



# Improving Knowledge Graph Entity Alignment with Graph Augmentation

Feng Xie, Xiang Zeng, Bin Zhou<sup>(✉)</sup>, and Yusong Tan

College of Computer, National University of Defense Technology, Changsha, China  
{xiefeng,zengxiang,binzhou,ystan}@nudt.edu.cn

**Abstract.** Entity alignment (EA) which links equivalent entities across different knowledge graphs (KGs) plays a crucial role in knowledge fusion. In recent years, graph neural networks (GNNs) have been successfully applied in many embedding-based EA methods. However, existing GNN-based methods either suffer from the structural heterogeneity issue that especially appears in the real KG distributions or ignore the heterogeneous representation learning for unseen (unlabeled) entities, which would lead the model to overfit on few alignment seeds (i.e., training data) and thus cause unsatisfactory alignment performance. To enhance the EA ability, we propose GAEA, a novel EA approach based on graph augmentation. In this model, we design a simple Entity-Relation (ER) Encoder to generate latent representations for entities via jointly modeling comprehensive structural information and rich relation semantics. Moreover, we use graph augmentation to create two graph views for margin-based alignment learning and contrastive entity representation learning, thus mitigating the negative influence caused by structural heterogeneity and sparse seeds. Extensive experiments conducted on benchmark datasets demonstrate the effectiveness of our method. Our codes are available at <https://github.com/Xiefeng69/GAEA>.

**Keywords:** Knowledge Graph · Entity Alignment · Graph Neural Networks · Graph Augmentation · Knowledge Representation

## 1 Introduction

Knowledge graphs (KGs) can effectively organize and represent facts about the world in a structured fashion. More and more KGs have been constructed based on different data sources or for different purposes. Therefore, the knowledge contained in different KGs is far from complete yet complementary [22]. Entity alignment (EA) which aims to link semantically equivalent entities located on different KGs has attracted increasing attention since it could facilitate knowledge integration and thus promote knowledge-driven applications, such as question answering, recommender systems, and semantic search.

In recent years, embedding-based EA methods [2, 3, 11, 14, 20, 22, 26, 30] have achieved decent results. The general pipeline can be summarized into two steps: (I) generating low-dimensional embeddings (latent representations) for entities

via KG encoder (e.g., TransE [1]), and then (II) pulling two KGs into a unified embedding space through prior alignment seeds and pairing each entity by distance metrics (e.g., Euclidean distance). Moreover, some works further improve the EA performance by introducing extra information, such as entity names [29], attributes [5, 12], and literal descriptions [24], while these discriminative features are usually privacy sensitive, noise polluted, and hard to collect [8].

Due to the powerful structure learning capability, Graph Neural Networks (GNNs) like GCN [4] and GAT [17] have been employed as the encoder with Siamese architecture (i.e., shared-parameter) for many embedding-based models [6, 14, 20, 26]. KGs are heterogeneous, especially in real KG distributions, which means entities that have the same referent in different KGs usually have dissimilar relational neighborhood. To address this problem, existing GNN-based models modify and improve GNN variants to better capture structural information in KGs, e.g., AliNet [14] adopts multi-hop aggregation with gating mechanism to expand neighborhood ranges and RDGCN [21] incorporates relation features via attention interactions for embedding learning. However, these models introduce a large number of neural network operations and ignore representation learning for unseen entities, which will tend to make the models overfit on few alignment seeds and thus undermine their generalization and performance.

In this paper, we propose GAEA, a novel knowledge graph entity alignment model based on graph augmentation. Firstly, we design an Entity-Relation (ER) Encoder to generate entity representations via jointly leveraging neighborhood structures and relation semantics in KGs. Then, we apply graph augmentation to increase the structural diversity of input KG in the alignment learning process, which encourages the model to capture the semantic importance of different neighbors and enforces the model to obtain stable representations against structure perturbation, thus mitigating overfitting issue to some extent. Moreover, since graph augmentation can inherently generate two distinct graph views without extra parameters, we can let the model perceive structural differences and further improve the feature learning for (unseen) entities by applying contrastive entity representation learning to maximize the consistency between the original KG and augmented KG [19, 25]. Our experiments on benchmark datasets OpenEA [15] show that GAEA outperforms the existing state-of-the-art embedding-based EA methods. We also conduct thorough auxiliary analyses to demonstrate the effectiveness of incorporating graph augmentation techniques.

