

Error detection in Knowledge Graphs: Path Ranking, Embeddings or both?

Romanos Fasoulis · Konstantinos Bougiatiotis ·
Fotis Aisopos · Anastasios Nentidis · Georgios
Paliouras

the date of receipt and acceptance should be inserted later

Abstract This paper attempts to compare and combine different approaches for detecting errors in Knowledge Graphs. Knowledge Graphs constitute a mainstream approach for the representation of relational information on big heterogeneous data, however, they may contain a big amount of imputed noise when constructed automatically. To address this problem, different error detection methodologies have been proposed, mainly focusing on path ranking and representation learning. This work presents various mainstream approaches and proposes a novel hybrid and modular methodology for the task. We compare these methods on two benchmarks and one real-world biomedical publications dataset, showcasing the potential of our approach and drawing insights regarding the state-of-art in error detection in Knowledge Graphs.

1 Introduction

A Knowledge Graph (KG) is a construct for representing relational information between entities that can be extracted either manually or automatically (e.g. from text found online). Each piece of information is usually presented as a triple (s, r, o) , where s is the subject, o is the object, and r is the relation connecting them. Every such triple is also called a fact (Wang et al. 2017).

In the last decade, as the Natural Language Processing (NLP) domain is growing rapidly, we have seen a surge of growth in automatic knowledge graph construction and development. Knowledge graphs like DBpedia (Auer et al. 2007), Wikidata (Tanon et al. 2016), NELL (Carlson et al. 2010) and YAGO (Suchanek et al. 2007)

Romanos Fasoulis · Konstantinos Bougiatiotis · Fotis Aisopos · Anastasios Nentidis · Georgios Paliouras
Institute of Informatics and Telecommunications, National Center Scientific Research Demokritos, Athens,
Greece

E-mail: Romas-92@hotmail.com

Konstantinos Bougiatiotis

Department of Informatics and Telecommunications, National and Kapodistrian University of Athens,
Greece,

E-mail: kbogas@di.uoa.gr

are automatically created, without any manual supervision, which was the preferred method for constructing KG's until recently (Bollacker et al. 2008). While the aforementioned knowledge graphs are bigger and more detailed than ever before, it is apparent that errors and noise cannot be avoided, as automatic extraction tools are not perfect. Consequently, it is important to know how much noise exists in automatically constructed KG's, how prominent it is and how does this noise affects any further downstream tasks that will be performed on the KG (e.g. link prediction or nodes classification).

To address the problem of noise and errors in knowledge graphs, we focus on comparing and contrasting various methods that use different techniques concerning error detection. Although many papers focus on KG completion and link prediction tasks, few actually deal with the problem of noise, while many are making the assumption that the KG is free of noise, something that is usually far from the truth.

Our objective is to compare and combine different approaches, mainly stemming from path ranking analysis and graph embeddings. The latter have gained significant attention because of their ability to preserve KG structure while simplifying manipulation, as well as their performance in downstream tasks (Wang et al. 2017). In addition, we propose a hybrid of the two approaches. We suggest that a specific combination of path ranking algorithms and embeddings can achieve better results on error detection tasks in some cases, while also creating robust-to-noise embeddings that can later be used for further analysis and tasks.

The main contributions of this work are concluded as follows:

- Quantitative and qualitative comparison and assessment of different error detection methods based on path ranking, representation learning and hybrid techniques.
- Development of a generic hybrid error detection approach that combines other error detection methods to enhance results.
- A framework that extends the use of path ranking algorithms to assist the generation of error-robust embeddings that can also be used in other embedding related tasks.

The rest of the paper is structured as follows: Section 2 formulates the problem addressed in the rest of the document. Section 3 presents the basic approaches used for error detection in Knowledge Graphs, while Section 4 analyses the methods employed in the current work, as well as the hybrid approach proposed. Lastly, Section 5 presents the experiments performed and the related results and finally Section 6 the conclusions of the current work.

2 Problem Formulation

Before presenting the various existing approaches aiming for error detection in Knowledge Graphs, it is important to provide a formal definition of the problem.

As a starting point, a knowledge graph G is defined as a set of triples. Each triple follows the form of (s, r, o) , where $(s, o) \in E$ are the entities and $r \in R$ is the relation that binds them. E is the set that contains all entities that exist in the knowledge graph G and R is the set that contains all relations. We assume that the knowledge graph

G also contains some ratio of noise $N\%$, denoting that $N\%$ of the triples in G are erroneous. These erroneous triples are essentially wrong edges between subjects s and objects o (both $\in E$) with a relation $r \in R$ connecting them. Thus, our objective is to find a way to pinpoint these errors in G .

3 Related Work

Error detection tasks in knowledge graphs become more and more prominent as modern, automated ways of constructing knowledge graphs create higher demands regarding data integrity (Heindorf et al. 2016). There are a handful of methods for error detection in knowledge graphs and each one may target various types of information (Paulheim 2016). This information can be internal and present in knowledge graph (e.g. density, structure, etc.) or external (e.g. textual information). A good example of an internal method is SDValidate (Paulheim and Bizer 2014), which uses the characteristic distribution of types (concepts) and relations. On the contrary, methods like DeFacto (Lehmann et al. 2012), which specialize in finding erroneous relations, uses external information in the form of lexicalizations. In this paper, we will mainly focus on two types of methods: path ranking methods and graph embedding methods, as well as various combinations of these two approaches.

3.1 Path Ranking Methods

The Path Ranking Algorithm (PRA) (Lao and Cohen 2010) is a method that can discover complex patterns in relational data, applying logistic regression over paths between nodes that are used as features. These paths are extracted through feature selection (random walks over the graph). For each triple in the KG, each path is assigned a weight that reflects the probability of arriving to the triple’s targeted object, given the triple’s subject and the path.

Sub-graph Feature Extraction (SFE) is an improvement of PRA, proposed by Gardner and Mitchell (Gardner and Mitchell 2015), aiming to reduce overall complexity, run-time and achieve statistical superiority. Novel improvements in SFE comprise the replacement of PRA’s path probability with binary values that reflect the ability to go to the triple’s object from the triple’s subject, as well as the replacement of the Random Walks method with a Breadth-First Search (BFS) algorithm.

