assigns to the triple. The labels that were added in the previous step, which indicate if the triple is a positive or a negative example, are converted into a numeric label.

It is important to note that, to compute features on a triple, we temporarily remove said triple from the KG (in case it is a positive example and thus exists in the KG), since not doing so would result in trivial prediction models such as “a person plays a character if there exists a triple in the KG stating that the person plays that character”.

4.4. Grouping feature vectors

The labeled feature vectors that result from the previous step are grouped by the relation that was present in the original triples, resulting in a group of labeled vectors for every relation for which we train a prediction model. For example, the feature vector produced by the triple (*Daniel Radcliffe, plays, Harry Potter*) is added to the set of feature vectors for the relation *plays* alongside its label, which denotes whether it is a positive or a negative example. These groups are then split into training and test sets, resulting in two sets of labeled feature vectors per relation.

4.5. Training and evaluating the models

For every relation that we predict, we create one or more neural models (whose hyperparameters can be found in Section 5.2), where each model is fed the fea-

tures that are obtained from a certain neighborhood size. Thus, using only neighborhood subgraphs of size 1 results in one model, using neighborhood subgraphs of size of up to 2 results in two models, and so on. This allows each model to capture the specific information that every neighborhood size may yield. To combine two or more models, we use an additional layer of three neurons and an output layer of one neuron.

Said models are trained using the labeled feature vectors in the training split for the desired relation, where each model receives only the features corresponding to its assigned neighborhood size, and the label or ground truth is shared among them. Prior to training our models, we first remove any individual features that have the exact same value in every feature vector and thus lack any predictive power. This is due to some features being relatively sparse (especially path-based features, since only a small subset of all possible paths of fixed length occur between two given entities), and thus helps the model focus only on potentially useful features.

5. Evaluation

In this section, we present our experimental results for several datasets, and we compare our proposal against other state-of-the-art fact checking techniques. In the following, we introduce the datasets that we used for evaluation, our experimental settings and, finally, our results.

5.1. Datasets

We evaluated our proposal using four datasets for KG completion provided by the freely available AYNEC-DataGen [36] tool: FB13-A, WN11-AR, WN18-AR and NELL-AR. These datasets are based on the well-known FB13, WN11 [15], WN18 [37] and a subset of NELL proposed in [24]. However, they have been processed to remove reciprocal relations detected by AYNEC (that is, relations r and r' where if (s, r, t) exists then (t, r', s) also exists very frequently). Additionally, relations that make up for less than 5% of the total amount of triples in the graph have been removed.

These datasets originally contained one negative example per each positive one in both their training and testing splits. In order to study how the KG completion techniques perform when presented with a much higher number of negative evidence, we created versions of these datasets whose testing splits contain, whenever possible, 10 negative examples per each positive one using the AYNEC-DataGen tool. Exceptions may apply for relations where the set of possible target entities is very small, e.g., the relation *gender*. In these cases we generated as many negative examples as possible, up to 10 per positive example. We believe that this is a more realistic scenario, since a real-world KG completion task would be presented with a much higher number of false candidate triples than true ones [21]. To avoid confusion, we denote these versions as FB13-A-10, WN11-AR-10, WN18-AR-10 and NELL-AR-10. For the sake of reproducibility, our source code is publicly available on GitHub¹.

For FB13-A-10, WN11-AR-10 and WN18-AR-10, we aimed to predict all possible relations, and for NELL-AR-10 we focused on the same subset of 10 relations that were used to evaluate SFE [24]. However, in the later dataset, one relation was removed by AYNEC for being the reciprocal of another relation, leaving 9 relations for evaluation. In the specific case of FB13-A-10, we transferred some training triples over to the testing set in order to provide testing examples for some relations, as they were not available in the original dataset as introduced in [15]. This was done in a random manner, and 1 in 4 such triples was transferred to the testing set. Table 1 provides an overview on the aforementioned datasets.