## 2 Related Works

Entity alignment is a fundamental task to identify the same entities across different KGs, which has attracted increasing attention in recent years. The existing embedding-based methods can be roughly divided into two categories:

1. **Structure-based models.** These models solely rely on the original structure information of KGs (i.e., triples) to align entities. Previous methods mainly use knowledge representation learning to generate low-dimensional embeddings for entities [2, 7, 30]. For example, MTransE [2] applies TransE [1] to

embed different KGs into independent vector spaces and constructs transitions via proposed alignment modules. Inspired by the powerful structure learning ability of Graph Neural Networks (GNNs), a large body of works begin to focus on employing GNNs as the encoder. GCN-Align [20] incorporates GCN [4] to capture entities’ neighborhood structures for the first time and achieves promising results. Subsequent works not only apply various GNN variants, like GAT [17], but also improve the structure awareness by overcoming heterogeneity of different KGs [3, 14, 21], capturing multi-context structural features [22], and infusing relation semantics [6, 11].

2. **Enhancement-based models.** These models aim to build a high-accuracy alignment system using designed alignment strategies or extra information. BootEA [13] applies iterative learning to find potential alignments and adds them to the training set for data augmentation. CEA [28] formulates alignment inference as a stable matching problem to model collective signals, successfully guaranteeing 1-to-1 alignment. Other effective models introduce extra information to enhance the alignment performance, including entity names [29], attributes [5, 12], and literal descriptions [24].

In this work, we aim to improve the performance and efficiency of entity alignment only utilizing structural contexts which are abundant and always available without privacy issues in the real-world KGs.

### 3 Preliminaries

**Knowledge Graph.** A knowledge graph (KG) is formalized as  $G = (E, R, T)$ , where  $E$  and  $R$  refer to the set of entities and the set of relations, respectively.  $T = E \times R \times E = \{(h, r, t) | h, t \in E \wedge r \in R\}$  is the set of triples, where  $h$ ,  $r$ , and  $t$  denote the head entity, connected relation, tail entity, respectively.

**Entity Alignment.** Given two KGs:  $G_s = (E_s, R_s, T_s)$  as the source KG and  $G_t = (E_t, R_t, T_t)$  as the target KG, and few alignment seeds (aka pre-aligned entity pairs)  $S = \{(e_i, e_j) | e_i \in E_s \wedge e_j \in E_t \wedge e_i \equiv e_j\}$ , where  $\equiv$  means equivalence relationship, entity alignment (EA) aims to seek remaining equivalent entities located on different KGs via entity representations.

**Augmented Graph.** Graph augmentation techniques will generate a perturbed version of the original graph, i.e., augmented graph, by augmentation strategies (e.g., node dropping, edge perturbation). In order not to introduce wrong facts, we only choose edge dropping in this work. At each training iteration, we randomly drop out some triples based on the deletion ratio  $r \sim \text{uniform}(0, pr)$ , where  $pr$  is a preset upper bound of the deletion ratio. The augmented graphs for  $G_s$  and  $G_t$  are denoted as  $G_s^{aug}$  and  $G_t^{aug}$ , respectively. Note that we do not consider deleting the triples associated with entities whose degree is less than 2, because these long-tail entities have sparse neighborhood structures inherently.

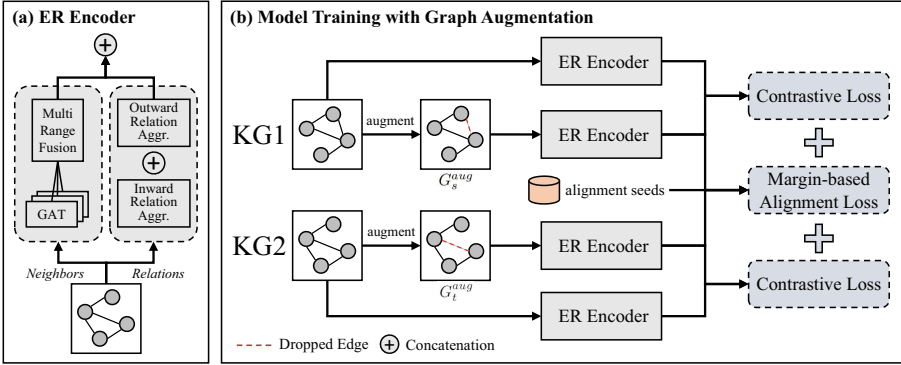


Fig. 1. The framework of our proposed GAEA.

## 4 Methodology

This section details our proposed method, termed as GAEA, which is drawn in Fig. 1: (a) Entity-Relation (ER) Encoder which generates latent representations for entities by capturing neighborhood structures and relation semantics jointly; (b) the training process of GAEA can be decomposed into multiple epochs, and in each epoch, we incorporate graph augmentation to conduct margin-based alignment learning and contrastive entity representation learning.

**Initialization.** At the beginning, we randomly initialize entity embeddings  $\mathbf{H}^{ent} \in \mathbb{R}^{(|E_s|+|E_t|) \times d_{ent}}$  and relation embeddings  $\mathbf{H}^{rel} \in \mathbb{R}^{|R_s \cup R_t| \times d_{rel}}$ , where  $d_{ent}$  and  $d_{rel}$  are the embedding dimension of entities and relations, respectively.