PaTyBRED (Paths and Types with Binary Relevance for Error Detection) (Melo and Paulheim 2017) is also a path ranking approach, that improves on PRA and SFE. Specifically, the authors propose the use of Random Forests as a classifier instead of logistic regression. Additionally, they introduce a K-best selection method before training the classifiers and some heuristic measures for quicker and more robust feature extraction.

3.2 Embedding methods

While triples are considered to be a very effective structural representation of KGs, the need to better manage and gain access to underlying symbolic information of these triples led to the representation known as knowledge graph embeddings (Cai et al. 2017). This involves the transformation of entities and relations of knowledge graphs into lower-dimensional, fixed-size vectors. These vectors can afterward be employed for further downstream tasks such as link prediction, node classification, as well as error detection, which is the main focus of this paper.

The first family of embedding models is the Translational Distance Models (Wang et al. 2017), with TransE being one of the first to be introduced (Bordes et al. 2013). TransE reflects the idea that the embedding vector of the subject plus the embedding vector of the relation is very close to the vector of the object for a specific triple, meaning $s + r \approx o$. This relationship between the triple's components defines the Translational Requirement.

TransH (Wang et al. 2014) and TransR (Lin et al. 2015b) expand upon TransE model's relation-specific approach, by modeling a relation as a translating operation on a hyper-plane or even modeling relations and entities in two distinct spaces. This allows for the two models to deal with 1-to-N, N-to-1 and N-to-N relations that TransE cannot manage. TransM (Fan et al. 2014) and TransF (Feng et al. 2016) also try to deal with this problem by relaxing the Translation Requirement mentioned above, by pre-calculating the relational mapping property and allowing flexible translations, respectively.

The second category of embedding models is the Semantic Matching models, which match latent semantics of entities and relations to measure the confidence of different facts (Wang et al. 2017). One of the first semantic matching models to be introduced was RESCAL (Nickel et al. 2011), which deals with the factorization of the tensor produced by the KG. DistMult (Yang et al. 2014), ComplEx (Trouillon et al. 2016), HolE (Nickel et al. 2015) and ProjE (Shi and Weninger 2016) all concern different simplifications or extensions of the RESCAL model. A detailed description and formulation of the different embedding models can be found in (Wang et al. 2017).

3.3 Embeddings guided by internal/external info

It can be argued that all of the embedding methods mentioned in the previous section use topological characteristics of the KG to construct the embeddings of its entities and relations. However, there exist approaches that utilize additional information regarding the entities and the relations found in the KG, to improve performance and expressiveness. Here we will present some of these approaches that have been used in error detection tasks.

One type of internal information that has already been discussed is the path, or paths, connecting a subject to an object. The PTransE model extends TransE, using paths in addition to relations (Lin et al. 2015a). It replaces the relation in the Translational Requirement with each path connecting a subject and the object, creating many more Translational Requirements to be satisfied and energy functions to be

minimized. Similarly, the Confidence-aware KRL framework (CKRL) (Xie et al. 2017) introduces a triple confidence score that guides the loss function to pay attention to more "convincing" triples. This confidence score takes into account different aspects and characteristics of triples, both local and global. The Triple trustworthiness measurement model for knowledge graph (KGTm) (Jia et al. 2018) uses a crisscrossed neural network-based structure, combining different elements through a multi-layer perceptron fusioner to generate confidence scores for each triple.

Other methods may use other kinds of external information to guide the embeddings. TRESICAL extends the RESCAL model by employing additional external information like entity types/concepts, range and domain restrictions to improve performance (Nickel et al. 2011). The KALE (Guo et al. 2016) and RUGE (Guo et al. 2017) models, as well as the models proposed in (Wang et al. 2018) and (Wang et al. 2019) use logic rules and Horn clauses to guide embeddings creation and optimization. In contrast to external resources like textual information, rules are internal information, therefore, always available with extraction tools like AMIE+ (Galárraga et al. 2015).

4 Methods Employed

In this section, we briefly present the methods that we will examine, and we describe how we use these methods to improve results by generating embeddings more robust to noise.

4.1 PaTyBRED

PaTyBRED (Melo and Paulheim 2017) is a PRA-inspired algorithm that was developed with the task of error detection in mind. Therefore, we opt to use this PRA-variant in the context of error detection. PaTyBRED uses paths as features, with a path being defined as a sequence of relations $r_1 \rightarrow r_2 \rightarrow \dots \rightarrow r_n$. A subject s and an object o can be connected by a path $P(s, o)$ if there exist entities x_1, x_2, \dots, x_{n-1} such that $P(s, o) = r_1(s, x_1) \rightarrow \dots \rightarrow r_i(x_{i-1}, x_i) \rightarrow \dots \rightarrow r_n(x_{n-1}, o)$. The concept of the algorithm is to use these paths as features to decide whether a given triple is noise or not. Using heuristics, some of these paths are pruned and not taken into consideration, in order to improve complexity, while simultaneously discarding any irrelevant features. Whenever available, types/concepts of triples are also added as features.

Once the paths are pruned, N_R feature tables are populated, where N_R is the number of relations. The rows of each feature table correspond to the (s, o) tuples, linked with the corresponding relation. Columns represent the extracted/pruned paths and the types/concepts used as features. Feature tables are populated with 0 and 1 indicating whether the specific path connects the (s, o) tuple, or whether a triple is of a certain type/concept. After using N_R different classifiers, one for each relation, a confidence score with values [0-1] is decided for each triple, with low scores indicating noise.

