

Contrastive Knowledge Graph Error Detection

Qinggang Zhang
The Hong Kong Polytechnic
University
Hung Hom, Hong Kong SAR
qinggangg.zhang@connect.polyu.hk

Junnan Dong
The Hong Kong Polytechnic
University
Hung Hom, Hong Kong SAR
hanson.dong@connect.polyu.hk

Keyu Duan
The Hong Kong Polytechnic
University
Hung Hom, Hong Kong SAR
k.duane9902@gmail.com

Xiao Huang
The Hong Kong Polytechnic
University
Hung Hom, Hong Kong SAR
xiaohuang@comp.polyu.edu.hk

Yezi Liu
University of California Irvine
Irvine, CA, USA
yezil3@uci.edu

Linchuan Xu
The Hong Kong Polytechnic
University
Hung Hom, Hong Kong SAR
linch.xu@polyu.edu.hk

ABSTRACT

Knowledge Graph (KG) errors introduce non-negligible noise, severely affecting KG-related downstream tasks. Detecting errors in KGs is challenging since the patterns of errors are unknown and diverse, while ground-truth labels are rare or even unavailable. A traditional solution is to construct logical rules to verify triples, but it is not generalizable since different KGs have distinct rules with domain knowledge involved. Recent studies focus on designing tailored detectors or ranking triples based on KG embedding loss. However, they all rely on negative samples for training, which are generated by randomly replacing the head or tail entity of existing triples. Such a negative sampling strategy is not enough for prototyping practical KG errors, e.g., (*Bruce Lee*, *place_of_birth*, *China*), in which the three elements are often relevant, although mismatched. We desire a more effective unsupervised learning mechanism tailored for KG error detection. To this end, we propose a novel framework - ContrActive knowledge Graph Error Detection (CAGED). It introduces contrastive learning into KG learning and provides a novel way of modeling KG. Instead of following the traditional setting, i.e., considering entities as nodes and relations as semantic edges, CAGED augments a KG into different hyper-views, by regarding each relational triple as a node. After joint training with KG embedding and contrastive learning loss, CAGED assesses the trustworthiness of each triple based on two learning signals, i.e., the consistency of triple representations across multi-views and the self-consistency within the triple. Extensive experiments on three real-world KGs show that CAGED outperforms state-of-the-art methods in KG error detection. Our codes and datasets are available at <https://github.com/Qing145/CAGED.git>.

CCS CONCEPTS

• **Computing methodologies** → **Anomaly detection.**

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '22, October 17–21, 2022, Atlanta, GA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9236-5/22/10...\$15.00

<https://doi.org/10.1145/3511808.3557264>

KEYWORDS

Knowledge graph, anomaly detection, contrastive learning

ACM Reference Format:

Qinggang Zhang, Junnan Dong, Keyu Duan, Xiao Huang, Yezi Liu, and Linchuan Xu. 2022. Contrastive Knowledge Graph Error Detection. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management (CIKM '22)*, October 17–21, 2022, Atlanta, GA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3511808.3557264>

1 INTRODUCTION

Knowledge graph (KG) error detection becomes increasingly essential as more KG-based systems have been being deployed, such as conversational agents, and recommender systems [1, 2]. KGs reformulate a realistic assertion as a triple, i.e., (*head entity*, *relation*, *tail entity*), and could effectively organize information and knowledge in a structured and scalable way. Most KGs are extracted from web corpora by using heuristic algorithms. Amounts of noisy triples were inevitably introduced into KGs due to the noises in the original sources and the imperfect extraction algorithms [3–6]. For example, a widely-used KG, NELL [7] has a precision of 74%, corresponding to around 0.6 million inaccurate triples. Existing KG driven studies ignore the significant impact of KG errors. They assume that all triples in KGs are correct and therefore leads to severe performance degradation in downstream tasks. To mitigate the roadblocks, there is an urgent need for the development of effective KG error detection algorithms.

One of the key challenges in KG error detection is that the patterns of errors are unknown and diverse, while ground-truth labels are rare or even unavailable. In practice, trivial erroneous triples were often corrected during the construction of KGs [8]. Therefore, labeled KG errors, i.e., the ground truth, are often not available, and remaining errors in KGs are nontrivial. Briefly, the taxonomy of previous works encompasses two branches: (i) *Rule-based* and (ii) *Embedding-based*. Namely, the former one conducts KG error detection based on rules [9, 10]. They define KG errors as assertions that violate any one of the pre-defined rules [11–13]. But these methods are not generalizable since different KGs have distinct rules with domain knowledge involved. More recently, several embedding-based KG error detection methods have been explored [14–18]. To perform unsupervised error detection, they employ a naive negative sampling to train models [15]. Such negative samples have also

Table 1: Major notations and definitions.

Notations	Descriptions
\mathcal{G}	a knowledge graph with errors
(h, r, t)	a triple with head entity h , relation r , tail entity t of all entities and all relation types
(e_h, e_r, e_t)	embedding representations of h , r , and t
\mathcal{T}	view I of \mathcal{G} , i.e., a graph with triples as nodes
\mathcal{T}'	view II of \mathcal{G} , i.e., a graph with triples as nodes
\mathbf{x}_i	representation of i^{th} triple learned from \mathcal{T}
\mathbf{z}_i	representation of i^{th} triple learned from \mathcal{T}'

been employed as synthetic labels to train detectors in supervised KG error detection methods [16, 18].

However, negative samples employed by existing models bring an inherent limitation. Given a positive triple (h, r, t) , the corresponding negative samples are generated by randomly replacing the entities, i.e. h or t . Such a negative sampling strategy is not enough for prototyping practical KG errors. For instance, given a true triple (*Newton, Nationality, England*), they may randomly generate (*Newton, Nationality, Google*). But real-world errors often have more complex patterns, e.g., (*Bruce_Lee, place_of_birth, China*), in which the three elements are mismatched but often relevant, since Lee was born in the Chinatown area of San Francisco. To this end, we desire a more appropriate way to enable effective detection.

As an alternative self-supervised learning technique [19–21], contrastive learning has proven effective in many network analysis tasks [22–26], e.g. classification [27], link prediction [28], and recommender systems [29]. Previous study [30, 31] have already extended the power of contrastive learning for error detection in computer vision and demonstrates its superiority under various error detection scenarios. Using data transformations, it first creates distinct views of the original data, and then learn the representations by maximizing agreement of each instance’s representation among different views. By contrasting and learning between the elaborate instance pairs, the model can capture informative features for distinguishing errors without manual labels.

In spite of the wide applications of contrastive learning, we argue that it is nontrivial to explore it for KG error detection. Basically, there are two major challenges. First, defining views on KGs is not a trivial task, given that KGs and errors have unique data structure. Most existing graph contrastive learning frameworks rely on graph-level augmentations such as node dropping, edge perturbation, subgraph sampling or matrix diffusion [24, 26] to create different views of graphs. However, these graph-level augmentation methods are not suitable for KG error detection. A pure KG is actually a set of triples, thus detecting errors on the KG is equal to identifying noisy triples. Existing graph contrastive learning frameworks, that focus on entity or graph-level contrasting, benefit little for learning rich and distinguishing representations at triple level. Additionally, existing graph augmentation methods are designed for network embedding, which could boost the robustness of the learned representations but at the price of damaging the underlying distribution of potential errors. Taking node dropping and edge perturbation as typical examples, they would change the structure of potential noisy triples or even introduce new erroneous triples, which makes the detection task more challenging. Second, existing

graph encoders would not take anomalies or errors into consideration. As a result, potential erroneous triples may also contribute significantly during the message-passing process, which lead to suboptimal performance in error detection. Thus, a robust error-aware graph encoder is desired, aiming to alleviate the impact of errors and perform effective learning.