### 4.1 Entity-Relation Encoder

Here, we present the Entity-Relation Encoder (ER Encoder for short), which aims to fully capture the contextual information of entities using two aspects jointly: (I) neighborhood structures and (II) relation semantics.

**Neighborhood Aggregator.** First, we aggregate neighbor entities’ information to the central entity. The rationality of neighborhood aggregator lies in the structure assumption that, equivalent entities tend to have similar neighbor structures [20]. Moreover, leveraging multi-range neighborhood structures is capable of providing more alignment evidence and mitigating the structural heterogeneity issue. In this work, we apply Graph Attention Network (GAT) [17] to allow the central entity to learn the importance of different neighbors and thus selectively aggregate surrounding information, and we then recursively capture multi-range neighbor information by stacking multiple layers:

$$\mathbf{h}_{e_i}^{(l)} = \sum_{e_j \in N_{e_i}} \alpha_{ij} \mathbf{h}_{e_j}^{(l-1)}, \quad (1)$$

$$\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(\mathbf{a}^\top [\mathbf{W}_g \mathbf{h}_{e_i} \oplus \mathbf{W}_g \mathbf{h}_{e_j}]))}{\sum_{e_k \in N_{e_i}} \exp(\text{LeakyReLU}(\mathbf{a}^\top [\mathbf{W}_g \mathbf{h}_{e_i} \oplus \mathbf{W}_g \mathbf{h}_{e_k}]))}, \quad (2)$$

where  $\top$  represents transposition,  $\oplus$  is the concatenation operation,  $\mathbf{W}_g$  and  $\mathbf{a}$  are the transformation parameter and attention transformation vector, respectively.  $N_{e_i}$  means the neighbor set of entity  $e_i$  in KG, and  $\alpha_{ij}$  indicates the learned importance of entity  $e_j$  to entity  $e_i$ .  $\mathbf{h}_{e_i}^{(l)}$  denotes the embedding of  $e_i$  at  $l$ -th layer (total  $L$  layers) with  $\mathbf{H}^{(0)} = \mathbf{H}^{ent}$ . Note that here we remove the feature transformation and nonlinear activation that act on input embeddings in vanilla GAT since we mainly focus on information aggregation. We only use  $\mathbf{W}_g$  and  $\mathbf{a}$  to make each entity aware of its neighborhood contexts.

After multi-layer GAT, we obtain the multi-range neighborhood structural representation matrix for each entity, i.e.,  $\mathbf{H}_{e_i}^m = [\mathbf{h}_{e_i}^{(1)}, \dots, \mathbf{h}_{e_i}^{(L)}] \in \mathbb{R}^{L \times d_{ent}}$  for  $e_i$ . Since different neighborhood ranges have different contributions to characterize the central entity, it is necessary to employ a mechanism to adaptively control the flow of each range and thus reduce noise. Inspired by the skipping connections in neural networks [10, 14, 23], we firstly utilize a Scaled Dot-Product Attention mechanism [16] to learn the importance of each range, and then fuse small-range and wide-range representations by weighted average:

$$[\hat{\mathbf{h}}_{e_i}^{(1)}, \dots, \hat{\mathbf{h}}_{e_i}^{(L)}] = \text{softmax}\left(\frac{(\mathbf{H}_{e_i}^m \mathbf{W}_q)(\mathbf{H}_{e_i}^m \mathbf{W}_k)^\top}{\sqrt{d_{ent}}}\right) \mathbf{H}_{e_i}^m \quad (3)$$

$$\mathbf{h}_{e_i}^n = \frac{1}{L} \sum_{l=1}^L \hat{\mathbf{h}}_{e_i}^{(l)}, \quad (4)$$

where  $1/\sqrt{d_{ent}}$  is the scaling factor,  $\mathbf{W}_q$  and  $\mathbf{W}_k$  are the learnable parameter matrices, and  $\mathbf{h}_{e_i}^n$  is the output of neighborhood aggregator.

**Relation Aggregator.** Relation-level information which carries rich semantics is vital to align entities in KGs [24, 29] because two equivalent entities may share overlapping relations. MRAEA [6] pointed out that relation directions impose extra but delicate constraints on the head and tail entity individually. Therefore, in this work, we directly use two mean aggregators to gather outward relation semantics and inward relation semantics separately to provide supplementary alignment signals for heterogeneous KGs:

$$\mathbf{h}_{e_i}^r = \frac{1}{|N_{e_i}^{r+}|} \sum_{r \in N_{e_i}^{r+}} \mathbf{h}_r^{rel} \oplus \frac{1}{|N_{e_i}^{r-}|} \sum_{r \in N_{e_i}^{r-}} \mathbf{h}_r^{rel}, \quad (5)$$

where  $N_{e_i}^{r+}$  and  $N_{e_i}^{r-}$  are the outward and inward relation set of  $e_i$ , respectively.