4.2 TransE

The basic idea behind the TransE (Bordes et al. 2013) model is that, given a triple (s, r, o) that is true, the subject s and the relation r can be connected with the object o with low error, meaning $s + r \approx o$. Each component of entity (s, r, o) is represented through an embedding, a vector of fixed size d . Thus, the energy function $E(s, r, o)$ of TransE is:

$$E(s, r, o) = \|s + r - o\|_{l_{1|2}} \quad (1)$$

where $l_{1|2}$ denotes the L_1 norm or the L_2 norm respectively. The higher the degree of fitness between a subject, a relation and an object, the smaller the value of the energy function. The embeddings of the entities and relation are learned through training. Specifically, TransE recursively minimizes a pairwise scoring function that uses the aforementioned energy function and negative sampling for training:

$$Loss = \sum_{(s,r,o) \in S} \sum_{(s',r,o') \in S'} [\gamma + E^+ - E^-]_+ \quad (2)$$

where $E^+ = E(s, r, o)$ is the energy function score of a positive triple from dataset S^+ , $E^- = E(s', r, o')$ is the energy function score of a negative triple from the negative set S' generated by random sampling and γ is the hyper-parameter of margin. $[x]_+$ denotes the positive part of x as this loss function is a *max-margin* one.

4.3 Confidence-aware KRL (CKRL)

The Confidence-aware KRL framework - (CKRL) (Xie et al. 2017) injects a triple confidence measure $C(s, r, o)$ in the pairwise loss function of the TransE model, intending to learn better knowledge representations. The proposed pairwise function now becomes:

$$Loss = \sum_{(s,r,o) \in S} C(s, r, o) \sum_{(s',r',o') \in S'} [\gamma + E^+ - E^-]_+ \quad (3)$$

The triple confidence $C(s, r, o)$ measure reflects the model's ability to pay attention to triples that are more likely to be true. Specifically, when $C(s, r, o)$ is large the loss function is greatly affected by the specific triple as opposed to a triple with a small score.

The value of $C(s, r, o)$ captures local characteristics through a Local Triple Confidence (LT) measure, and global ones through Prior Path Confidence (PP) and Adaptive Path Confidence (AP). (Xie et al. 2017) describes in great detail these measures, as well as the whole mathematical formulation of the CKRL method. The final confidence score is defined using 3 different parameters λ_1 , λ_2 and λ_3 that reflect the importance of each measure:

$$C(s, r, o) = \lambda_1 LT(s, r, o) + \lambda_2 PP(s, r, o) + \lambda_3 AP(s, r, o) \quad (4)$$

As a side note, in our evaluation, in addition to CKRL, we also use the PTransE method (Lin et al. 2015a), a predecessor to the CKRL method using paths to guide embeddings, albeit in a different way than CKRL.




Fig. 1 PRGE method outline.

4.4 Path Ranking Guided Embeddings (PRGE)

As described above, any path ranking algorithm outputs a confidence score for each triple, given the path features. The CKRL algorithm uses different and more sophisticated measures, to estimate a confidence score that is used afterward in the loss function. Consequently, path ranking methods and the CKRL confidence score $C(s, r, o)$, although not similar in the range of values and derivation, score the triples in a similar manner.

Given that the two algorithms use some form of confidence measures, we can replace the dynamic triple confidence measure derived by the CKRL method with just a weight that is extracted through a path ranking method. This action is very analogous to the improvement that PaTyBRED brought to the PRA algorithm. By transforming probabilities to 0-1 values (PRA to SFE) and pruning paths/change classifiers (SFE to PaTyBRED), the whole procedure is not only simplified, but it also produces improved results (Gardner and Mitchell 2015; Melo and Paulheim 2017). Therefore, we choose to use that kind of a simplification to the above loss function version of CKRL, adopting a much simpler version of the derived confidence score.

As described above, CKRL uses 3 different measures: semantic similarity, local confidence and global confidence measures to guide the construction of the embeddings. Each of these measures is weighted differently, with many and different parameters affecting the score. Moreover, the confidence score used by CKRL is being altered during the training process, as it depends on the whole graph embeddings in each training epoch. We chose to use a path ranking score instead of the CKRL confidence score $C(s, r, o)$, as it is simplified and constant for any triple throughout the training process and rejected the option of using any probabilistic measures, path/relation embeddings or local relationships.

We also leverage the role that the confidence score and the pairwise max-margin loss function (2) play in training the embeddings. From the CKRL loss function, it is apparent that the TransE energy function and the CKRL confidence measure are in the same order of magnitude. By using an additional exponential parameter in the confidence score, we can determine how much the confidence score value will affect each subject, object and relation embedding during training.

These two ideas are thus applied to the pair-wise loss function, replacing the triple confidence $C(s, r, o)$ measure of CKRL with the confidence measure $P(s, r, o)$ of a path ranking method, while adding a parameter λ to scale the importance of the path ranking value. Thus, we propose this hybrid approach of Path Ranking Guided Embedding (PRGE). The pair-wise loss function of PRGE evolves into the following function:

$$Loss = \sum_{(s,r,o) \in S} \sum_{(s',r',o') \in S'} ([\gamma + E^+ - E^-]_+) \cdot P(s, r, o)^\lambda \quad (5)$$

While there are no restrictions to what path ranking method should be employed for $P(s, r, o)$, we opted on using the PaTyBRED method to calculate these scores, as it produces more accurate scores, while being the most simplified and robust of the PRA methods (Melo and Paulheim 2017).

5 Experiments

5.1 Datasets

We evaluate the aforementioned methods on the task of error detection. We perform experiments on two commonly used KG datasets and a KG created for a real-world application. Statistics for all the datasets are presented in Table 1:

WN18: Wordnet (Fellbaum 2005) is an English database, which can be seen as a dictionary, as well as a thesaurus. Syn-sets, a term which reflects sets of nouns, verbs, adjectives, etc., are actually the database’s entities. Relations define lexical connections between these entities. The WN18 dataset that is employed in the current experiments is a subset of Wordnet and is used as a benchmark in multiple studies.

FB15k: Freebase (Bollacker et al. 2008) is a large-scale, collaborative knowledge base that contains general facts about the real-world. While the whole database contains some millions of entities and at least a billion of triplets, we are using a sub-graph of Freebase, named FB15k. FB15k is pretty dense, and all of its entities are present in the Wiki-links database.