¹<https://github.com/tdg-seville/CAFE>

5.2. Evaluation method

A neural prediction model is created for every relation of interest and trained using its corresponding training set. Then, the model is applied to all feature vectors in the test set, and we compare the expected label (which denotes whether it represents a valid triple or not) against the label that was produced by our model. We report our results in terms of precision, recall and F1, in order to determine how effective our proposal is when it determines that a triple should be added to the KG. We computed these metrics using the AYNEC-ResTest tool [36].

We evaluated three versions of CAFE, denoted CAFE₁ to CAFE₃, which were limited to using feature instances that exploited neighborhood subgraphs and paths of a maximum size of 1, 2, and 3, respectively. This was done in order to study how using larger neighborhoods affects the performance of CAFE. In the specific case of the NELL dataset, we only used CAFE₁ and CAFE₂. The reason for this is that the larger amount of total relations made it impractical to reach neighborhoods of size 3, as the number of features that would have to be computed was large.

We use TransE [38], TransD [20], TransH [25], TransR [18] and ComplEx [39] as baselines for our evaluation, since they are well-known state-of-the-art KG completion proposals. In order to evaluate these different embedding-based proposals in a fair and level manner, we use the OpenKE [31] framework to train and evaluate these proposals using our datasets. Additionally, since these baselines proposals usually employ metrics such as MRR and Precision@N, we used the utilities provided by the OpenKE tool to obtain binary labels for the testing triples by allowing it to set an automatic likelihood threshold. The overall metrics evaluation was carried out using AYNEC-ResTest.

To measure how using different neighborhood sizes influences the performance of CAFE, we evaluated CAFE₁, CAFE₂ and CAFE₃ independently against each other in all proposed metrics. Also, to determine if CAFE is comparable to the previously mentioned state-of-the-art proposals, we used the best-performing variant of CAFE (in terms of better average F1) for comparison against other proposals.

We selected the following values for the hyperparameters of our neural models: 3 layers with 1024, 512 and 256 neurons each, learning rate of 0.001, batch size of 8, dropout of 0.1 for all layers, 50 epochs and validation ratio of 10%. When two or more models were to be combined, we joined their results using a

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

Dataset	Training triples	Test triples	Entities	Relations
FB13-A-10	228,172	481,457	74,998	13
WN11-AR-10	77,948	198,231	38,195	9
WN18-AR-10	71,984	183,051	40,943	11
NELL-AR-10	86,971	219,374	53,934	148 (9)

Table 1

Metadata for the datasets we used for evaluation

hidden layer with 3 neurons and an output layer with a single neuron. These values for the hyperparameters were chosen using a hold-out or “dev” set for the FB13-A-10 dataset, and all datasets were then evaluated using these hyperparameters. We chose them because they provided satisfactory results in our empirical tests and, to the best of our knowledge, there are no proposals for automatically choosing the optimal parameters for a neural network. Our experiments were conducted on a computer equipped with an Intel Core i9-9900K CPU, 64GB of RAM and an Nvidia RTX 2080 Ti GPU.

5.3. Results and discussion

In Figure 4 we show the obtained results in terms of precision, recall and F1 for CAFE₁ through CAFE₃. Those obtained when comparing the embedding-based proposals to the best-performing variant of CAFE for every dataset are shown in Figure 5. This was done by selecting the CAFE variant that achieved a higher average F1 in each dataset. Additionally, in Figure 6 we show the detailed results per every relation in all datasets.