**Feature Fusion.** Finally, we concatenate two aspects of information:

$$\tilde{\mathbf{h}}_{e_i} = \mathbf{h}_{e_i}^n \oplus \mathbf{h}_{e_i}^r, \quad (6)$$

where  $\tilde{\mathbf{h}}_{e_i} \in \mathbb{R}^{d_{ent}+2 \times d_{rel}}$  is the final output representation of ER Encoder for  $e_i$ . In the following training process, the ER Encoder is shared for  $G_s$ ,  $G_t$ , and their augmented graphs, and given an entity  $e_i$ , we denote by  $\tilde{\mathbf{h}}_{e_i}$  its representation generated by ER Encoder with the original graph as input, and  $\tilde{\mathbf{h}}_{e_i}^{aug}$  its representation generated with the augmented graph as input.

## 4.2 Model Training with Graph Augmentation

Graph augmentation learning has been demonstrated to promote the performance of graph learning, such as overcoming overfitting and oversmoothing issues [9], and being used for graph contrastive learning [25]. We apply graph augmentation for EA and highlight two main enhancements contributed by it: (I) injecting perturbations into the original KG can increase the diversity of the structural differences, thus preventing the model from overfitting to the training data during alignment process to some extent as well as enforcing the model to produce robust entity representations against structural changes; (II) graph augmentation inherently generates two graph views without extra parameters, which facilitates conducting contrastive learning to promote heterogeneous representation learning for (unseen) entities by contrasting different views.

**Margin-Based Alignment Loss.** In order to make equivalent entities close to each other and unmatched entities pull away from each other in a unified embedding space. Following previous works [5, 6, 20], we apply the margin-based alignment loss supervised by pre-aligned entity pairs  $S$ . Notably, here, we use the output of ER Encoder based on augmented graphs to make the model avoid overfitting and behave durable against edge changes:

$$\mathcal{L}_a = \sum_{(e_i, e_j) \in S} \sum_{(\bar{e}_i, \bar{e}_j) \in \bar{S}_{(e_i, e_j)}} \left[ \|\tilde{\mathbf{h}}_{e_i}^{aug} - \tilde{\mathbf{h}}_{e_j}^{aug}\|_{L2} + \rho - \|\tilde{\mathbf{h}}_{\bar{e}_i}^{aug} - \tilde{\mathbf{h}}_{\bar{e}_j}^{aug}\|_{L2} \right]_+, \quad (7)$$

where  $\rho$  is a hyper-parameter of margin,  $[x]_+ = \max\{0, x\}$  is to ensure non-negative output, and  $\bar{S}_{(e_i, e_j)}$  denotes the set of negative entity alignments constructed by corrupting the ground-truth alignment  $(e_i, e_j)$ , i.e., replacing  $e_i$  or  $e_j$  with another entity in  $G_s$  or  $G_t$  via negative sampling strategy.

**Contrastive Loss.** Contrastive learning is a good means to explore supervision signals from the vast unlabeled data. Many graph learning works [18, 19, 25] apply it to learn representations by contrasting different views and then maximizing feature consistency between them. RAC [27] is an effective EA model which incorporates contrastive learning to ameliorate the alignment performance. However, RAC needs to employ two separate graph encoders with the same architecture to model different views of the structural features of entities, which will bring twice the parameters and damage the diversity of graph views. Graph augmentation inherently provides two different views (i.e., original graph view and augmented graph view) without extra parameters. Therefore, we define the contrastive loss to improve entity representation learning by maximizing the feature consistency between the original structure and augmented structure:

$$\mathcal{L}_c = \sum_{z=\{s,t\}} \frac{1}{2|E_z|} \sum_{e_i \in E_z} (\mathcal{L}_{c, e_i}^{(G_z, G_z^{aug})} + \mathcal{L}_{c, e_i}^{(G_z^{aug}, G_z)}), \quad (8)$$

$$\mathcal{L}_{c, e_i}^{(G_z, G_z^{aug})} = -\log \frac{\exp(\langle \text{proj}(\tilde{\mathbf{h}}_{e_i}), \text{proj}(\tilde{\mathbf{h}}_{e_i}^{aug}) \rangle)}{\sum_{e_k \in E_z} \exp(\langle \text{proj}(\tilde{\mathbf{h}}_{e_i}), \text{proj}(\tilde{\mathbf{h}}_{e_k}^{aug}) \rangle)}, \quad (9)$$

where  $\langle \cdot \rangle$  means inner product, and  $\text{proj}(\cdot)$  is a shared projection head consisting of a linear layer and a ReLU activation function to map entity representations to low-dimensional vector space [25]. The definition of the symmetric contrastive loss term  $\mathcal{L}_{c,e_i}^{(G_z^{aug}, G_z^*)}$  is similar with Eq. (9).