To this end, we formally define the problem of KG error detection, and investigate two research questions. ❶ How to construct appropriate different views of a KG to conduct effective contrastive learning? ❷ How to perform error-aware KG encoding that is tailored for error detection? To answer these questions, we propose an effective framework named ContrActive knowledge Graph Error Detection (CAGED). Our major contributions are listed as follows.

- We develop a framework - CAGED, which could perform effective error detection by integrating contrastive learning and KG embedding.
- We design a new KG augmentation mechanism to enable appropriate contrastive learning in CAGED. It generates different views of a KG at the triple level and avoids undermining the underlying structure of potential errors.
- We present an error-aware graph encoder - EaGNN, to model triple representations for each view. It adopts a tailored gated attention mechanism to block the information propagation from erroneous triples.
- We conduct empirical studies on three real-world KGs. Experimental results show that CAGED outperforms state-of-the-art error detection algorithms.

2 PROBLEM STATEMENT

Notations: We use an uppercase bold alphabet (e.g., \mathbf{W}) to denote a matrix and a lowercase bold alphabet (e.g., \mathbf{x}) to represent a vector. The transpose of a matrix is denoted as \mathbf{W}^T . We use $\|\mathbf{x}\|_2$ to represent the ℓ^2 norm of a vector. The operation $\mathbf{x} = [\mathbf{h}; \mathbf{r}; \mathbf{t}]$ denotes concatenating column vectors \mathbf{h} , \mathbf{r} , and \mathbf{t} into a new column vector \mathbf{x} . We list the major symbols in this paper in Table 1. Let \mathcal{G} denote a KG that contains a set of triples. Each triple is composed of a head entity h , a relation r , and a tail entity t , represented as (h, r, t) . We formally define KG errors at the triple level.

DEFINITION 1. Knowledge Graph Errors. Given a triple in a KG (h, r, t) , if there is a mismatch between head/tail entities and its relation r , then this triple (h, r, t) is an error. E.g., (*Bruce_Lee, place_of_birth, China*) and (*Bill_Gates, CEO, Google*) are errors.

In this paper, we employ a metric named confidence scores [32], ranging from 0 to 1, to indicate the degree of triples being correct. False triples should have confidence scores close to 0. We formally define the problem of KG error detection as follows.

Given a KG $\mathcal{G} = \{(h, r, t)\}$, we aim to find an optimal rating model $p_\omega : (h, r, t) \rightarrow \mathbb{R}^{(0,1)}$ with our proposed training framework such that $p_\omega(h, r, t)$ accurately denotes the probability of triple (h, r, t) being a defined error in Def 1. The performance of the framework is evaluated by *Precision* and *Recall* of the triples with the top lowest confidence scores.




Figure 1: Two separate augmentation operators are applied to the original KG, generating two triple graphs as congruent views, i.e., View I and View II. After training, we estimate the confidence score by measuring the consistency of triple representations across multi-views, i.e., x_A and z_A , and the self-consistency within the triple, i.e., $(\tilde{e}_h, \tilde{e}_r, \tilde{e}_t)$.

3 METHODOLOGY

In traditional KG representation learning, a KG is usually modeled as a heterogeneous graph by centering the entities as nodes and regarding relations as semantic edges. However, with such modeling, it is notoriously hard to capture the complex correlation across triples. In this work, we propose a novel framework named Contrastive Knowledge Graph Error Detection (CAGED) to effectively detect errors on large-scale KGs. Its core idea is to transform the target KG into different hyper-views, i.e. triple graphs, by regarding each relational triple as nodes, and then assess the trustworthiness of each triple based on the consistency between its representations learned from multiple views. The intuition behind this approach is that contrastive learning can learn a rich representation of key semantics in normal instances then the absence of such feature in anomalous triples can force their representations to lie far in the latent vector space, thus how well a triple’s representations from different views can be pulled close to each other serves as a good learning signal for error detection.

As illustrated in Figure 1, our proposed CAGED consists of three components, i.e., KG augmentation, error-aware encoder, and joint confidence estimation. First, we propose a novel KG augmentation method to generate two triple-level graphs by regarding each relational triple as nodes and treat them as two congruent views, i.e. View I and View II, of our contrastive framework. Second, we design a tailored error-aware knowledge graph neural network (EaGNN) as the encoder to block the information propagation from potential errors during the representation learning phase. Third, to learn rich and distinguishable across-view triple representations, we integrate translation-based KG embedding loss and contrastive learning loss to jointly optimize the model. After training, we estimate the confidence score of each triple based on two learning signals, i.e., the consistency of triple representations across multi-views and the self-consistency within the triple.

3.1 Knowledge Graph Augmentation

Data augmentation is the most crucial part in contrastive learning [22]. However, unlike image data in which views can be created

by standard augmentations [24, 25], e.g., rotating, cropping, distorting, etc., defining views on KGs is not a trivial task due to its unique data structure. Most existing graph contrastive learning frameworks rely on graph-level augmentations such as node dropping, edge perturbation, subgraph sampling or matrix diffusion [26] to create different views of graphs, and achieve considerable performance improvements in some classical graph related downstream tasks, such as node classification, graph classification, and link prediction.

However, these graph-level augmentation methods are not suitable for KG error detection. A pure KG is actually a set of triples, thus detecting errors on the KG is equal to identifying noisy triples. Existing graph contrastive learning framework, that focus on entity or graph-level contrasting, benefit little for learning rich and distinguishing representations at triple level. Additionally, existing graph augmentation methods are designed for network embedding, which could boost the robustness of the learned representations but at the price of damaging the underlying structure of potential errors. Specifically, existing augmentation methods such as node dropping, edge perturbation would change the distribution of potential errors or even introduce new errors, which makes the detection task more challenging.

DEFINITION 2. Linking Pattern. For any two triples sharing entities, i.e. $T_1 = (h_1, r_1, t_1) \cap T_2 = (h_2, r_2, t_2)$, we have two linking patterns: (i) sharing head entity ($h_1 = h_2 \oplus h_1 = t_2$), (ii) sharing tail entity ($t_1 = h_2 \oplus t_1 = t_2$). For our construction criterion, we build two triple graphs with these two linking patterns, based on the rationale that these two patterns possess different semantics.