Regarding the question of how using different neighborhood sizes affects the performance of CAFE, Fig. 4 shows that performance metrics generally seem to improve for CAFE₂ with respect to CAFE₁, but the same cannot be said for CAFE₃ and CAFE₂. Indeed, metrics appear to remain stagnant or even decrease when using larger neighborhoods in some cases. For example, in FB13-A-10 (Fig. 4(a)) we do not observe a significant increase in effectiveness when using neighborhood subgraphs of size 3 in contrast of those of size 2. Similar remarks can be made for WN11-AR-10 (Fig. 4(c)) and WN18-AR-10 (Fig. 4(d)). A possible conclusion for this is that, at a certain point, larger neighborhood subgraphs do not provide additional value or predictive power over smaller ones. In a worst case scenario, the number of features with little to no predictive power would greatly increase by using

larger neighborhood sizes, negatively affecting the performance of CAFE. This effect actually occurs in the case of the NELL-AR-10 dataset, where performance worsens when increasing the size of the neighborhood subgraphs for which we compute features. We believe that a possible explanation for this is that NELL is a noisier dataset [40], and thus taking larger neighborhoods into account significantly increases the amount of noise that our classification model has to deal with, effectively hindering its performance.

Our results also show that CAFE is able to match the performance of state-of-the-art embedding-based proposals, and in many cases achieve higher values on the metrics under evaluation. In the case of FB13-A-10 (Fig. 4(a)), CAFE can achieve a wide range of values for the different metrics, however, it is able to reach higher performance values than the embedding-based proposals under evaluation for all metrics (Fig. 5(a)). In other datasets, such as WN11-AR-10 (Fig. 5(c)) and WN18-AR-10 (Fig. 5(d)), CAFE obtains a more satisfactory performance in a more consistent manner. These results show that CAFE can be more effective than other proposals when presented with a high amount of negative evidence to evaluate.

Figure 6 displays that, in general, both a satisfactory precision and recall can be achieved, and thus we consider that CAFE is generally effective. However, there exists a number of relations for which a very high precision value is obtained at the expense of a lower recall, resulting in a typical precision-recall trade-off. We have also observed that, in the specific case of the FB13-A-10 dataset (Fig. 5(a)), one finds an unusually higher number of classification results with a low precision value. We hypothesize that this is due to the dataset in question having relations that are hard to predict using only other information present in it. Such is the case of the relation *cause_of_death*, since learning to predict the cause of the death of a person with a very high effectiveness would be a remarkable achievement that unfortunately falls out of the scope of this work.