Dementia PubMed (Dementia): In order to demonstrate the need for error detection methodologies in real-world applications, we experimented with a Knowledge Graph created in the context of the iASiS Project (Krithara et al. 2019). For the needs of the project, we extracted relations between biomedical entities from abstracts of publications related to Dementia in PubMed¹ using automatic tools. Specifically,

¹ <https://www.ncbi.nlm.nih.gov/pubmed/>

Table 1 Datasets info

Dataset	# Rel	# Ent	# Triples
WN18	18	40,943	141,442
FB15k	1345	14,951	483,142
Dementia	64	48,008	135,000

after fetching abstracts related to Dementia, through semantic MeSH queries², we use SemRep (Rindflesch and Fiszman 2003) for extracting biomedical predications, i.e. semantic triples in the form of subject-predicate-object, from unstructured text. The subject and object arguments in these predications are concepts from the Unified Medical Language System (Bodenreider 2004) and the predicate is one of the semantic relations of the Semantic Network (McCray 2003), connecting the semantic types of the subject and object in the context of the specific sentence. For constructing the graph, 68,791 publications were fetched using the related MeSH term *Dementia*. The statistics of the final dataset generated are visible in Table 1. More details regarding the exact procedure followed for the creation of this knowledge graph and the extraction process can be found in (Nentidis et al. 2019).

5.2 Error Imputation Protocol

To assess the methodologies presented and proposed, we need noise to be present in the KG. However, there are no explicitly-labeled noisy triples in FB15K or WN18. Therefore, we generated new datasets with different percentages of noise levels to simulate real-world knowledge graphs constructed automatically. In order to do so, we construct negative triples following different approaches. The basic idea behind the error imputation process is that for each positive triple (s, r, o) in the dataset, we generate a noisy one by corrupting either s or o . For the FB15K knowledge graph, we follow the procedure described in (Xie et al. 2017), where the generation of noise is constrained, in that the new subject s' or object o' should have appeared in the dataset with the same relation r . This constraint focuses on generating harder and more confusing noise for any method. On the contrary, negative sampling on WN18 and Dementia KGs was performed randomly, without any constraint, to compare and contrast different methods and datasets on different noise types.

It is also important to note that, all these errors which were imbued to the 3 datasets are labeled as positives for training purposes. This means that the evaluation of the methods will be based on how effective they are in finding these hidden errors in every KG.

5.3 Evaluation Protocol

Following the same steps as (Socher et al. 2013), we compute the energy function $E(s, r, o) = ||s + r - o||$ for each triple in the dataset. Then, we generate a ranking

² <https://www.nlm.nih.gov/mesh/>

Table 2 Number of imputed errors based on ratio for each dataset

Dataset	N1 (10%)	N2 (20%)	N3 (40%)
WN18	14,144	28,288	56,445
FB15k	46,408	93,782	187,925
Dementia	13,500	27,000	54,000

for all triples based on this energy function score. The smaller the value of the energy value of the triple, the more valid the triple is. As such, we would hope that the erroneous triples would have much greater value than the initial correct ones. To measure this we use the filtered mean rank (fMR) and the filtered mean reciprocal rank ($fMRR$) (Melo and Paulheim 2017):

$$fMR = \frac{1}{N} \sum_{i=1}^N rank_i - i + 1 \quad (6)$$

$$fMRR = \frac{1}{N} \sum_{i=1}^N \frac{1}{rank_i - i + 1} \quad (7)$$

Additionally, after normalizing the energy function score in the [0-1] interval, we also use the Area Under the ROC Curve (AUC) to further examine how well algorithms classify the noise as an error. Values close to 0 indicate a correct triple, while values close to 1 indicate an erroneous triple. For fMR , lower is better while for $fMRR$ and AUC, higher is better.

5.4 Parameters and Settings

For all methods, we used the settings and parameters suggested by the corresponding authors on the two benchmark datasets (i.e. WN18 and FB15k). Concerning PaTy-BRED, as the authors underline, maximum path length (the maximum number of hops needed to go from a subject to and object) is set to 2. The maximum number of paths per length is set to 1000. As far as the heuristic measures for best path extraction are concerned, the authors employ the $m2$ heuristic, as it has the best performance between all other heuristic measures of relevance. Due to better results overall, random forests are the preferred choice of classifiers on the FB15k and WN18 datasets by the authors, with $k = 10$ best paths to select from a χ^2 procedure. For consistency, we also used these parameters on the Dementia dataset.

In all embeddings methods (including our own implementation) we used $d = 50$ as the dimension of the embeddings. Melo and Paulheim (Melo and Paulheim 2017) point out that on FB15k, $d = 50$ is the best value, adding that, at least for error detection, dimensionality should generally be low. In CKRL the authors also state that $d = 50$ is the best value for the FB15k dataset (although on a different task). The margin γ was set to 1.0, as CKRL and TransE use this specific value, and the learning rate was tested with the values $\{0.001, 0.01, 0.05\}$. For both datasets, training was limited to 1000 epochs, as further training didn't improve performance substantially.

Early stopping was used to determine the best model during these epochs. It is also stated (Bordes et al. 2013) that the L_1 norm works best on the loss function on both the WN18 and FB15k datasets, hence we also use the L_1 norm for each embedding method and dataset. Regarding the scaling value λ of the PRGE method, we use $\lambda = 5$, which yielded the best results on all datasets, after searching over a small subset of possible values.

5.5 Results and Discussion

Error Detection Experiments

Tables 3, 4 and 5 demonstrate the results of all approaches for the error detection task on all datasets. Some interesting observations and insights stemming from the results are presented here:

Table 3 Error Detection results for WN18 (Random Errors)

Dataset	WN18-N1			WN18-N2			WN18-N3		
	fMR	$fMRR$	AUC	fMR	$fMRR$	AUC	fMR	$fMRR$	AUC
PaTyBRED	4593	0.0008	0.9673	4694	0.0009	0.9668	4703	0.0007	0.9668
TransE	38942	0.0002	0.7247	39339	0.0003	0.7219	44464	0.0005	0.6857
PTransE	45721	0.0007	0.6768	45392	0.0003	0.6791	46412	0.0002	0.6719
CKRL	15738	0.0009	0.8887	16969	0.0007	0.8800	39253	0.0011	0.7225
PRGE	9913	0.0006	0.9299	12450	0.0004	0.9120	19956	0.0004	0.8589
PRGE-Scaled	3681	0.0009	0.9740	3870	0.0009	0.9727	3673	0.0008	0.9740

1) WN18 Dataset: Regarding the WN18 dataset, it is evident from Table 3 that our proposed PRGE-Scaled method outperforms all other methods. CKRL and PaTyBRED perform similarly on the $fMRR$ metric but are outperformed on the other evaluation metrics. It is also evident that, compared to the other methods, PRGE-Scaled performs better as the noise ratio goes up, with fMR and AUC score values being non-decreasing from N1 to N3 datasets.