**Model training.** We combine the margin-based alignment loss and the contrastive loss, arriving at the final objective of our model:

$$\mathcal{L} = \mathcal{L}_a + \lambda \mathcal{L}_c, \quad (10)$$

where  $\lambda \geq 0$  is a tunable parameter weighting the two objectives. The training process of GAEA is outlined in Algorithm 1, where negative sample set and augmented graphs will be updated every iteration (10 epochs as an iteration).

---

### Algorithm 1: Training Procedure of GAEA

---

**Input:** Knowledge graph  $G_s$  and  $G_t$ , pre-aligned entity pairs  $S$ .

```

1 Initialize entity embeddings and relation embeddings;
2 while Not Converge do
3   for each Epoch do
4     if Epoch % 10 == 0 then // 10 epochs as an iteration
5       Generate augmented graphs  $G_s^{aug}$  and  $G_t^{aug}$  for  $G_s$  and  $G_t$ ;
6       Generate negative sample set  $\tilde{S}$  based on  $S$ ;
7       Generate entity representations using ER Encoder;
8       Calculate  $\mathcal{L}_a$  using  $S$  and  $\tilde{S}$  via Eq.(7);
9       Calculate  $\mathcal{L}_c$  using Eq.(8) and Eq.(9);
10       $\Theta \leftarrow \text{BackProp}(\mathcal{L}_a + \lambda \mathcal{L}_c)$ ; ▷ Adam step
11 return Model parameters  $\Theta$ ;
```

---

### 4.3 Alignment Inference

After pulling embeddings from two KGs into a unified vector space and making them comparable, alignment relationships can be inferred by measuring the distance between two entities. In this work, we use Euclidean Distance to be the distance metric, i.e., for  $e_i \in E_s$  and  $e_j \in E_t$ , the distance between entity pair  $(e_i, e_j)$  is calculated by  $\|\tilde{\mathbf{h}}_{e_i} - \tilde{\mathbf{h}}_{e_j}\|_{L2}$ . In order to find  $e_i$ ' alignment relationship, we calculate its distance to all entities belonging to  $G_t$  and perform the nearest neighbor (NN) search to identify  $e_i$ ' counterpart entity in  $G_t$ :

$$e_j = \arg \min_{e'_j \in E_t} \|\tilde{\mathbf{h}}_{e_i} - \tilde{\mathbf{h}}_{e'_j}\|_{L2}. \quad (11)$$

Notably, we use the original KG structures in the inference phase instead of augmented versions to generate final entity representation  $\tilde{\mathbf{h}}$ . We apply Faiss<sup>1</sup> to accelerate the alignment inference process.

<sup>1</sup> <https://github.com/facebookresearch/faiss>.

## 5 Experimental Setup

### 5.1 Experimental Setup

**Datasets.** We use the 15K benchmark dataset (V1) in OpenEA [15] for evaluation since the entities thereof follow the degree distribution in real-world KGs. It contains two cross-lingual settings, i.e., EN-FR-15K (English-to-French) and EN-DE-15K (English-to-German), and two monolingual settings, i.e., D-W-15K (DBpedia-to-Wikidata) and D-Y-15K (DBpedia-to-YAGO). Following the data splits in OpenEA, we use the same split setting where 20%, 10%, and 70% alignments are harnessed for training, validation, and testing, respectively.

**Metrics.** We adopt  $Hits@k$  ( $k = 1, 5$ ) and *Mean Reciprocal Rank* ( $MRR$ ) as the evaluation metrics.  $Hits@k$  is to measure the alignment accuracy, while  $MRR$  measures the average performance of ranking over all test samples. The higher the  $Hits@k$  and  $MRR$ , the better the alignment performance.

**Baselines.** We choose some GNN variants and several existing state-of-the-art embedding-based EA models as baselines: GCN [4] and GAT [17] are the classic variants of GNNs; MTransE [2] and SEA [7] are triple-based methods that capture the local semantics information of relation triples via knowledge representation learning; GCN-Align [20], AliNet [14], HyperKA [11], and KE-GCN [26] are the neighborhood-based methods which apply GNNs to explore neighborhood structure information; IPTransE [30] and RSNs [3] both are path-based methods that extract the long-term dependencies across relation paths; IMEA [22] is the recent strong baseline which uses Transformer-like architecture to capture multiple structural contexts in an end-to-end manner.

We should note here that our model and the above baselines all mainly focus on the structural information of KGs. Therefore, for a fair comparison, we do not consider the models which utilize extra information (e.g., attributes, literals) for enhancement, such as AttrGNN [5], HMAN [24], MultiKE [29].