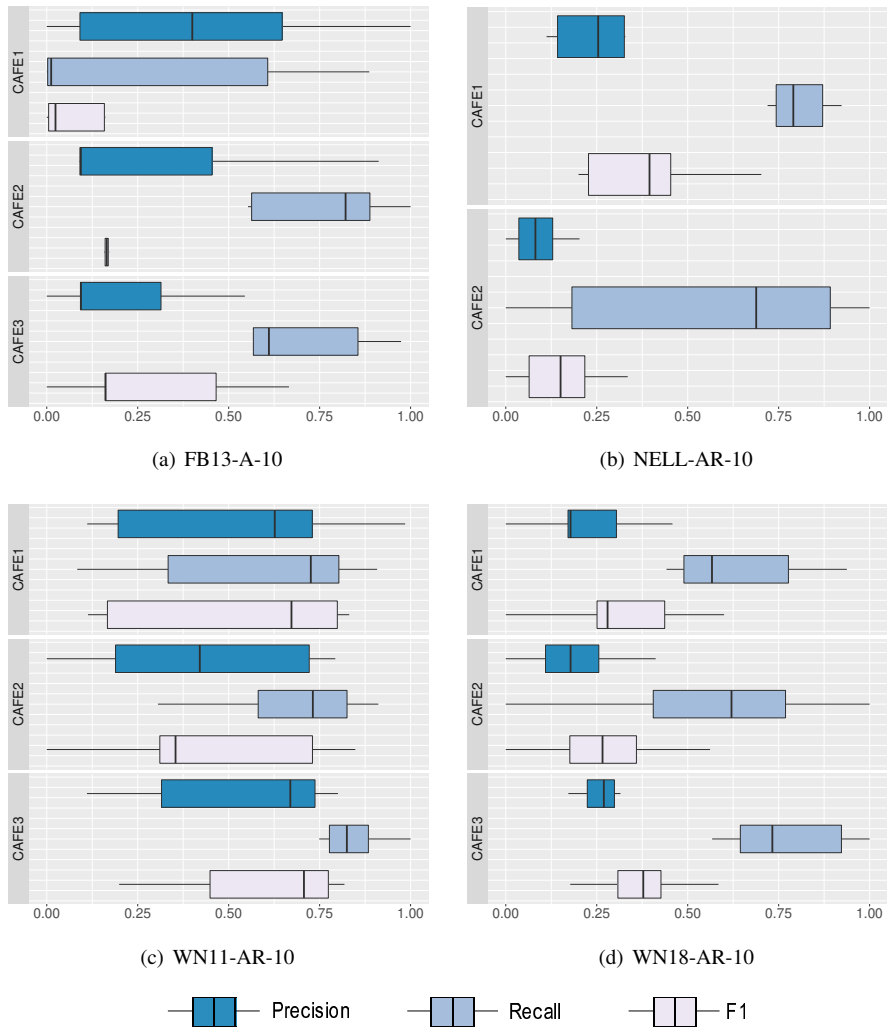


Fig. 4. Performance comparison between the different CAFE variants

5.4. Limitations

Despite CAFE being generally effective, it has limitations. Since the number of features associated with every neighborhood size does not increase linearly, but rather at least quadratically, the computational cost associated with using larger neighborhoods to compute features can escalate quickly. Additionally, CAFE does not work well for relations where useful predictive information cannot always be found in the neighborhoods of the entities in a triple (for example, certain values of 0.00 in Fig. 6). In these cases, even if some information is present in the neighborhoods under consideration, it may be not successfully captured by CAFE due to the high amount of irrelevant data surrounding it. Also, the NELL-AR-10 dataset shows

that larger neighborhood sizes are not always better for predictive purposes, since the amount of noise they may include can be detrimental for the performance of CAFE. This can be due to some source or target entities being present in many triples (for example, countries), and thus larger neighborhoods of these entities introduce many other entities that are not relevant at all to the triple under evaluation. In these cases, it is up to the user to decide which maximum neighborhood size best caters to their interests.

6. Conclusions

In this paper we introduced CAFE, a proposal for fact checking in Knowledge Graphs that uses a novel



Fig. 5. Performance comparison between CAFE and other state-of-the-art proposals

feature set based upon neighborhoods and neural prediction models. This feature set, and by extension our proposal, is applicable to all KGs, and it is completely deterministic. Additionally, CAFE does not need any manually supplied or external information, nor does it need any previous preprocessing of the KG, and it is able to operate using a KG as its only input.

We have performed a number of experiments to evaluate the performance of CAFE on four well-known Knowledge Graphs, identifying candidate triples that should be added to each KG. To provide a more realistic evaluation scenario, we performed our experiments

using testing splits that contain a much lower amount of true triples than false ones. Our results show that applying the proposed feature set yields a performance result that is comparable to state-of-the-art embedding-based techniques, and in many cases even higher. In all datasets under study, CAFE achieves higher average precision values, thus leading to a more trustworthy fact checking process. We also show that using information present in the neighborhoods of the two entities in a triple provides satisfactory results, and thus local information existing around any given entity appears to be of great value.