2) FB15k Dataset: Here, PaTyBRED performs better than almost any base embedding method in error detection, indicating that potentially big factors here are the dataset size (see Table 1) and the different error imputation method. However, our PRGE-Scaled method fares better on the $fMRR$ metric, indicating that it separates better the obvious erroneous triples from the others. In addition, our PRGE-Scaled method fares better than all other embedding based methods.

3) Dementia Dataset: Firstly, as seen from Table 1 and Section 5.1, the knowledge graph is very sparse given the number of entities and relations available. Moreover, this dataset had noise present even before the noise imputation process, due to the automatic extraction process during its creation. As such, the actual noise level is much higher than the other datasets. Hence, given the distorted connectivity and much higher actual noise level, it is expected that the Dementia dataset will pose a more challenging error detection task. This can be seen in Table 5 where we can see

that the error detection is very hard for all methods, independently of approach and methodology. In spite of PaTyBRED being slightly better at the ranking metrics, our PRGE-Scaled method achieves better AUC scores, indicating that on average it can perform better than other models when comparing between an actual and a noisy triple. It can also scale better with the increasing noise ratio, something also seen in the WN18 dataset. In the N3 dataset, our method can achieve better fMR score than every method, indicating that in the presence of much noise, something that is almost a given in most automatically generated KGs, it can fair better than state-of-the-art methods.

4) Effect of noise: As expected, when the noise level rises from N1 to N3, the performance of all models deteriorates regardless of the dataset as seen in all Tables. However, our model is the most robust, especially when compared to the other embedding methods, showing much smaller fluctuations in performance.

Table 4 Error Detection results for FB15K (Same Relation Errors)

Dataset	FB15K-N1			FB15K-N2			FB15K-N3		
	fMR	$fMRR$	AUC	fMR	$fMRR$	AUC	fMR	$fMRR$	AUC
PaTyBRED	41785	0.0005	0.9064	46046	0.0003	0.8907	53320	0.0002	0.8694
TransE	127940	0.0002	0.7352	133763	0.0001	0.7231	169488	0.0000	0.6492
PTransE	166349	0.0000	0.6557	167997	0.0000	0.6523	173643	0.0000	0.6406
CKRL	96113	0.0001	0.8011	101583	0.0001	0.7897	112325	0.0001	0.7675
PRGE	89058	0.0004	0.8157	103167	0.0002	0.7865	106907	0.0001	0.7787
PRGE-Scaled	73994	0.0006	0.8469	89164	0.0005	0.8155	86347	0.0002	0.8213

Table 5 Error Detection results for Dementia (Random Errors)

Dataset	Dementia-N1			Dementia-N2			Dementia-N3		
	fMR	$fMRR$	AUC	fMR	$fMRR$	AUC	fMR	$fMRR$	AUC
PaTyBRED	56485	0.0006	0.5674	55749	0.0007	0.5604	59817	0.0003	0.5552
TransE	58014	0.0001	0.5702	59421	0.0001	0.5599	59835	0.0000	0.5568
PTransE	59718	0.0002	0.5576	61518	0.0001	0.5443	65533	0.0000	0.5146
CKRL	60584	0.0001	0.5512	61034	0.0001	0.5479	61089	0.0001	0.5475
PRGE	58049	0.0001	0.5700	58510	0.0001	0.5666	59844	0.0002	0.5567
PRGE-Scaled	57642	0.0001	0.5730	58258	0.0001	0.5685	59314	0.0001	0.5606

5) PRGE scaling effect: Regarding our proposed method, we can see that the λ scaled PRGE method works better than the unscaled method. This is true for all different noise imputation ratios and datasets, reflecting the importance of tuning the confidence score for each triple during training. Scaling the CKRL confidence score to achieve better performance is a possible future experiment and work, although the fact that the unscaled PRGE method performs better than CKRL makes it unlikely for a scaled CKRL version to perform better than a scaled PRGE version.

An extension to the scaling methodology here could be the additional scaling of the pairwise max-margin score function part of the loss function $[\gamma + E^+ - E^-]_+$. This part could be scaled with an additional θ parameter, leveraging more holistically

the role that both the confidence score and the pairwise max-margin loss function have when training the embeddings.

6) PTransE performance: Another interesting result is that PTransE performs worse than TransE on error detection on all datasets. This is also mentioned in (Xie et al. 2017) and is reconfirmed here by using the same energy function ($\|s + r - o\|$) for all embeddings methods. This is rather unexpected, as PTransE uses additional path information to guide embeddings similar to CKRL, which in turn performs better in all aspects than both TransE and PTransE.

At this point it is also important to stress two main advantages of the proposed methodology, alongside the superior performance results:

- **Modularity:** The proposed PRGE method is agnostic of the underlying energy function and triple-scoring mechanism. Since the path ranking score only acts in a multiplicative manner to the energy function, one can deduce that other embedding energy functions can be used to improve performance. Specifically, methods like TransH and TransR or even Semantic Matching methods like RESCAL can also be used in conjunction with the path ranking score and could improve results. Respectively, instead of the PaTyBRED score, other scoring mechanisms could be employed as well. This makes the proposed PRGE method generic and flexible, allowing for different combinations of techniques for the energy function and the confidence score that could enhance results for the task at hand.
- **Robust Embeddings:** Contrary to the PRA methods where only a confidence score for each triple is provided in the end, the PRGE method produces embeddings trained and guided by this confidence score. This results in embeddings robust to the inherent noise of the knowledge graph, which can be further used in other embedding-based downstream tasks, such as link prediction, triple classification, clustering, etc.

