Degree-Aware Alignment for Entities in Tail

Weixin Zeng
National University of Defense
Technology
zengweixin13@nudt.edu.cn

Xiang Zhao*
National University of Defense
Technology
xiangzhao@nudt.edu.cn

Wei Wang
The University of New South Wales
weiw@cse.unsw.edu.au

Jiuyang Tang
National University of Defense
Technology
jiuyang_tang@nudt.edu.cn

Zhen Tan
National University of Defense
Technology
tanzhen08a@nudt.edu.cn

ABSTRACT

Entity alignment (EA) is to discover equivalent entities in knowledge graphs (KGs), which bridges heterogeneous sources of information and facilitates the integration of knowledge. Existing EA solutions mainly rely on structural information to align entities, typically through KG embedding. Nonetheless, in real-life KGs, only a few entities are densely connected to others, and the rest majority possess rather sparse neighborhood structure. We refer to the latter as *long-tail entities*, and observe that such phenomenon arguably limits the use of structural information for EA.

To mitigate the issue, we revisit and investigate into the conventional EA pipeline in pursuit of elegant performance. For pre-alignment, we propose to amplify long-tail entities, which are of relatively weak structural information, with *entity name* information that is generally available (but overlooked) in the form of concatenated power mean word embeddings. For alignment, under a novel complementary framework of consolidating structural and name signals, we identify entity's *degree* as important guidance to effectively fuse two different sources of information. To this end, a degree-aware *co-attention* network is conceived, which dynamically adjusts the significance of features in a degree-aware manner. For post-alignment, we propose to complement original KGs with facts from their counterparts by using confident EA results as anchors via *iterative* training. Comprehensive experimental evaluations validate the superiority of our proposed techniques.

KEYWORDS

Entity alignment; Long-tail; Co-attention; Iterative training

---

*Corresponding author.

---

1 INTRODUCTION

Over recent years, a large number of Knowledge Graphs (KGs), e.g., YAGO [18], DBpedia [1], Knowledge Vault [8] and NELL [4], have been constructed, which contribute significantly to the development of intelligent information services. In addition, there is a growing number of domain-specific KGs, such as medical KG [1] and scientific KG [2]. Nevertheless, a single KG can never reach perfect coverage or being 100% correct, due to the inevitable trade-off that the KG construction process needs to take between *coverage* and *correctness* [13].

An effective way to automatically and efficiently increase the coverage and correctness of KGs is by integrating knowledge from other KGs. This is because KGs constructed independently and/or from independent sources generally provide redundancy and complementary information, so that the integrated KG is expected to be better in both coverage and correctness. For instance, a general KG constructed from Web pages might only contain brand names of medicine, while more details of it (e.g., generic names of medicine) can be found in a medical KG constructed from medical literature. To incorporate knowledge from external KGs into the original KG, the first and the most crucial step, is to align KGs. As such, recent efforts have been devoted to entity alignment (EA) [22, 26, 27], and these aligned entities serve as pivots to connect KGs and lay foundation for downstream tasks.

Current solutions to EA mainly rely on graph structure of KGs [6, 10, 14, 20, 23], which take for granted that equivalent entities possess similar neighborhood structures. On some synthetic datasets extracted from large-scale KGs, these methods have achieved state-of-the-art performance [3, 20, 29]. However, recent study pointed out that those *synthetic* datasets are much denser than KGs in real life, and existing EA methods are *not* able to yield satisfactory results on datasets with real-life distributions [10].

In particular, Guo et al. [10] revealed that nearly half of the entities in real-life KGs are connected to less than three other entities. We refer to them as *long-tail entities*, and these entities in tail make the KG a relatively sparse graph. This follows our perception of real-life KGs that only very few entities are frequently accessed, and possess detailed attributes and rich connections, whereas the majority are left under-explored and come with little structural information. As a consequence, existing EA methods that solely

---

[1] https://flowhealth.com/
[2] https://www.aminer.cn/scikg

rely on structural information tend to take a toll on their capability of accurate alignment, as demonstrated by Example 1.