In this paper, we propose a novel KG augmentation method that generates different views of KG at triple level and avoids undermining the distribution of potential errors. Specifically, we create two triple graphs, i.e. \mathcal{T} and \mathcal{T}' , by regarding each relational triple as nodes and connect them based on different linking patterns, as shown in Def 2. These two triple graphs can be regarded as congruent views of the target KG. Concretely, the triple graph \mathcal{T} models the correlation between relational triples from the perspective of head entity, while the triple graph \mathcal{T}' represent the distribution of relational triple from the perspective of tail entities. Considering

that the connected triples that share the same entity are always semantically relevant, for normal triples in KGs, we can easily find enough relevant neighboring triples in \mathcal{T} , i.e. View I, to reconstruct its semantics learned from \mathcal{T}' , i.e. View II. So, modeling and measuring the consistency between its representations learned from these two views can help us assess the trustworthiness of each triple in the original KG.

3.2 Error-aware Encoder - EaGNN

The existence of errors may significantly jeopardize the expressiveness of the learned representations and consequently degrade the performance of KG error detection, since erroneous triples will also propagate noisy information to its neighbors during the representation learning phase. Thus, a robust encoder is desired to block the information propagation from potential errors during the message passing process. In this paper, we design a tailored error-aware knowledge graph neural network (EaGNN) as an encoder to capture the rich semantics and structure information of triples in each view. Our EaGNN tightly integrates a local information modeling layer and a global error-aware attention layer that work jointly to get the reliable KG embeddings.

3.2.1 Local information modeling layer. Constructing triple graph from the original knowledge graph, in some way, may lose some local structure information, which means a structure of $h \rightarrow r \rightarrow t$. Since every instance in the triple graphs is transformed from a corresponding triple (h, r, t) in the original KG, we first randomly initialize the embedding of entities and relations in the original KG, and then adopt a local information modeling layer, i.e. a set of Bi-LSTM units, to learn the local relational structure within each triple. Taking the triple (h, r, t) as an example, the local information modeling layer is formulated as follows.

$$\tilde{e}_h, \tilde{e}_r, \tilde{e}_t = \text{Bi-LSTM}(e_h, e_r, e_t), \quad (1)$$

$$\mathbf{q}_i = [\tilde{e}_h; \tilde{e}_r; \tilde{e}_t]. \quad (2)$$

The output triple embedding \mathbf{q}_i could well capture the relational structure within the input triple. Thus, we used them as the initial triple embedding in triple graphs.

3.2.2 Global error-aware attention layer. When evaluating the trustworthiness of a triple, apart from the local structures within triples, the global contextual information from the neighbor triples would also provide crucial hints. To model such global contextual information, an advanced approach is to use knowledge graph neural networks (KGNNs). However, existing graph neural networks are not effective for KG error detection because they ignore the existence of anomalous instance in KGs when learning the representations. In other words, potential noisy instances may get the same level of attention as normal ones during the message-passing process, which leads to final triple representations containing noisy facts. To reduce the adverse effect from potential errors, we adopt a novel triple-level attention mechanism to learn global triple representations for each view.

Given an anchor triple $\mathbf{q}_i \in \mathbb{R}^d$ in View I, we update its embedding representation based on the weighted aggregation of its

neighbor triples, e.g. $\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_m\}$. The weight between the anchor triple i and its neighbor triple j is calculated as follows.

$$\hat{\alpha}_{ij} = \mathcal{A}(\mathbf{W}\mathbf{q}_i, \mathbf{W}\mathbf{q}_j). \quad (3)$$

$\hat{\alpha}_{ij}$ indicates the importance of triple j to triple i . $\mathbf{W} \in \mathbb{R}^{n \times d}$ is a learnable linear augmentation matrix to project the initial triple representations into the same vector space. \mathcal{A} is the attentional function: $\mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$. To make attention coefficients easily comparable across different triples, we normalize them by applying a softmax function.

$$\bar{\alpha}_{ij} = \frac{\exp(\hat{\alpha}_{ij})}{\sum_{k=1}^m \exp(\hat{\alpha}_{ik})}. \quad (4)$$

Since erroneous triples are relatively irrelevant with its neighbors, such an attention mechanism can adaptively assign smaller attention weights to potential abnormal triples. Based on this observation, we introduce a hyperparameter $\mu \in \mathbb{R}$ as threshold value to block information comes from potential anomalous triples. To validate our gated attention mechanism and the effect of μ , we conduct a detailed ablation study in Sec 4.3 and 4.4.

$$\alpha_{ij} = \begin{cases} \bar{\alpha}_{ij}, & \bar{\alpha}_{ij} > \mu, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

The final error-aware triple representation in View I can be calculated with a sigmoid function, as depicted in Eq.(6):

$$\mathbf{x}_i = \sigma\left(\sum_{j=1}^m \alpha_{ij} \mathbf{W}\mathbf{q}_j\right). \quad (6)$$

Similarly, given an anchor triple \mathbf{q}_i in View II, we also update its embedding by attending over its neighborhoods' features through:

$$\mathbf{z}_i = \sigma\left(\sum_{j=1}^m \alpha'_{ij} \mathbf{W}\mathbf{q}_j\right). \quad (7)$$

3.3 Joint Confidence Estimation

To make sure that our model can learn rich and distinguishable across-view representations, we integrate KG embedding loss and standard contrastive learning loss to jointly train the model.

3.3.1 KG embedding loss. In triple level, relations can be interpreted as translations operating on the low-dimensional embeddings of the entities. So, the more a triple fits the translation assumption, i.e. $\mathbf{h} + \mathbf{r} \simeq \mathbf{t}$, the more convincing this triple should be considered. Existing KG embedding algorithms have developed various energy functions to model the translational structure for better learning embeddings. In this paper, we take a simple squared euclidean distance to measure the unconformity of each triple with translation assumption and define the KG embedding loss as follows.

$$\mathcal{L}_{kge} = \sum_{(h,r,t) \in \mathcal{G}} \sum_{(\hat{h}, \hat{r}, \hat{t}) \in \hat{\mathcal{G}}} \max\left(0, \gamma + E(h, r, t) - E(\hat{h}, \hat{r}, \hat{t})\right), \quad (8)$$

where $E(h, r, t) = \|e_h + e_r - e_t\|_2$ is the traditional energy score for translational embedding models following translation assumption. $\gamma > 0$ is the hyperparameter of margin, and \mathcal{G} represents the sampled positive triple set. Since there are no explicit negative triples in KGs, we generate them by replacing the head h or tail t entity of a triple with a random one \hat{h} or \hat{t} .

$$\hat{\mathcal{G}} = \{(\hat{h}, r, t) \mid \hat{h} \in \mathcal{G}\} \cup \{(h, r, \hat{t}) \mid \hat{t} \in \mathcal{G}\}. \quad (9)$$

Algorithm 1: Contrastive KG Error Detection

Input: Knowledge graph \mathcal{G} with errors.
Output: KG embeddings and triple’s confidence score.