- [6] A. Bordes, S. Chopra and J. Weston, Question Answering with Subgraph Embeddings, in: *EMNLP, ACL*, 2014, pp. 615–620. <http://aclweb.org/anthology/D/D14/D14-1067.pdf>.
- [7] M. Glass and A. Gliozzo, A Dataset for Web-Scale Knowledge Base Population, in: *ESWC*, Springer, 2018, pp. 256–271. doi:10.1007/978-3-319-93417-4_17.
- [8] G.A. Miller, WordNet: A Lexical Database for English, *Commun. ACM* **38**(11) (1995), 39–41. doi:10.1145/219717.219748.
- [9] D. Ayala, I. Hernández, D. Ruiz and M. Toro, TAPON: A two-phase machine learning approach for semantic labelling, *Knowledge-Based Systems* **163** (2019), 931–943. doi:10.1016/j.knosys.2018.10.017.
- [10] S. Neumaier, J. Umbrich, J.X. Parreira and A. Polleres, Multi-level Semantic Labelling of Numerical Values, in: *ISWC*, Vol. 9981, 2016, pp. 428–445. doi:10.1007/978-3-319-46523-4_26.
- [11] A. Bordes and E. Gabrilovich, Constructing and Mining Web-scale Knowledge Graphs, in: *SIGKDD*, ACM, 2014, pp. 1967–1967. doi:10.1145/2623330.2630803.
- [12] L. Galárraga, C. Teflioudi, K. Hose and F.M. Suchanek, Fast rule mining in ontological knowledge bases with AMIE+, *VLDB J.* **24**(6) (2015), 707–730. doi:10.1007/s00778-015-0394-1.
- [13] H. Paulheim, Knowledge graph refinement: A survey of approaches and evaluation methods, *Semantic Web* **8**(3) (2017), 489–508. doi:10.3233/SW-160218.
- [14] A. Neelakantan and M.-W. Chang, Inferring Missing Entity Type Instances for Knowledge Base Completion: New Dataset and Methods, in: *HLT-NAACL*, The Association for Computational Linguistics, 2015, pp. 515–525. <https://aclweb.org/anthology/N/N15/N15-1054.pdf>.
- [15] R. Socher, D. Chen, C.D. Manning and A. Ng, Reasoning with neural tensor networks for knowledge base completion, in: *NIPS*, 2013, pp. 926–934.
- [16] L.A. Galárraga, C. Teflioudi, K. Hose and F. Suchanek, AMIE: Association rule mining under incomplete evidence in ontological knowledge bases, in: *WWW*, ACM, 2013, pp. 413–422. doi:10.1145/2488388.2488425.
- [17] N. Lao and W.W. Cohen, Relational retrieval using a combination of path-constrained random walks, *Machine Learning* **81**(1) (2010), 53–67. doi:10.1007/s10994-010-5205-8.
- [18] Y. Lin, Z. Liu, M. Sun, Y. Liu and X. Zhu, Learning entity and relation embeddings for knowledge graph completion., in: *AAAI*, Vol. 15, AAAI Press, 2015, pp. 2181–2187. <https://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9571>.
- [19] L. Drumond, S. Rendle and L. Schmidt-Thieme, Predicting RDF triples in incomplete knowledge bases with tensor factorization, in: *SAC*, ACM, 2012, pp. 326–331. doi:10.1145/2245276.2245341.
- [20] G. Ji, S. He, L. Xu, K. Liu and J. Zhao, Knowledge graph embedding via dynamic mapping matrix, in: *ACL (1)*, The Association for Computer Linguistics, 2015, pp. 687–696. <https://aclweb.org/anthology/P/P15/P15-1067.pdf>.
- [21] A. Borrego, D. Ayala, I. Hernández, C.R. Rivero and D. Ruiz, Generating Rules to Filter Candidate Triples for their Correctness Checking by Knowledge Graph Completion Techniques, in: *K-CAP*, 2019, pp. 115–122. doi:10.1145/3360901.3364418.
- [22] P. Lin, Q. Song and Y. Wu, Fact Checking in Knowledge Graphs with Ontological Subgraph Patterns, *Data Science and Engineering* **3** (2018), 341–358. doi:10.1007/s41019-018-0082-4.
- [23] M. Nentwig, M. Hartung, A.-C. Ngonga Ngomo and E. Rahm, A survey of current link discovery frameworks, *Semantic Web* **8**(3) (2017), 419–436. doi:10.3233/SW-150210.