Example 1. *In Figure 1 is a partial English KG (KG$_{EN}$) and a partial Spanish KG (KG$_{ES}$) concerning the film* Summer 1993. *Note that the entities* The Bookshop *and* La Librería *in grey describe the original novel, while those in white depict the film.*

During aligning entities of high degrees, e.g., Spain *and* España, *structural information is of great help; however, as to long-tail entities, e.g.,* Carla Simón *in KG$_{EN}$, structural information may suggest* Laia Artigas *in KG$_{ES}$ as its match, since they have a single link to* Summer 1993 *and* Verano 1993, *respectively.*

The example unveils the shortcoming of solely relying on structural information for EA, which renders existing EA methods suboptimal, even infeasible for long-tail entities. Hence, we are motivated to revisit the key phases of EA pipeline, and address the challenge of EA when structure information is insufficient.

For *pre-alignment* phase, we seek additional signals that can benefit EA, and discover a source of information from entity names. It is generally available among real-life entities, yet has been overlooked by existing research. For instance, for the long-tail entity Carla Simón in KG$_{EN}$, introducing entity name information would easily help locate the correct mapping—Carla Simón in KG$_{ES}$. Thus, entity name information tends to serve as a *complementing* view to the commonly used structural information. In fact, name embeddings (in the form of averaged word embedding) have been incorporated before [24–26], which were plainly used to fill up the initial feature matrices for learning *structural representation*; in contrast, we explicitly harness entity names as another channel of signal, parallel to structure information [27], by encoding via concatenated power mean word embeddings.

For *alignment* phase, we explore to judiciously fuse the two aforementioned signals, based on the observation that for entities of different degrees, the importance of structural and name information varies. In Example 1, when aligning long-tail entity Carla Simón in KG$_{EN}$, its limited neighboring structure is less useful than entity name information; distinctively, structure plays a more significant role for mapping popular entities like the film La Librería, typically in the presence of ambiguous entity names (i.e., both the film La Librería and the novel La Librería share the same name). In general, a plausible intuition is that *the lower (resp. higher) the degree, the more important the signal from entity name (resp. neighboring structure).* To precisely capture the non-linear dynamics between these two signals, we devise a co-attention network to determine the weights of different signals under the guidance of entity degrees. Noteworthily, Pei et al. [14] introduced degrees with the intention of correcting the undesirable bias from *structural embedding* methods that place entities with similar degrees closely; our motivation is different in that degrees are leveraged to compute pair-wise similarities (rather than individual embeddings).

For *post-alignment* phase, we propose to substantially enhance structural information of KGs by looking at and referencing each other recurrently. Long-tail entities fall short of structural information in their original KG (source KG), but the KG to be aligned (target KG) may possess such information complementarily. For example, for the entity Carla Simón, the fact that Carla Simón is from España is missing in KG$_{EN}$, which is yet seen in KG$_{ES}$. If the source KG can harvest it from the counterpart of the target KG (after pairing the surrounding entities), alignment could be ameliorated. Inspired by the positive effect of initial completion of KGs using rules [3], we conceive an *iterative* training procedure with KG completion embedded, which takes confident EA results in each round as anchors and replenishes possible missing relations to enhance the current KGs. As a net effect, these KGs are enriched, from which better structural embeddings could be learned, and the matching signal could propagate to long-tail entities—previously difficult in a single shot but now may become easier to align.

Contribution. In short, our contribution can be summarized as

- We identify the deficiency of existing EA methods in aligning long-tail entities, largely due to sole reliance on structure. We approach the limit by (1) introducing a complementary signal from entity names in the form of concatenated power mean word embeddings; and (2) conceiving an effective way via degree-aware co-attention mechanism to dynamically fuse name and structural signals.
- We propose to reduce long-tail entities through augmenting relational structure via KG completion embedded into an iterative self-training strategy, which is realized by taking confident EA results as anchors and using each other KG as references. The strategy enhances the performance of EA and the coverage of KGs simultaneously.
- The techniques constitute a novel framework, namely, DAT. We empirically evaluate the implementation of DAT on both mono-lingual and cross-lingual EA tasks against state-of-the-art methods, and the comparative results and ablation analysis demonstrate the superiority of DAT.