- 1 Initialize network parameters;
- 2 Construct two triple graphs \mathcal{T} and \mathcal{T}' based on augmentation rules, as shown in Sec 3.1;
- 3 **while not converged do**
- 4 **for** $(h, r, t) \in \mathcal{G}$ **do**
- 5 /* Local information modeling layer: */
- 5 Compute the local triple embedding representations, defined in Eqs.(1)-(2);
- 6 /* Global error-aware attention layer: */
- 6 Get global error-aware triple representations, defined in Eqs.(3)-(7);
- 7 /* Joint optimization: */
- 7 Jointly optimize the model via a combination of translation-based KG embedding loss in Eq.(8) and contrastive loss in Eq.(10);
- /* confidence estimation module: */
- 8 Calculate the confidence score of each triple (h, r, t) as defined in Eq. (11);

3.3.2 *Contrastive loss.* The negative sampling used in KG embedding loss can guide the model to learn rich structure and semantics information within triples. But such kind of local feature alone is not enough to enable effective error detection on the whole KG. To fill the gap, we jointly optimize it along with the contrastive learning loss. By learning and contrasting between the two triple graphs, our model can learn a rich representation of key semantics in normal instances then the absence of such feature in anomalous triples can force them to lie far in the latent vector space. Therefore, how well a triple’s representations from different views can be pulled close to each other serves as a good learning signal for error detection.

To integrate contrastive learning into KG embedding, in this paper, we employ the normalized temperature-scaled cross entropy loss as our contrastive learning objective to jointly train the model. Specifically, we randomly sample a minibatch of n examples from the original KG and define the training task on pairs of augmented examples derived from the minibatch, resulting in $2n$ triple representations, i.e. $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ from View I and $\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n\}$ from View II. Assuming that $(\mathbf{x}_i, \mathbf{z}_i)$ are the representations of triple i learned from different views. The loss function is defined as follows:

$$\mathcal{L}_{con}(\mathbf{x}_i, \mathbf{z}_i) = -\log \frac{\exp(\text{sim}(\mathbf{x}_i, \mathbf{z}_i) / \tau)}{\sum_{j \in \{1, 2, \dots, n\} \setminus \{i\}} \exp(\text{sim}(\mathbf{x}_i, \mathbf{z}_j) / \tau)}, \quad (10)$$

where $\text{sim}(\mathbf{x}_i, \mathbf{z}_i)$ denotes the cosine similarity of multi-view triple representations, i.e. \mathbf{x}_i and \mathbf{z}_i . The final contrastive loss is computed across all triples in a mini-batch.

3.3.3 *Triple confidence score assessment.* Upon the features learned by our proposed framework, we define the final confidence score function as follows:

$$C(h, r, t) = \sigma(\text{sim}(\mathbf{x}_i, \mathbf{z}_i) - \lambda \cdot E(h, r, t)), \quad (11)$$

Table 2: The statistical information of the datasets.

Dataset	Entities	Relations	Triples	Mean in-degree
FB15K	14,541	237	310,116	18.71
WN18RR	40,943	11	93,003	2.12
NELL-995	75,492	200	154,213	1.98

where $E(h, r, t) = \|\mathbf{e}_h + \mathbf{e}_r - \mathbf{e}_t\|_2$ reflects the degree of self contradictory within the triple, while $\text{sim}(\mathbf{x}_i, \mathbf{z}_i)$ measures the consistency of triple representations of the same sample from View I and View II. In practice, we adopt a trade-off parameter λ to balance the contribution of these two learning signals for error detection and then a sigmoid function is used to map the score into the range of $[0, 1]$. The learning process of CAGED is summarized in Algorithm 1.

4 EXPERIMENTS

To verify the effectiveness of the proposed framework CAGED, in this section, we conduct comprehensive experiments on three real-world KGs. Specifically, we aim to answer the following questions.

- **Q1 (Effectiveness):** How effective is CAGED compared with the state-of-the-art KG error detection methods?
- **Q2 (Ablation study):** How does each component of CAGED contribute to its performance?
- **Q3 (Parameter analysis):** How do the hyperparameters influence the performance of CAGED?
- **Q4 (Case study):** Is our proposed CAGED able to detect errors in real-world domain-specific KGs?

4.1 Experimental Settings

In this section, we introduce the detailed experimental settings, including the benchmark datasets, compared baseline methods and evaluation metrics.

4.1.1 *Datasets.* Following the previous study [15, 32], we employ three real-world datasets that are constructed with noisy triples to be 5%, 10% and 15% of the whole KGs based on the popular benchmarks, i.e. FB15k, WN18RR and NELL-995. Their statistical information is summarized in Table 2.

FB15K is a dataset adapted from Freebase Knowledge Base. It is expanded with textual relations, which are converted from textual mentions of ClueWeb12 corpus and annotations in Freebase.

WN18RR is a subset of WordNet, making up for the weakness of WN18. To avoid test leakage, it also excludes inverse relations using the same procedure as the derivation of FB15K. It contains 11 types of relations, and is less complex.

NELL-995 is derived from the 995th iteration of NELL, which is a huge and growing knowledge base. When constructing this dataset, the triples without reasoning values are eliminated before the Top-200 unique relations are selected.

4.1.2 *Baseline methods.* In the experiments, we include two categories of baselines. First, KG-embedding based methods, including **TransE** [33], **DistMult** [34], **Complex** [35]. To perform error detection, after we have learned the embedding representations, we access the confidence of triples based on the corresponding score functions. For example, in TransE, we use $\|\mathbf{e}_h + \mathbf{e}_r - \mathbf{e}_t\|_2$. Second, state-of-the-art KG error detection methods, i.e., **CKRL** [32], **KGtTm** [15], and **KGtst** [36]. **CKRL** advances TransE by taking

Table 3: Error detection results of Precision@K and Recall@K based on the three datasets with anomaly ratio = 5%.

		FB15K					WN18RR					NELL-995				
		K=1%	K=2%	K=3%	K=4%	K=5% ^a	K=1%	K=2%	K=3%	K=4%	K=5% ^a	K=1%	K=2%	K=3%	K=4%	K=5% ^a
Precision@K	TransE	0.756	0.674	0.605	0.546	0.488	0.581	0.488	0.371	0.345	0.331	0.659	0.550	0.476	0.423	0.383
	ComplEx	0.718	0.651	0.590	0.534	0.485	0.518	0.444	0.382	0.341	0.307	0.627	0.538	0.472	0.427	0.378
	DistMult	0.709	0.646	0.582	0.529	0.483	0.574	0.451	0.390	0.349	0.322	0.630	0.553	0.493	0.446	0.408
	CKRL	0.789	0.736	0.684	<u>0.630</u>	0.574	0.675	0.526	0.436	0.389	0.349	0.735	0.642	0.559	0.498	0.450
	KGTtm	0.815	<u>0.767</u>	<u>0.713</u>	0.612	<u>0.579</u>	<u>0.770</u>	<u>0.628</u>	<u>0.516</u>	<u>0.444</u>	<u>0.396</u>	<u>0.808</u>	<u>0.691</u>	<u>0.602</u>	<u>0.535</u>	0.481
	KGlst	<u>0.825</u>	0.754	0.703	0.617	0.569	0.747	0.599	0.476	0.407	0.379	0.782	0.678	0.584	0.528	<u>0.485</u>
	CAGED	0.852	0.796	0.735	0.665	0.595	0.826	0.726	0.632	0.541	0.469	0.850	0.736	0.644	0.573	0.516
Recall@K	TransE	0.151	0.270	0.363	0.437	0.488	0.116	0.195	0.223	0.276	0.331	0.132	0.220	0.285	0.338	0.383
	ComplEx	0.143	0.260	0.354	0.427	0.485	0.103	0.177	0.229	0.273	0.307	0.125	0.215	0.283	0.341	0.378
	DistMult	0.141	0.258	0.349	0.423	0.483	0.114	0.180	0.234	0.279	0.322	0.126	0.221	0.295	0.357	0.408
	CKRL	0.158	0.294	0.410	<u>0.504</u>	0.574	0.135	0.210	0.261	0.311	0.349	0.147	0.256	0.335	0.398	0.450
	KGTtm	0.163	<u>0.307</u>	<u>0.428</u>	0.490	<u>0.579</u>	<u>0.154</u>	<u>0.251</u>	<u>0.309</u>	<u>0.355</u>	<u>0.396</u>	<u>0.161</u>	<u>0.276</u>	<u>0.361</u>	<u>0.428</u>	0.481
	KGlst	<u>0.165</u>	0.302	0.422	0.494	0.569	0.149	0.240	0.285	0.325	0.379	0.156	0.271	0.350	0.422	<u>0.485</u>
	CAGED	0.171	0.318	0.441	0.532	0.595	0.165	0.290	0.379	0.433	0.469	0.170	0.294	0.386	0.458	0.516