- [24] M. Gardner and T. Mitchell, Efficient and expressive knowledge base completion using subgraph feature extraction, in: *EMNLP*, The Association for Computational Linguistics, 2015, pp. 1488–1498. <https://aclweb.org/anthology/D/D15/D15-1173.pdf>.
- [25] Z. Wang, J. Zhang, J. Feng and Z. Chen, Knowledge Graph Embedding by Translating on Hyperplanes, in: *AAAI*, Vol. 14, AAAI Press, 2014, pp. 1112–1119. <https://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/view/8531>.
- [26] K. Kolthoff and A. Dutta, Semantic relation composition in large scale knowledge bases, in: *LD4IE@ISWC*, Vol. 1467, CEUR-WS.org, 2015, pp. 34–47. http://ceur-ws.org/Vol-1467/LD4IE2015_Kolthoff.pdf.
- [27] B. Shi and T. Weninger, Open-World Knowledge Graph Completion, in: *AAAI*, 2018, pp. 1957–1964. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16055>.
- [28] B. Shi and T. Weninger, ProjE: Embedding Projection for Knowledge Graph Completion, in: *AAAI*, 2017, pp. 1236–1242. <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14279>.
- [29] M. Nickel, V. Tresp and H.-P. Kriegel, Factorizing YAGO: scalable machine learning for linked data, in: *WWW*, ACM, 2012, pp. 271–280. doi:10.1145/2187836.2187874.
- [30] Z. Wang, Z. Ren, C. He, P. Zhang and Y. Hu, Robust embedding with multi-level structures for link prediction, in: *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, AAAI Press, 2019, pp. 5240–5246.
- [31] X. Han, S. Cao, X. Lv, Y. Lin, Z. Liu, M. Sun and J. Li, OpenKE: An Open Toolkit for Knowledge Embedding, in: *EMNLP*, 2018, pp. 139–144.
- [32] S. Mazumder and B. Liu, Context-aware Path Ranking for Knowledge Base Completion, in: *IJCAI*, AAAI Press, 2017, pp. 1195–1201. doi:10.24963/ijcai.2017/166.
- [33] W. Xiong, T. Hoang and W.Y. Wang, DeepPath: A reinforcement learning method for knowledge graph reasoning, in: *EMNLP*, 2017.
- [34] Y. Shen, D. Wen, Y. Li, N. Du, H.-t. Zheng and M. Yang, Path-based attribute-aware representation learning for relation prediction, in: *Proceedings of the 2019 SIAM International Conference on Data Mining*, SIAM, 2019, pp. 639–647.
- [35] V.T. Ho, D. Stepanova, M.H. Gad-Elrab, E. Kharlamov and G. Weikum, Learning rules from incomplete kgs using embeddings, in: *The 17th International Semantic Web Conference*, ceur. ws. org, 2018.
- [36] D. Ayala, A. Borrego, I. Hernández, C.R. Rivero and D. Ruiz, AYNEC: All You Need for Evaluating Completion Techniques in Knowledge Graphs, in: *ESWC*, 2019, pp. 397–411. doi:10.1007/978-3-030-21348-0_26.
- [37] A. Bordes, X. Glorot, J. Weston and Y. Bengio, A semantic matching energy function for learning with multi-relational data - Application to word-sense disambiguation, *Machine Learning* **94**(2) (2014), 233–259. doi:10.1007/s10994-013-5363-6.

- [38] A. Bordes, N. Usunier, A. García-Durán, J. Weston and O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: *NIPS*, 2013, pp. 2787–2795.
- [39] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier and G. Bouchard, Complex Embeddings for Simple Link Prediction, in: *ICML*, Vol. 48, 2016, pp. 2071–2080. <http://proceedings.mlr.press/v48/trouillon16.html>.
- [40] H. Paulheim and C. Bizer, Improving the quality of linked data using statistical distributions, *Int. J. Semantic Web Inf. Syst.* **10**(2) (2014), 63–86. doi:10.4018/ijswis.2014040104.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51