**Implementation Details.** All programs are implemented using *Python* 3.6.13 and *PyTorch* 1.10.2 with *CUDA* 11.3 on an *NVIDIA GeForce RTX 3090* GPU. Following OpenEA [15], we report the average results of five-fold cross-validation. We initialize trainable parameters with the Xavier initializer, and we train the model using Adam optimizer with weight decay  $1e-5$  and perform early stopping to terminate training based on the  $MRR$  score tested every 10 epochs on the validation data. As for hyper-parameters, the learning rate is set to 0.001, the dropout rate is 0.2, the layer number of GAT  $L$  is 2, the number of negative samples for each entity is 5, the negative sampling strategy is  $\epsilon$ -Truncated Uniform Negative Sampling [13] with  $\epsilon = 0.9$ , the margin  $\rho$  is 1, the balance parameter  $\lambda$  is 100, and the embedding dimension of entities  $d_{ent}$  and relations  $d_{rel}$  are set to 256 and 128, respectively. The  $pr$  is searched in  $\{0.05, 0.1, 0.15\}$ . Following the convention, the default alignment direction is from left to right. Taking D-W-15K as an example, we regard DBpedia as the source KG and seek to find the counterparts of source entities in the target KG Wikidata.



**Table 1.** Entity alignment results in cross-lingual and monolingual settings. The results with † are retrieved from [15], and ‡ from [22]. Results labeled by \* are reproduced using the released source codes. The **boldface** indicates the best result of each column and underlined the second-best.

Models	EN-FR-15K			EN-DE-15K			D-W-15K			D-Y-15K		
	Hit@1	Hit@5	MRR	Hit@1	Hit@5	MRR	Hit@1	Hit@5	MRR	Hit@1	Hit@5	MRR
GCN*	.210	.414	.304	.304	.497	.394	.208	.367	.284	.343	.503	.416
GAT*	.297	.585	.426	.542	.737	.630	.383	.622	.489	.468	.707	.573
MTransE <sup>†</sup>	.247	.467	.351	.307	.518	.407	.259	.461	.354	.463	.675	.559
SEA <sup>†</sup>	.280	.530	.397	.530	.718	.617	.360	.572	.458	.500	.706	.591
IPTransE <sup>†</sup>	.169	.320	.243	.350	.515	.430	.232	.380	.303	.313	.456	.378
RSNs <sup>†</sup>	.393	.595	.487	.587	.752	.662	.441	.615	.521	.514	.655	.580
GCN-Align <sup>†</sup>	.338	.589	.451	.481	.679	.571	.364	.580	.461	.465	.626	.536
AliNet <sup>‡</sup>	.364	.597	.467	.604	.759	.673	.440	.628	.522	.559	.690	.617
HyperKA <sup>‡</sup>	.353	.630	.477	.560	.780	.656	.440	.686	.548	.568	.777	.659
KE-GCN <sup>‡</sup>	.408	.670	.524	<u>.658</u>	.822	<u>.730</u>	.519	.727	.608	.560	.750	.644
IMEA <sup>‡</sup>	<u>.458</u>	<u>.720</u>	<u>.574</u>	.639	<u>.827</u>	.724	<u>.527</u>	<u>.753</u>	<u>.626</u>	<b>.639</b>	<b>.804</b>	<b>.712</b>
GAEA	<b>.486</b>	<b>.746</b>	<b>.602</b>	<b>.684</b>	<b>.854</b>	<b>.760</b>	<b>.562</b>	<b>.768</b>	<b>.654</b>	<u>.608</u>	<u>.791</u>	<u>.688</u>
w/o <i>rel.</i>	.324	.626	.458	.593	.785	.678	.409	.666	.521	.502	.743	.605

## 5.2 Experimental Results

**Performance Comparison.** Table 1 reports the comparison results on the OpenEA 15K datasets. Experimental results show that our proposed GAEA outperforms other models in most tasks, especially in cross-lingual settings. There is a phenomenon that the performance of models utilizing knowledge representation learning as the encoder, e.g., MTransE, SEA, and IPTransE, are inferior compared with the models applying GNNs as the encoder like AliNet and KE-GCN, and have on-par or even worse performance than vanilla GCN and GAT, which demonstrates the GNNs’ powerful representation ability in EA. We also notice that, compared with some methods applying GCN as the encoder (e.g., GCN-Align, AliNet), the vanilla GCN fails to surpass them, which shows the significance of designing a more effective encoder for representing entities in KGs. IMEA is a strong baseline that captures abundant structure contexts and it obtains excellent results on D-Y-15K task. However, IMEA introduces carefully designed data processing (e.g., entity paths encoding) and becomes a complicated network due to the Transformer-like architecture, which will inevitably increase the training difficulty and overfitting risk. Additionally, we compare the model size (denoted as #Params) in Table 2. GAEA greatly reduces the number of parameters compared to IMEA while acquiring decent alignment performance. This is because GAEA designs a simple Entity-Relation Encoder to capture multi-range neighborhood structures to mitigate heterogeneity and infuse relation semantics to provide more comprehensive signals for alignment. Moreover, GAEA further facilitates producing expressive and robust entity representations by integrating graph augmentation to achieve alignment learning supervised by alignment seeds and contrastive representation learning for unseen entities. In summary, our proposed GAEA is a light and powerful solution for EA.