^a Please note that according to Eq.(12), when K equals the anomaly ratio (i.e. 5%), $Precision@K$ and $Recall@K$ are equal.

all paths between the head entity and tail entity into consideration. **KGTtm** further enhances CKRL by integrating the global graph structure of the KG. **KGlst** proposes an unsupervised way to learn soft rules and detect errors based on these rules.

4.1.3 Evaluation metrics. We use ranking measures to evaluate the performance of all the compared approaches. Specifically, we rank all the triples in the target KG according to their confidence score in ascending order. A triple with a lower score would be more likely to be erroneous. To fairly evaluate the performance of the KG validation, we use the following two evaluation measures.

Precision@K represents the percentage of false triples found among the triples with top K lowest confidence scores.

Recall@K denotes the percentage of identified false triples over the total erroneous triples.

$$Precision@K = \frac{|\text{Errors Discovered in Top } K \text{ Ranking List}|}{|K|} \quad (12)$$

$$Recall@K = \frac{|\text{Errors Discovered in Top } K \text{ Ranking List}|}{|\text{Total Number of Errors in KG}|}$$

4.1.4 Implementation details. We implement all the baselines and the proposed framework with PyTorch. For the baseline methods, we use the released codes to conduct experiments. We use Nvidia RTX 3090 GPU server to train our proposed framework as well as all the baselines. To be more specific, we optimize all models with Adam optimizer, where the batch size is fixed at 256. We use the default Xavier initializer to initialize our model parameters, and the initial learning rate is 0.01. The embedding size is fixed to 100 for all models. We apply a grid search for hyperparameter tuning, the attention threshold parameter μ is searched from 0.001 to 0.2, the margin parameter γ from 0 to 1, and the trade-off coefficient λ between 0.001 to 1000. The number of neighbors is computed by taking the average number of neighbors of all triples in a dataset to make sure our model could adapt to datasets with different densities of neighbors, and thus make the best of the neighborhood information. To reduce the randomness, we use the random seed and report the average results of ten runs.

4.2 Effectiveness of CAGED (Q1)

To answer **Q1**, we conduct comprehensive experiments on three real-world KGs. The experimental results with anomaly ratio equals to 5% are summarized in Table 3. The results with other anomaly ratios are similar as shown in Table 4.

In summary, we have three observations. First, KG error detection methods, i.e., CAGED, CKRL, KGTtm, and KGlst are superior to KG-embedding based methods, i.e., TransE, ComplEx, and DistMult. The reason is that these KG embedding frameworks do not consider the errors in KG, thus can not learn discriminative representations for normal and noisy triples. Second, CAGED achieves the best performance compared with all baselines on the three datasets in terms of recall and precision. For example, comparing with the second best results, CAGED achieves 1.6%, 7.3% and 3.1% improvements in FB15K, WN18RR, and NELL-995, respectively, when K equals 5% and anomaly ratio equals 5%. It is because all baselines use synthetic false triples to train the model, and thus is hard to detect nontrivial errors with complex patterns. While our CAGED introduces the multi-view contrasting to perform optimization. So, it is capable of exploiting the relational structure within the triples and cross-view features for error detection. Third, CAGED outperforms all baselines more significantly when the K is smaller. To further verify the capability of CAGED in detecting errors in KGs, we conduct a supplementary experiment on NELL-995 with noisy triples to be different ratios of 5%, 10% and 15%. As shown in Table 4, our model consistently achieves the best performance in all cases, which demonstrates the capability and stability of our model in detecting errors under different scenarios.

4.3 Ablation Study (Q2)

We now investigate the question **Q2**. CAGED has three components, i.e., KG augmentation, error-aware encoder (EaGNN), joint estimation. Four pairs of variants of CAGED are used for this ablation study, as shown in Table 5. Their detection performance on

Table 4: Error detection results on NELL-995 with different anomaly ratios.

	Ratio	5%			10%			15%		
		$K=5\%^a$	$K=10\%$	$K=15\%$	$K=5\%$	$K=10\%^a$	$K=15\%$	$K=5\%$	$K=10\%$	$K=15\%^a$
<i>Precision@K</i>	TransE	0.383	0.285	0.225	0.626	0.499	0.407	0.702	0.621	0.535
	ComplEx	0.378	0.289	0.231	0.614	0.507	0.402	0.696	0.589	0.528
	DistMult	0.408	0.298	0.227	0.633	0.510	0.414	0.718	0.618	0.548
	CKRL	0.450	0.306	0.236	0.679	0.524	0.421	0.745	0.646	0.560
	KGTtm	0.481	<u>0.320</u>	0.242	0.713	0.527	0.437	0.788	<u>0.673</u>	<u>0.576</u>
	KG1st	<u>0.485</u>	0.317	<u>0.244</u>	<u>0.748</u>	<u>0.552</u>	<u>0.440</u>	<u>0.791</u>	0.663	0.569
	CAGED	0.516	0.325	0.251	0.799	0.585	0.458	0.823	0.729	0.599
<i>Recall@K</i>	TransE	0.383	0.57	0.675	0.313	0.499	0.612	0.234	0.414	0.535
	ComplEx	0.378	0.578	0.693	0.307	0.507	0.603	0.232	0.393	0.528
	DistMult	0.408	0.596	0.681	0.317	0.510	0.621	0.239	0.412	0.548
	CKRL	0.450	0.612	0.708	0.340	0.524	0.632	0.248	0.431	0.560
	KGTtm	0.481	<u>0.640</u>	0.726	0.357	0.527	0.656	0.263	<u>0.449</u>	<u>0.576</u>
	KG1st	<u>0.485</u>	0.634	<u>0.732</u>	<u>0.374</u>	<u>0.552</u>	<u>0.660</u>	<u>0.264</u>	0.442	0.569
	CAGED	0.516	0.650	0.753	0.400	0.585	0.687	0.274	0.486	0.599

^a It is noted that when K equals anomaly ratio, *Precision@K* and *Recall@K* both have the same value.