Organization. Section 2 overviews related work. In Section 3, we analyse the long-tail phenomenon in EA. DAT and its components are elaborated in Section 4. Section 5 introduces experimental settings, evaluation results and detailed analysis, followed by conclusion in Section 6.
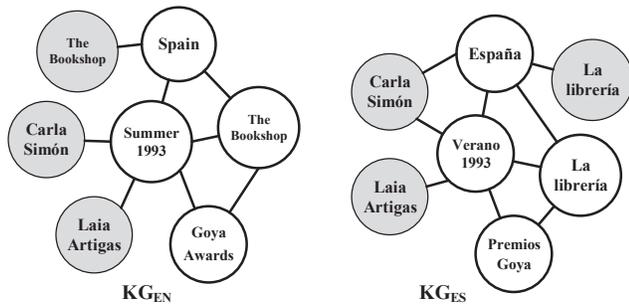

Figure 1: An example of EA. Nodes in grey (resp. white) are long-tail (resp. popular) entities (relation names and other entities are omitted in the interest of space).

2 RELATED WORK

Conventional EA framework. The contributions of state-of-the-art methods can be described from the angle of a phased pipeline for EA. For pre-alignment, one utilizes KG representation methods,

Table 1: Degree distribution of entities in test set (the first KG in each KG pair) and results of RSNs

Degree	EN-FR			EN-DE			DBP-WD			DBP-YG		
	#Total	#Correct	Accuracy	#Total	#Correct	Accuracy	#Total	#Correct	Accuracy	#Total	#Correct	Accuracy
1	2,660	380	14.29%	1,978	453	22.90%	1,341	276	20.58%	3,327	538	16.17%
2	2,540	699	27.52%	2,504	1,005	40.14%	2,979	800	26.85%	2,187	688	31.46%
3	1,130	408	36.11%	1,514	820	54.16%	1,789	600	33.54%	1,143	563	49.26%
$\geq$=4	3,120	1,803	57.79%	3,454	2,416	69.95%	3,341	2,093	62.65%	2,793	2,008	71.89%
All	9,450	3,290	34.81%	9,450	4,694	49.67%	9,450	3,769	39.88%	9,450	3,797	40.18%

e.g., TransE [5, 6, 28] and GCN [23], to encode structural information and embed KGs into individual low-dimensional spaces. Then for alignment, under the supervision of seed entity pairs, alignment results can be derived by evaluating and comparing the embedding spaces. Some methods [10, 19, 20] directly project entities in different KGs into the same embedding space by fusing the training corpus in pre-alignment; then for alignment, according to some distance in the unified embedding space, equivalence in different KGs can also be found. For post-alignment, iterative strategies [20, 28] are employed to augment supervision signals by exploiting the results of alignment phase. In this way, structural embeddings are updated, and alignment can be recursively executed until reaching some stopping criterion. These techniques can be roughly summarized into a framework, depicted by Figure 2.

Recent advancement on EA. There are some very recent efforts aiming to overcome structural heterogeneity by devising more advanced structure learning models, e.g., topic graph matching [26] and multi-channel graph neural network [3]. Particularly, Pei et al. [14] improved the structural embedding with awareness of the degree difference by performing adversarial training. Albeit, it still relies on structure for aligning signals, which tends to be ineffective when entities in both KGs are in tail; and moreover, degree information here is used to better learn structural embeddings, whereas our use of degrees is to better fuse two different and useful aligning signals—structural and name information.

Iterative strategies are beneficial to improving EA, but can be either time-consuming and biased towards one KG [20], or susceptible to introducing many false positive instances [28], which can barely satisfy the real-life requirement. To strike a balance between precision and computational overhead, we implement the idea of iterative training with the inclusion of a KG completion module, such that structure is actually updated in each round around confident anchoring entity pairs. The strategy is of light weight and keeps the possible inclusion of incorrect pairs at a small number.
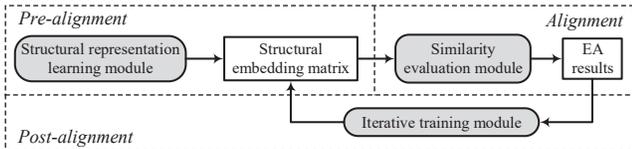

Figure 2: Conventional framework of EA

It can be seen that almost all the aforementioned embeddings learn from structural information, which can be sometimes insufficient, especially for long-tail entities. In this connection, some proposed to incorporate *attributes* [19, 22, 23, 27] to potentially compensate the shortage. Nevertheless, between 69% and 99% of instances in popular KGs lack at least one attribute that other entities in the same class have [9]. Similarly, *entity descriptions* [5] could be utilized to provide extra information, which is nonetheless often missing in many KGs. In short, these efforts possibly strengthen the overall EA performance but tend to fall short in aligning entities in tail. Entity name has also been tried, either as initial features for learning *structural representation* [24–26], or together with other information for representation learning [27]. In contrast, we consolidate the features on top of separate similarity matrices learned from structure and name information; empirical evaluation of different strategies will be described in Section 5.2.