Triple Classification Experiments

To prove the usefulness of noise-robust embeddings in downstream tasks, we also performed a triple classification experiment following (Socher et al. 2013). Triple classification as a task revolves around predicting whether a triple belongs to a graph or not. Ultimately, our goal is to predict correct facts in the form of relations (s, r, o) from the test data, using the score $\|s + r - o\|$ by utilizing the embeddings generated from the various models.

More specifically, to classify whether a triple is valid or not, a threshold τ_r for every relation r is introduced. This threshold is chosen based on performance on the validation set. The threshold value for every relation is the one that maximizes the classification accuracy on the validation set. Then, using these thresholds, the performance of the model is estimated on the test set.

Both the validation and the test set are imputed with noise. Again, we create multiple instances of different noise levels to evaluate the effect of noise in performance. The imputation method is the same as described in Section 5.2.

Results on the FB15k can be seen in Table 6. We can see that as the noise ratio gets larger, PRGE methods perform better than the other methods. In addition, the

Table 6 Triple Classification on the FB15k dataset

Method	FB15K-N1	FB15K-N2	FB15K-N3
TransE	0.717	0.703	0.671
PTransE	0.686	0.678	0.67
CKRL	0.639	0.709	0.691
PRGE	0.712	0.715	0.681
PRGE-Scaled	0.715	0.712	0.702

PRGE-Scaled method consistently outperforms CKRL and PTransE on all noise ratios, indicating that using the path ranking score to train embeddings yields better results. The same was observed on the Dementia dataset, where PRGE-Scaled outperformed CKRL, PTransE and PRGE on all noise levels. PRGE-Scaled did not fare well on the WN18 dataset though, with CKRL and PTransE being better across all noise levels, though the unscaled version of PRGE outperformed any other method on the WN18 dataset. This indicates that the scaling parameter λ plays a big role in producing robust embeddings across all tasks and its choice depends on the knowledge graph structure and form. Conclusively, we can see that utilizing the PRGE framework to incorporate an error estimation score during the training process of the embeddings, actually helps in other downstream tasks with the generation of noise-robust embeddings.

Qualitative Results on the Dementia Dataset

Since our final goal is to detect the erroneous triples already present in a Knowledge Graph, we also performed a qualitative analysis of the predictions given by the model. We are interested in seeing if the lowest scoring triples, the ones that our model deemed most probable to be errors, should be removed. We focused on the Dementia dataset to showcase what the PRGE model predicts as erroneous in a real-world application. The PRGE model was trained on the initial Dementia dataset, before the imputation process. Thus, we are trying to detect the actual noise present in the Knowledge Graph. To assess the validity of the predictions we devised the following annotation task:

Firstly, we fetched the top-100 lowest-scoring triples, as predicted by our model. We also fetched the exact textual snippets from the publications that these triples were found in. Three human experts on natural language processing and bioinformatics were presented with these triples alongside their corresponding text. The triples were presented in the form (Entity1 Relation Entity2). They were asked to assess the quality of the triple given the corresponding textual content and how useful was the piece of information that was extracted. Specifically, the annotators could select one of the following labels for each triple:

- *Correct*, meaning the triple could be extracted from the text snippet, was sound and useful
- *Unsure*, meaning the annotator was not sure about the validity of the fact given the text snippet
- *Extraction Error (ER)*, meaning the triple was wrongly extracted from the text snippet




Fig. 2 Annotators' decisions on the 100 most erroneous triples for the Dementia Dataset, according to the PRGE-scaled model

- *Too General (TG)*, meaning the triple was actually found in the text but was deemed too general to be useful in the context of our KG (e.g. "*Taste Buds - PART OF - Homo sapiens*")
- *Too Specific (TS)*, meaning the triple was actually found in the text but was deemed too specific to be useful in the context of our KG (e.g. "*Brain - LOCATION OF - Decreased plasmalogens*")

The last three labels *Extraction Error*, *Too Specific* and *Too General* are all errors according to this evaluation scheme. We devised multiple labels for the errors because it is important to have a qualitative analysis of the errors made. For example, high *Extraction Errors* would indicate errors made from the relation extraction tool working directly on the text and would support research towards enhancing that part of the pipeline to reduce the error propagation. On the other hand, higher *Too General/Too Specific* errors would insinuate that the relation extraction tool works correctly, however, the extracted triples are not important for the task at hand. In that case, we could devise and apply a post-processing step to keep meaningful triples.

The results of the annotators evaluation can be seen in Figure 2. We can see that for all annotators more than 85 % of these triples seemed to be erroneous given the context. On the other hand, out of these lowest-scoring triples only $\approx 8\%$ is actually correct, across all annotators. This indicates the high precision of the prediction and allows us to be fairly confident on the scoring of the model.

As an added example, some of the manually assessed triples can be seen in Table 7. There, we present the two lowest-scoring triples from each one of the erroneous categories (i.e. *ER*, *TG*, *TS*) and the corresponding text they were derived from. The triples presented were all unanimously labeled as the specific type of error from the annotators. As we can see, the triple *Denmark - PART_OF - Neurons* is an exemplary case where the extraction tool failed to identify the correct entities and their relation.

On the other hand, the *TS* and *TG* errors showcased, such as *Entire thumb - PART_OF - Patients*, may be correctly extracted from the text but they provide little to no added value to the aggregated knowledge. Having multiple such facts deteriorates performance both qualitatively, as each model/algorithm has to encapsulate a lot of “useless” information and technically, as the volume of the data greatly expands. Thus, these examples showcase the importance of the discrimination regarding the type of error made, as well as, the added value of performing such an analysis in noisy graphs.

Table 7 The lowest-scoring triples of the erroneous categories, two from each one, along with the type of error and the initial text they were extracted from.