Table 5: Four pairs of variants of CAGED.

	Original component	Replacement
Var_DN	KG augmentation	Node dropping
Var_EP	KG augmentation	Edge perturbation
Var_Concat	Bi-LSTM units	Only concatenation
Var_LSTM	Bi-LSTM units	LSTM units
Var_GCN	EaGNN	R-GCN
Var_GAT	EaGNN	KGAT
Var_Local	Joint optimization	Only negative sampling
Var_Global	Joint optimization	Only contrastive learning

NELL-995 is included in Table 6. We omit the results on FB15K and WN18RR since they show similar trends.

4.3.1 Analysis of component 1 - tailored KG augmentation. To demonstrate the effectiveness of the KG augmentation method, we compare it with two other general augmentation approaches, i.e., node dropping and edge perturbation. The corresponding two variants are Var_DN and Var_EP. As shown in Table 6, Var_DN and Var_EP perform worse, even comparing to some embedding-based baselines, such as TransE. It is because directly applying node dropping and edge perturbation will destroy the triple structure and introduce noise to KGs, and thus influence training a correct error detector. The huge gap of performance between these two variants and our proposed model supports the assumption that our tailor augmentation mechanism could generate different views on the premise of avoiding undermining the distribution of potential errors, which enables our contrastive framework to perform effective error detection on KGs.

4.3.2 Analysis of component 2 - EaGNN. To validate our encoder EaGNN, we replace it with SOTA KGNNs, i.e., R-GCN [37] and KGAT [38]. The corresponding variants are Var_GCN and Var_GAT. From Table 6, we can see that the Var_GAT shows superior performance compared to Var_GCN. But there is still a significant gap

in comparison with CAGED. It is because that R-GCN is a typical message-passing model designed for KG representation learning. It assumes that all observed triples are correct, and may overfit noisy triples and fail to detect such errors. KGAT adopts an attention mechanism to learn KG representations which may filter out part of noisy information. Nevertheless, Var_GAT still shows no excellent performance on error detection task. It is because that KGAT applies attention mechanism from the perspective of entities and relations, while KG errors always occur at the triple level, i.e., the entities naturally exist in KGs, the error is usually a mismatch of the head entity, tail entity, and the corresponding relation. Our tailored error-aware attention layer in EaGNN takes the triple-level embeddings as input can effectively filter out noisy triples.

To evaluate the effectiveness of EaGNN in capturing translation structure within each triple, we introduce Var_LSTM and Var_Concat. Var_LSTM replaces Bi-LSTM units in our local information modeling layers by LSTM. Var_Concat removes the Bi-LSTM units, and directly concatenates the randomly initialized embeddings of head entity, relation, and tail entity. From the results in Table 6, we observe that the performance of Var_LSTM and Var_Concat drops significantly while Var_LSTM shows better performance compared to Var_Concat. It is because concatenation ignores the local structure information.

4.3.3 Analysis of component 3 - joint estimation. To evaluate the effectiveness of our joint optimization approach, we only use translation based KG embedding loss, i.e., negative sampling, and contrastive learning loss to train the model in Var_Local and Var_Global, correspondingly. From Table 6, we can observe that both these two variants are inferior to our proposed model while Var_Local performs worse, even comparing to some straightforward baseline methods, such as DistMult. It is because that negative sampling could only help our model learn rich structure and semantics information within triples, and such kind of local feature is not enough to enable effective error detection on the whole KG. As a complement to the negative sampling, contrastive learning could help the model

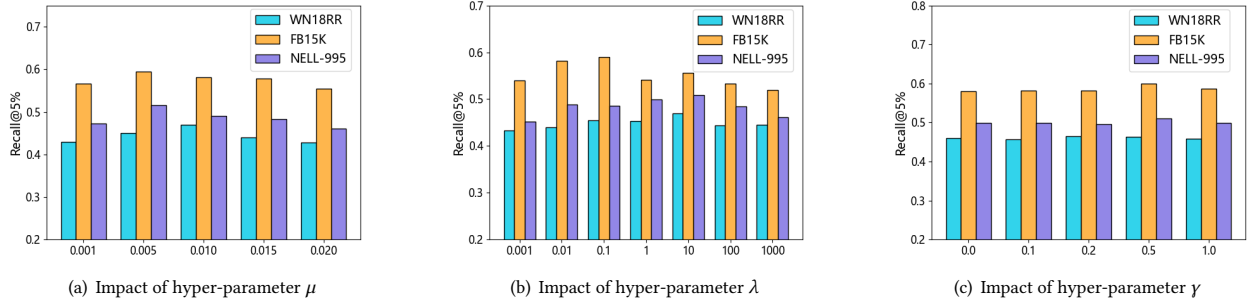


Figure 2: Impact of hyperparameters on the three datasets.

Table 6: The comparisons of CAGED and its variants on NELL-995 with ratio of errors equals 5%.

	Precision@K					Recall@K				
	1%	2%	3%	4%	5%	1%	2%	3%	4%	5%
CAGED	.850	.736	.644	.573	.516	.170	.294	.386	.458	.516
Var_DN	.613	.490	.412	.358	.325	.122	.196	.247	.286	.325
Var_EP	.593	.500	.416	.354	.313	.118	.200	.249	.283	.313
Var_Concat	.767	.648	.532	.468	.440	.153	.259	.319	.374	.440
Var_LSTM	.797	.662	.564	.500	.458	.159	.264	.338	.400	.458
Var_GCN	.760	.623	.529	.464	.414	.152	.249	.317	.371	.414
Var_GAT	.782	.629	.551	.471	.426	.156	.251	.330	.377	.426
Var_Local	.724	.603	.504	.433	.383	.144	.241	.302	.346	.383
Var_Global	.783	.652	.560	.492	.443	.156	.261	.336	.393	.443

learn distinguishable across-view representations. With the rich and distinguishable across-view representations, error detection on KGs is much more effective.

4.4 Parameter Analysis (Q3)

In this section, we mainly investigate the impact of three key parameters in our framework and report the results in Figure 2.

As discussed in Eq.(5), μ is the parameter that controls the attention threshold. The smaller μ allows more irrelevant neighbors to be attended over when learning the global triple embedding representations, while a larger μ forces our encoder to focus on a few but closely related neighbor triples. To investigate the impact of μ , we vary it from 0.001 to 0.020. Figure 2(a) shows the detection performance in terms of Recall@K as the evaluation metrics. When $\mu = 0.020$, performance on all three datasets is not ideal enough. This happens because only few global contextual features have been mined when μ is large. On the missing of enough contextual information, the error detection task will be very challenging. As μ decreases, our encoder starts to attend over more neighbor triples to learn much richer global information, so the performance keeps improving. As shown in Figure 2(a), performance on FB15K and NELL-995 achieves optimal when μ is close to 0.005 while WN18RR gets the best performance at $\mu = 0.01$. However, as μ continues to fall down, the contrary result happens, i.e., the performance on these three datasets markedly decreases. This is because smaller μ allows more irrelevant neighbors involving in the representation

Table 7: The average running time for one iteration.