**Table 2.** #Params comparison.


Models	#Params (M)
GCN	~7.81M
AliNet	~16.18M
IMEA	~20.44M
GAEA (ours)	~8.10M

**Table 3.** Ablation study results.

Models	EN-DE-15K			D-W-15K		
	Hit@1	Hit@5	MRR	Hit@1	Hit@5	MRR
GAEA	<b>.684</b>	<b>.854</b>	<b>.760</b>	<b>.562</b>	<b>.768</b>	<b>.654</b>
<i>-gaal.</i>	.674	.848	.751	.557	.764	.650
$-\mathcal{L}_c$	.665	.841	.744	.544	.755	.639

**Ablation Study.** In the above experiments, the overall effectiveness of GAEA is proved. In this section, we conduct ablation analyses to demonstrate the validity of each component of GAEA. First, Table 1 also gives the results of a variant of GAEA (denoted as *w/o rel.*), which means the original GAEA eliminates relation injection. The ablation results clearly show the effectiveness of relation embedding learning, which identifies the relation semantics can help in enriching the expressiveness of entity representations. Next, Table 3 gives the ablation results about graph augmentation. *-gaal.* and  $-\mathcal{L}_c$  represent the variants by removing graph augmentation in alignment learning (i.e., Eq. (7)) or removing contrastive objective (i.e., Eq. (8)), respectively (the results of removing graph augmentation are illustrated in the next section). The results show that utilizing graph augmentation can have positive impacts on EA and consistently get better performance. By introducing graph augmentation into EA training process, the model not only is encouraged to learn useful and robust entity representations but also lets the scarce yet valuable alignment seeds and vast unlabeled entities in KGs jointly provide abundant supervision for model learning.

**Parameter Analysis.** Considering that our model employs edge dropping to generate augmented graphs for margin-based alignment learning and contrastive entity representation learning. We investigate how the alignment performance varies with the upper bound of the deletion ratio. We evaluate upper bound  $pr$  in  $\{0, 0.05, 0.1, 0.15\}$ , and the results measured by *Hit@1* and *MRR* are drawn in Fig. 2. The performance is worst on all three tasks when  $pr = 0$ , i.e., without any graph augmentation enhancement, indicating that graph augmentation can



**Fig. 2.** Parameter analysis results of  $pr$  measured by Hit@1 (green bar with left axis) and MRR (yellow bar with right axis). (Color figure online)

do benefit for alignment learning. We can see that the alignment effect is best when  $pr$  equals 0.05 or 0.1, increasing  $pr$  to 0.15 will not further improve the performance, and even bring performance drops. One potential reason is that when  $pr$  becomes large, edge dropping will lead to losing more semantic knowledge and structural information, thus bringing an adverse impact on neighborhood aggregation and model training. Therefore, we need to set  $pr$  as a suitably small value to ensure information retention as well as performance improvement.

## 6 Discussion and Conclusion

In this paper, we propose GAEA, a novel entity alignment method based on graph augmentation. Specifically, we design an Entity-Relation (ER) Encoder to generate latent representations for entities via jointly capturing neighborhood structures and relation semantics. Meanwhile, we apply graph augmentation to create two graph views for margin-based alignment learning and contrastive entity representation learning, thus improving the model’s alignment performance. Finally, experimental results verified the effectiveness of our method.

Although GAEA achieves promising results, it still has limitations that need further investigation. First, our experimental results show that graph augmentation learning can bring some performance gains, but the supervision signals provide key performance bases in the alignment learning process. Thus, it is worth further studying how to amplify the improvement brought by graph augmentation when there no alignment seeds are given. Besides, we currently apply edge dropping as the only graph augmentation strategy, which exposes a new problem, that is, how to conduct graph augmentation learning in a highly structured KG to improve performance without introducing logic errors.

**Acknowledgments.** We thank reviewers for their helpful feedback. This work is supported by the National Natural Science Foundation of China No. 62172428.