Erroneous triple (s - r - o)	Type	Supporting text
acetonitrile - AUGMENTS - 80%	ER	... The acetonitrile concentration was increased to 80% in 5 min and then held in 80% acetonitrile for an additional 5 min ...
Denmark - PART_OF - Neurons	ER	... Neurons transfected with DA-GFP were found to have dendritic spines that had significantly lengthened necks compared to ...
Cells - PART_OF - Medial geniculate body	TG	... Dendritic spines of the polyhedral and elongated cells of the medial geniculate bodies were decreased in number ...
Entire thumb - PART_OF - Patients	TG	... The patient is asked to hold the ruler with his thumb and forefinger and to release the ruler while the investigator continues to ...
DDMS - PART_OF - Homo sapiens	TS	... HT-22 hippocampal cells and confirms observations using brain extracts from monkey, mouse, rat and human DDM ...
Brain - LOCATION_OF - Decreased plasmalogens	TS	... data suggest that long-term alterations in plasmalogen synthesis degradation result in decreased brain plasmalogen levels, a hallmark feature of AD ...

6 Conclusion and Future Work

In this paper, we evaluated and compared different methods for error detection in knowledge graphs. The methods we compared use path ranking elements, embedding structures or a combination of both. We also proposed a general framework for combining the path ranking score of a triple with the graph embedding framework, resulting in embeddings robust towards noise present in the graph. We assessed both the quantitative and qualitative performance of the framework. Utilizing the score from the best-performing path ranking algorithm (PaTyBRED) to train the embeddings, we have managed to overcome other state-of-the-art hybrid methods (PTransE, CKRL) on all datasets and enhance the classification results of PaTyBRED on two of the three datasets. Moreover, by combining these two frameworks we extended the PRA methods with the ability to generate embeddings for the entities, that could be used further. Thus, we have proposed a generic framework to generate embeddings resilient to noise and we proved that they can also be used in multiple downstream tasks

enhancing performance in the presence of noise³. Finally, we performed a qualitative evaluation of the possible errors detected in a real-world dataset, to showcase the importance of such approaches in actual applications.

We have showcased the efficiency of this approach focusing on the confidence score of PaTyBRED and the TransE energy function. However, a plethora of other combinations could be also explored. In the future, the following directions will be researched: Firstly, energy functions from other embedding methods, such as TransH (Wang et al. 2014) and TransR (Lin et al. 2015b) will be used. Secondly, as already mentioned in the error detection results section, we will try to further expand the PRGE-Scaled method by including an additional parameter θ to the pairwise max-margin part of the loss function. Lastly, inspired by rule-guided embedding frameworks such as KALE (Guo et al. 2016) and RUGE (Guo et al. 2017), we plan to combine path ranking scores with rule-based scores or rule-based training to enhance the results on the error detection task and propose a more unifying framework.

Acknowledgement

This paper is supported by European Union’s Horizon 2020 research and innovation programme under grant agreement No. 727658, project IASIS (Integration and analysis of heterogeneous big data for precision medicine and suggested treatments for different types of patients).

References

- Auer S, Bizer C, Kobilarov G, Lehmann J, Cyganiak R, Ives Z (2007) Dbpedia: A nucleus for a web of open data. In: Aberer K, Choi KS, Noy N, Allemang D, Lee KI, Nixon L, Golbeck J, Mika P, Maynard D, Mizoguchi R, Schreiber G, Cudré-Mauroux P (eds) *The Semantic Web*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 722–735
- Bodenreider O (2004) The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research* 32(suppl.1):D267–D270
- Bollacker K, Evans C, Paritosh P, Sturge T, Taylor J (2008) Freebase: A collaboratively created graph database for structuring human knowledge. In: *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, ACM, New York, NY, USA, SIGMOD ’08, pp 1247–1250, DOI 10.1145/1376616.1376746, URL <http://doi.acm.org/10.1145/1376616.1376746>
- Bordes A, Usunier N, Garcia-Duran A, Weston J, Yakhnenko O (2013) Translating embeddings for modeling multi-relational data. In: Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ (eds) *Advances in Neural Information Processing Systems 26*, Curran Associates, Inc., pp 2787–2795, URL <http://papers.nips.cc/paper/5071-translating-embeddings-for-modeling-multi-relational-data.pdf>
- Cai H, Zheng VW, Chang KC (2017) A comprehensive survey of graph embedding: Problems, techniques and applications. CoRR abs/1709.07604, URL <http://arxiv.org/abs/1709.07604>, 1709.07604
- Carlson A, Betteridge J, Wang RC, Hruschka ER Jr, Mitchell TM (2010) Coupled semi-supervised learning for information extraction. In: *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, ACM, New York, NY, USA, WSDM ’10, pp 101–110, DOI 10.1145/1718487.1718501, URL <http://doi.acm.org/10.1145/1718487.1718501>