Method	FB15K	WN18RR	NELL-995
TransE	0.05	0.07	0.09
ComplEx	0.06	0.09	0.12
DistMult	0.07	0.11	0.14
CKRL	0.45	0.29	0.32
KGTtm	0.44	0.32	0.42
KG1st	0.57	0.49	0.49
CAGED	0.41	0.22	0.38

learning process, correspondingly, some potential noisy triples may swoop in, which leads to overfitting issues for error detection.

As shown in Eq.(11), λ is the trade-off coefficient that balances the contribution of two learning signals for error detection, i.e., the inconsistency of triple representations across multi-views (e.g., \mathbf{x}_A and \mathbf{z}_A) and the self-contradictory within the triple embedding ($\mathbf{e}_h, \mathbf{e}_r, \mathbf{e}_t$). The larger λ indicates that effective error detection relies more on the signal of the self-contradictory, and vice versa. To investigate the impact of λ , we vary it from 10^{-3} to 10^3 with a multiple 10. We perform detection task on all datasets and the detection performance in terms of Recall@K is shown in Figure 2(b). From the results, we observe that 1) WN18RR achieves optimal result when λ is around 10, but, from the global perspective, the performance on WN18RR seems to be not sensitive with parameter λ since its value is stable basically without great fluctuate, similar results are observed on NELL-995. 2) For FB15K, the detection performance is much better when the value of λ is less than 1, with an optimal result at 0.1. After a detailed analysis of these three datasets, we figure out that the potential reason for this observation is that FB15K has a more dense data structure compared with the other datasets. Specifically, there exists more triple-level neighbors in such a dense data structure which assistant our model to learn more comprehensive global triple embedding representations from different views. With richer triple representations, the learning signal of the inconsistency across multi-views can play a more important role in detection task.

γ serves as a margin to control the distance between the representations of negative and positive pairs in the translation-based KG embedding loss function, i.e., Eq.(8). To investigate the impact of γ , we vary it from 0.1 to 1.0. Figure 2(c) shows the detection performance is stable basically without great fluctuation as γ varies.




Figure 3: The scatter plot of confidence scores distribution.

This is mainly because we adopt a joint training approach, avoiding the model from falling into suboptimal solutions.

4.5 Case Study (Q4)

The three benchmark datasets used in this paper are all general-purpose KGs. To further verify the effectiveness of our proposed CAGED, we conduct a case study on another real-world domain-specific knowledge graph, KnowLife [39]. KnowLife is a huge medical knowledge repository, which is automatically constructed based on scientific literature. To access ground-truth labels of errors, we manually label triples that are randomly selected from the official KnowLife website. Concretely, we first employ two participants to manually label each triple by cross-checking with the information presented on the original website. If the labels provided by these two participants are inconsistent, then this triple will be assigned to the third one for further investigation. Majority voting is applied to decide the final label of each triple. Following this procedure, we label more than 5000 triples, which are divided into two sets, i.e., the set of positive samples and the set of noisy samples.

To make the comparison more simple and intuitive, we display the triple effectiveness scores learned by our model and the best baseline method, i.e. KGTtm, in a centralized coordinate system, as shown in Figure 3. Specifically, the Figure 3(b) in the right visualizes the distribution of the values of the normal and noisy triples learned by KGTtm. It can be seen that a considerable proportion of erroneous triples have been wrongly classified with high confidence scores. But for our model, as shown in Figure 3(a), the confidence value of normal triples are mainly distributed in the upper region, while the values of the erroneous triples are mainly concentrated in the lower region, which verifies the superiority of our CAGED for KG error detection in real-world scenarios.

4.6 Efficiency Analysis of CAGED

To measure the efficiency of CAGED, we record the average computation time of one iteration for all models under the same setting, and report the results in Table 7. All experiments are conducted with one Nvidia RTX 3090 GPU. From Table 7, we can observe that embedding methods including TransE, ComplEx, and DistMult run faster than KG error detection methods in general since they only need to calculate the mean square losses. TransE costs less time than ComplEx and DistMult on all datasets. Comparing to other three error detection methods, our model shows comparable efficiency on average, while runs faster especially in the large-scale dataset, FB15K. This observation validates the efficiency of our model comparing to baseline methods. In summary, our model is more promising for KG error detection in real-world scenarios, due

to the fact that it outperforms state-of-the-art KG error detection algorithms, with comparable or even better running efficiency.

5 RELATED WORK

Recent studies in KG error detection could be roughly categorized into two classes. First, **classification-based** models. Melo and Paulheim [18] design a classifier based on entity categories and features of related paths. Jia et al. [15] integrate features from entity level, relation level, and global level to build a classifier. Ge et al. [16] explore using active learning to train a classifier based on a small number of labels. A remaining issue of classification-based models is the penalty of imbalanced classes, because there are far more correct triples than errors. Second, **embedding-based** models. TransE [33] is a typical KG embedding algorithm. It uses margin loss and negative sampling to induce $\mathbf{e}_h + \mathbf{e}_r \approx \mathbf{e}_t$. Several studies [40–42] propose to evaluate the confidence of triples based on the embedding loss, e.g., $\mathbf{e}_h + \mathbf{e}_r - \mathbf{e}_t$.

Efforts have also been devoted to employing extra information to facilitate KG error detection. In some KGs, categories of entities are (partially) available. Paulheim and Bizer [43] employ the statistical distributions of relations and entity categories to assess the suspiciousness of triples. Attempts have been made to leverage entity categories to perform clustering-based outlier detection [14] or construct path-based rules [44]. However, entity categories are not available in many KGs. Other beneficial extra information includes related webpages [45], crowdsourcing evaluations [46], and external human-curated KGs [42].

6 CONCLUSIONS AND FUTURE WORK

Existing KG error detection methods mainly rely on synthetic false triples to train their models, which are generated by replacing the head or tail entity in a triple with a random entity. These synthetic errors are not in line with the ones in practice, e.g., (Bruce_Lee, place_of_birth, Chinatown), in which the three elements in a triple are highly relevant. To this end, we propose an effective framework named CAGED, which integrates contrastive learning and KG error detection. CAGED provides a novel view towards knowledge graph modeling. It creates different hyper-views of the same graph structure by centering triples, and then learns the representation by maximizing agreement of each instance’s representation between these views. By learning to contrast elaborate instance pairs, the model can acquire informative features that can be used as an effective learning signal to detect errors in KGs. Experimental results on three real-world KGs demonstrates that our model outperforms the state-of-the-art KG error detection algorithms in terms of different evaluation metrics. In the future, we will develop error-aware KG embedding model to facilitate KG-related downstream tasks on the basis of CAGED.

ACKNOWLEDGMENTS

The authors gratefully acknowledge receipt of the following financial support for the research, authorship, and/or publication of this article. This work was supported in full by the Hong Kong Polytechnic University, Departmental General Research Fund (Project Number: P0036738).