## References

1. Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., Yakhnenko, O.: Translating embeddings for modeling multi-relational data. In: Proceedings of NIPS (2013)
2. Chen, M., Tian, Y., Yang, M., Zaniolo, C.: Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. In: Proceedings of IJCAI (2016)
3. Guo, L., Sun, Z., Hu, W.: Learning to exploit long-term relational dependencies in knowledge graphs. In: Proceedings of ICML (2019)
4. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint [arXiv:1609.02907](https://arxiv.org/abs/1609.02907) (2016)
5. Liu, Z., Cao, Y., Pan, L., Li, J., Chua, T.S.: Exploring and evaluating attributes, values, and structures for entity alignment. In: Proceedings of EMNLP (2020)
6. Mao, X., Wang, W., Xu, H., Lan, M., Wu, Y.: Mraea: an efficient and robust entity alignment approach for cross-lingual knowledge graph. In: Proceedings of WSDM (2020)

7. Pei, S., Yu, L., Hoehndorf, R., Zhang, X.: Semi-supervised entity alignment via knowledge graph embedding with awareness of degree difference. In: Proceedings of WWW (2019)
8. Pei, S., Yu, L., Yu, G., Zhang, X.: Graph alignment with noisy supervision. In: Proceedings of WWW (2022)
9. Rong, Y., Huang, W., Xu, T., Huang, J.: DropEdge: towards deep graph convolutional networks on node classification. arXiv preprint [arXiv:1907.10903](https://arxiv.org/abs/1907.10903) (2019)
10. Srivastava, R.K., Greff, K., Schmidhuber, J.: Highway networks. arXiv preprint [arXiv:1505.00387](https://arxiv.org/abs/1505.00387) (2015)
11. Sun, Z., Chen, M., Hu, W., Wang, C., Dai, J., Zhang, W.: Knowledge association with hyperbolic knowledge graph embeddings. In: Proceedings of EMNLP (2020)
12. Sun, Z., Hu, W., Wang, C., Wang, Y., Qu, Y.: Revisiting embedding-based entity alignment: a robust and adaptive method. IEEE TKDE (2022)
13. Sun, Z., Hu, W., Zhang, Q., Qu, Y.: Bootstrapping entity alignment with knowledge graph embedding. In: Proceedings of IJCAI (2018)
14. Sun, Z., et al.: Knowledge graph alignment network with gated multi-hop neighborhood aggregation. In: Proceedings of AAAI (2020)
15. Sun, Z., et al.: A benchmarking study of embedding-based entity alignment for knowledge graphs. arXiv preprint [arXiv:2003.07743](https://arxiv.org/abs/2003.07743) (2020)
16. Vaswani, A., et al.: Attention is all you need. In: Proceedings of NIPS (2017)
17. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint [arXiv:1710.10903](https://arxiv.org/abs/1710.10903) (2017)
18. Veličković, P., Fedus, W., Hamilton, W.L., Liò, P., Bengio, Y., Hjelm, R.D.: Deep graph infomax. In: Proceedings of ICLR (2018)
19. Wan, S., Pan, S., Yang, J., Gong, C.: Contrastive and generative graph convolutional networks for graph-based semi-supervised learning. In: Proceedings of AAAI (2021)
20. Wang, Z., Lv, Q., Lan, X., Zhang, Y.: Cross-lingual knowledge graph alignment via graph convolutional networks. In: Proceedings of EMNLP (2018)
21. Wu, Y., Liu, X., Feng, Y., Wang, Z., Yan, R., Zhao, D.: Relation-aware entity alignment for heterogeneous knowledge graphs. In: Proceedings of IJCAI (2019)
22. Xin, K., Sun, Z., Hua, W., Hu, W., Zhou, X.: Informed multi-context entity alignment. In: Proceedings of WSDM (2022)
23. Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K.I., Jegelka, S.: Representation learning on graphs with jumping knowledge networks. In: ICML (2018)
24. Yang, H.W., Zou, Y., Shi, P., Lu, W., Lin, J., Sun, X.: Aligning cross-lingual entities with multi-aspect information. In: Proceedings of EMNLP (2019)
25. You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., Shen, Y.: Graph contrastive learning with augmentations. In: Proceedings of NIPS (2020)
26. Yu, D., Yang, Y., Zhang, R., Wu, Y.: Knowledge embedding based graph convolutional network. In: Proceedings of WWW (2021)
27. Zeng, W., Zhao, X., Tang, J., Fan, C.: Reinforced active entity alignment. In: Proceedings of CIKM (2021)
28. Zeng, W., Zhao, X., Tang, J., Lin, X.: Collective entity alignment via adaptive features. In: Proceedings of ICDE (2020)
29. Zhang, Q., Sun, Z., Hu, W., Chen, M., Guo, L., Qu, Y.: Multi-view knowledge graph embedding for entity alignment. In: Proceedings of IJCAI (2019)
30. Zhu, H., Xie, R., Liu, Z., Sun, M.: Iterative entity alignment via knowledge embeddings. In: Proceedings of IJCAI (2017)