³ Code and material can be found here: <https://github.com/RomFas/PRGE/>

- Fan M, Zhou Q, Chang E, Zheng TF (2014) Transition-based knowledge graph embedding with relational mapping properties. In: Proceedings of the 28th Pacific Asia Conference on Language, Information and Computing, Department of Linguistics, Chulalongkorn University, Phuket, Thailand, pp 328–337, URL <https://www.aclweb.org/anthology/Y14-1039>
- Fellbaum C (2005) Wordnet and wordnets. In: Barber A (ed) Encyclopedia of Language and Linguistics, Elsevier, pp 2–665
- Feng J, Huang M, Wang M, Zhou M, Hao Y, Zhu X (2016) Knowledge graph embedding by flexible translation. URL <https://www.aaai.org/ocs/index.php/KR/KR16/paper/view/12887>
- Galárraga L, Teflioudi C, Hose K, Suchanek FM (2015) Fast rule mining in ontological knowledge bases with amie+. The VLDB Journal 24(6):707–730, DOI 10.1007/s00778-015-0394-1, URL <https://doi.org/10.1007/s00778-015-0394-1>
- Gardner M, Mitchell T (2015) Efficient and expressive knowledge base completion using subgraph feature extraction. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Lisbon, Portugal, pp 1488–1498, DOI 10.18653/v1/D15-1173, URL <https://www.aclweb.org/anthology/D15-1173>
- Guo S, Wang Q, Wang L, Wang B, Guo L (2016) Jointly embedding knowledge graphs and logical rules. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Austin, Texas, pp 192–202, DOI 10.18653/v1/D16-1019, URL <https://www.aclweb.org/anthology/D16-1019>
- Guo S, Wang Q, Wang L, Wang B, Guo L (2017) Knowledge graph embedding with iterative guidance from soft rules. CoRR abs/1711.11231, URL <http://arxiv.org/abs/1711.11231>, 1711.11231
- Heindorf S, Potthast M, Stein B, Engels G (2016) Vandalism detection in wikidata. In: Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, ACM, pp 327–336
- Jia S, Xiang Y, Chen X, E S (2018) TTMF: A triple trustworthiness measurement frame for knowledge graphs. CoRR abs/1809.09414, URL <http://arxiv.org/abs/1809.09414>, 1809.09414
- Krithara A, Aisopos F, Rentoumi V, Nentidis A, Bougatiotis K, Vidal ME, Menasalvas E, Rodriguez-Gonzalez A, Samaras E, Garrard P, et al. (2019) iasis: Towards heterogeneous big data analysis for personalized medicine. In: 2019 IEEE 32nd International Symposium on Computer-Based Medical Systems (CBMS), IEEE, pp 106–111
- Lao N, Cohen WW (2010) Relational retrieval using a combination of path-constrained random walks. Machine Learning 81(1):53–67, DOI 10.1007/s10994-010-5205-8, URL <https://doi.org/10.1007/s10994-010-5205-8>
- Lehmann J, Gerber D, Morsey M, Ngonga AC Ngomo (2012) Defacto - deep fact validation. In: The Semantic Web – ISWC 2012, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 312 – 327
- Lin Y, Liu Z, Luan H, Sun M, Rao S, Liu S (2015a) Modeling relation paths for representation learning of knowledge bases. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Lisbon, Portugal, pp 705–714, DOI 10.18653/v1/D15-1082, URL <https://www.aclweb.org/anthology/D15-1082>
- Lin Y, Liu Z, Sun M, Liu Y, Zhu X (2015b) Learning entity and relation embeddings for knowledge graph completion. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9571>
- McCray AT (2003) An upper-level ontology for the biomedical domain. Comparative and Functional Genomics 4(1):80–84
- Melo A, Paulheim H (2017) Detection of relation assertion errors in knowledge graphs. In: Proceedings of the Knowledge Capture Conference, ACM, New York, NY, USA, K-CAP 2017, pp 22:1–22:8, DOI 10.1145/3148011.3148033, URL <http://doi.acm.org/10.1145/3148011.3148033>
- Nentidis A, Bougatiotis K, Krithara A, Paliouras G (2019) Semantic integration of disease-specific knowledge 1912.08633
- Nickel M, Tresp V, Kriegel HP (2011) A three-way model for collective learning on multi-relational data. In: Proceedings of the 28th International Conference on International Conference on Machine Learning, Omnipress, USA, ICML’11, pp 809–816, URL <http://dl.acm.org/citation.cfm?id=3104482.3104584>
- Nickel M, Rosasco L, Poggio TA (2015) Holographic embeddings of knowledge graphs. CoRR abs/1510.04935, URL <http://arxiv.org/abs/1510.04935>, 1510.04935
- Paulheim H (2016) Knowledge graph refinement: A survey of approaches and methods. Semantic Web 8:489–508, DOI 10.3233/SW-160218

- Paulheim H, Bizer C (2014) Improving the quality of linked data using statistical distributions. *International Journal on Semantic Web and Information Systems (IJSWIS)* 10(2):63 – 86
- Rindflesch TC, Fiszman M (2003) The interaction of domain knowledge and linguistic structure in natural language processing: interpreting hypernymic propositions in biomedical text. *Journal of biomedical informatics* 36(6):462–477
- Shi B, Weninger T (2016) Proje: Embedding projection for knowledge graph completion. CoRR abs/1611.05425, URL <http://arxiv.org/abs/1611.05425>, 1611.05425
- Socher R, Chen D, Manning CD, Ng A (2013) Reasoning with neural tensor networks for knowledge base completion. In: Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ (eds) *Advances in Neural Information Processing Systems 26*, Curran Associates, Inc., pp 926–934, URL <http://papers.nips.cc/paper/5028-reasoning-with-neural-tensor-networks-for-knowledge-base-completion.pdf>
- Suchanek FM, Kasneci G, Weikum G (2007) Yago: A core of semantic knowledge. In: *Proceedings of the 16th International Conference on World Wide Web*, ACM, New York, NY, USA, WWW '07, pp 697–706, DOI 10.1145/1242572.1242667, URL <http://doi.acm.org/10.1145/1242572.1242667>
- Tanon TP, Vrandeic D, Schaffert S, Steiner T, Pintscher L (2016) From freebase to wikidata: The great migration. In: *World Wide Web Conference*
- Trouillon T, Welbl J, Riedel S, Gaussier É, Bouchard G (2016) Complex embeddings for simple link prediction. CoRR abs/1606.06357, URL <http://arxiv.org/abs/1606.06357>, 1606.06357
- Wang M, Rong E, Zhuo H, Zhu H (2018) Embedding knowledge graphs based on transitivity and asymmetry of rules. In: Phung D, Tseng VS, Webb GI, Ho B, Ganji M, Rashidi L (eds) *Advances in Knowledge Discovery and Data Mining*, Springer International Publishing, Cham, pp 141–153
- Wang P, Dou D, Wu F, de Silva N, Jin L (2019) Logic rules powered knowledge graph embedding. CoRR abs/1903.03772, URL <http://arxiv.org/abs/1903.03772>, 1903.03772
- Wang Q, Mao Z, Wang B, Guo L (2017) Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering* PP:1–1, DOI 10.1109/TKDE.2017.2754499
- Wang Z, Zhang J, Feng J, Chen Z (2014) Knowledge graph embedding by translating on hyperplanes URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/view/8531>
- Xie R, Liu Z, Sun M (2017) Does william shakespeare REALLY write hamlet? knowledge representation learning with confidence. CoRR abs/1705.03202, URL <http://arxiv.org/abs/1705.03202>, 1705.03202
- Yang B, Yih Wt, He X, Gao J, Deng l (2014) Learning multi-relational semantics using neural-embedding models