REFERENCES

- [1] Hongwei Wang, Miao Zhao, Xing Xie, Wenjie Li, and Minyi Guo. Knowledge graph convolutional networks for recommender systems. In *WWW*, 2019.
- [2] Xiao Huang, Jingyuan Zhang, Dingcheng Li, and Ping Li. Knowledge graph embedding based question answering. In *WSDM*, 2019.
- [3] Kurt D. Bollacker, Robert P. Cook, and Patrick Tufts. Freebase: A shared database of structured general human knowledge. In *AAAI*, 2007.
- [4] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. Dbpedia - A large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6(2):167–195, 2015.
- [5] Stefan Heindorf, Martin Potthast, Benno Stein, and Gregor Engels. Vandalism detection in wikidata. In *CIKM*, 2016.
- [6] Farzaneh Mahdisoltani, Joanna Biega, and Fabian M. Suchanek. YAGO3: A knowledge base from multilingual wikipedias. In *CIJR*, 2015.
- [7] Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., and Tom M. Mitchell. Toward an architecture for never-ending language learning. In *AAAI 2010*, 2010.
- [8] Heiko Paulheim. Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web*, 2017.
- [9] Rakesh Agrawal, Tomasz Imielinski, and Arun N. Swami. Mining association rules between sets of items in large databases. In Peter Buneman and Sushil Jajodia, editors, *SIGMOD*, pages 207–216, 1999.
- [10] Shu Guo, Quan Wang, Lihong Wang, Bin Wang, and Li Guo. Knowledge graph embedding with iterative guidance from soft rules. In *AAAI*, 2018.
- [11] Yurong Cheng, Lei Chen, Ye Yuan, and Guoren Wang. Rule-based graph repairing: Semantic and efficient repairing methods. In *ICDE*, 2018.
- [12] Luis Antonio Galárraga, Christina Teflioudi, Katja Hose, and Fabian M. Suchanek. AMIE: association rule mining under incomplete evidence in ontological knowledge bases. In *WWW*, 2013.
- [13] Thomas Pellissier Tanon, Daria Stepanova, Simon Razniewski, Paramita Mirza, and Gerhard Weikum. Completeness-aware rule learning from knowledge graphs. In *IJCAI*, 2018.
- [14] Heiko Paulheim and Aldo Gangemi. Serving dbpedia with DOLCE - more than just adding a cherry on top. In *ISWC*, 2015.
- [15] Shengbin Jia, Yang Xiang, Xiaojun Chen, Kun Wang, and Shijia E. Triple trustworthiness measurement for knowledge graph. In *WWW*, 2019.
- [16] Congcong Ge, Yunjun Gao, Honghui Weng, Chong Zhang, Xiaoye Miao, and Baihua Zheng. Kgclean: An embedding powered knowledge graph cleaning framework. *arXiv preprint arXiv:2004.14478*, 2020.
- [17] Zhiming Xu, Xiao Huang, Yue Zhao, Yushun Dong, and Jundong Li. Contrastive attributed network anomaly detection with data augmentation. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 444–457. Springer, 2022.
- [18] André Melo and Heiko Paulheim. Detection of relation assertion errors in knowledge graphs. In *Knowledge Capture Conference*, 2017.
- [19] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [20] Olivier J. Hénaff. Data-efficient image recognition with contrastive predictive coding. In *ICML*, 2020.
- [21] R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR*, 2019.
- [22] Thomas N. Kipf, Elise van der Pol, and Max Welling. Contrastive learning of structured world models. In *ICLR*, 2020.
- [23] Zhen Peng, Wenbing Huang, Minnan Luo, Qinghua Zheng, Yu Rong, Tingyang Xu, and Junzhou Huang. Graph representation learning via graphical mutual information maximization. In *WWW*, 2020.
- [24] Kaveh Hassani and Amir Hosein Khas Ahmadi. Contrastive multi-view representation learning on graphs. In *ICML*, 2020.
- [25] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. Deep graph contrastive representation learning. *arXiv preprint arXiv:2006.04131*, 2020.
- [26] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. In *NeurIPS*, 2020.
- [27] Jiaqi Zeng and Pengtao Xie. Contrastive self-supervised learning for graph classification. In *AAAI*, 2021.
- [28] Yingheng Wang, Yaosen Min, Xin Chen, and Ji Wu. Multi-view graph contrastive representation learning for drug-drug interaction prediction. In *WWW*, 2021.
- [29] Xin Xia, Hongzhi Yin, Junliang Yu, Qinyong Wang, Lizhen Cui, and Xiangliang Zhang. Self-supervised hypergraph convolutional networks for session-based recommendation. In *AAAI*, pages 4503–4511, 2021.
- [30] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *Advances in neural information processing systems*, 33:11839–11852, 2020.
- [31] Chengwei Chen, Yuan Xie, Shaohui Lin, Ruizhi Qiao, Jian Zhou, Xin Tan, Yi Zhang, and Lizhuang Ma. Novelty detection via contrastive learning with negative data augmentation. *arXiv preprint arXiv:2106.09958*, 2021.
- [32] Ruobing Xie, Zhiyuan Liu, Fen Lin, and Leyu Lin. Does william shakespeare REALLY write hamlet? knowledge representation learning with confidence. In *AAAI*, 2018.
- [33] Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *NeurIPS*, 2013.
- [34] Bishan Yang, Wen tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases, 2015.
- [35] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *ICML*, 2016.
- [36] Caleb Belth, Xinyi Zheng, Jilles Vreeken, and Danai Koutra. What is normal, what is strange, and what is missing in a knowledge graph: Unified characterization via inductive summarization. In *WWW*, 2020.
- [37] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pages 593–607. Springer, 2018.
- [38] Deepak Nathani, Jatin Chauhan, Charu Sharma, and Manohar Kaul. Learning attention-based embeddings for relation prediction in knowledge graphs. In *ACL*, 2019.
- [39] Patrick Ernst, Amy Siu, and Gerhard Weikum. Knowlife: a versatile approach for constructing a large knowledge graph for biomedical sciences. *BMC bioinformatics*, 16(1):1–13, 2015.
- [40] Aibo Guo, Zhen Tan, and Xiang Zhao. Measuring triplet trustworthiness in knowledge graphs via expanded relation detection. In *KSEM*, pages 65–76, 2020.
- [41] Yu Zhao and Ji Liu. SCEF: A support-confidence-aware embedding framework for knowledge graph refinement. *arXiv preprint arXiv:1902.06377*, 2019.
- [42] Yaqing Wang, Fenglong Ma, and Jing Gao. Efficient knowledge graph validation via cross-graph representation learning. In *CIKM*, pages 1595–1604, 2020.
- [43] Heiko Paulheim and Christian Bizer. Improving the quality of linked data using statistical distributions. *Int. J. Semantic Web Inf. Syst.*, 2014.
- [44] Baoxu Shi and Tim Wening. Discriminative predicate path mining for fact checking in knowledge graphs. *Knowl. Based Syst.*, 2016.
- [45] Jens Lehmann, Daniel Gerber, Mohamed Morsey, and Axel-Cyrille Ngonga Ngomo. Defacto - deep fact validation. In *ISWC*, 2012.
- [46] Manuela Nayantara Jeyaraj, Srinath Perera, Malith Jayasinghe, and Nadheesh Jihan. Probabilistic error detection model for knowledge graph refinement. In *CICLING*, 2019.