3 IMPACT OF LONG-TAIL PHENOMENON

Task definition. Given a source KG $G_1 = (E_1, R_1, T_1)$ and a target KG $G_2 = (E_2, R_2, T_2)$, where E_1 (resp. E_2) represents source (resp. target) entities, R denotes relations, $T \subseteq E \times R \times E$ represents triples. Denote the seed entity pairs as $S = \{(e_1^i, e_2^i)|e_1^i = e_2^i, e_1^i \in E_1, e_2^i \in E_2\}$, $i \in [1, |S|]$, where $|\cdot|$ denote the cardinality of a set. EA task is to find new EA pairs based on S and return the eventual results $S' = \{(e_1^i, e_2^i)|e_1^i = e_2^i, e_1^i \in E_1, e_2^i \in E_2\}$, $i \in [1, \min\{|E_1|, |E_2|\}]$, where = expresses that two entities are the same physical one.

A very recent work [10] points out that KGs in previous EA datasets are too dense and the degree distributions deviate from real-life KGs. Therefore, they establish a new EA benchmark that follows real-life distribution. The evaluation benchmark consists of both cross-lingual datasets, i.e., EN-FR, EN-DE, and mono-lingual datasets, i.e., DBP-WD, and DBP-YG. We show the degree distributions of entities in test sets in Table 1. The degree of an entity is defined as the number of relational triples in which the entity participates. We also report the amount of correctly aligned entities in different degrees yielded by RSNs, the best solution in [10].

It reads from Table 1 that in EN-FR and DBP-YG, entities with degree less than three account for over 50%, and nearly half of the entities in EN-DE and DBP-WD are with degree 1 or 2. It confirms that the majority of the entities in KG are long-tail and have very few connections to others. Moreover, it follows that the performance of long-tail entities is much worse than those of higher degrees (the accuracies of high-degree entities triple and even quadruple those of entities with degree 1), despite that RSNs is the leading method on the benchmark. This suggests that current methods fail
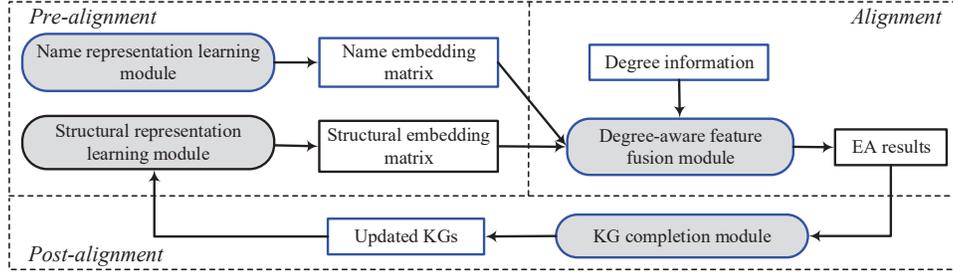
Figure 3: The framework of DAT

4 METHODOLOGY

To provide an overview, we summarize the main components of the framework DAT (degree-aware entity alignment in tail) in Figure 3, where purple blue highlights the new designs in this research. In pre-alignment, *Structural representation learning module* and *Name representation learning module* learn useful features of entities, i.e., name representation and structural representation; in alignment, these features are forwarded to *Degree-aware fusion module* for effective fusion and alignment under the guide of degree information. In post-alignment, *KG completion module* aims to complete KGs with confident EA pairs in the results, and the augmented KGs are then again utilized in the next round iteratively.

Since *Structural representation learning module* has been extensively studied, we adopt the state-of-the-art model RSNs [10] for this purpose. Given a structural embedding matrix $\mathbf{Z} \in \mathbb{R}^{n \times d_s}$, two entities $e_1 \in G_1$ and $e_2 \in G_2$, their structural similarity $Sim_s(e_1, e_2)$ is the cosine similarity between $\mathbf{Z}(e_1)$ and $\mathbf{Z}(e_2)$, where n denotes the number of all entities in two KGs, d_s is the dimension of structural embeddings, and $\mathbf{Z}(e)$ denotes the embedding vector for entity e (i.e., $\mathbf{Z}(e) = \mathbf{Z}\mathbf{e}$, where $\mathbf{e}$ is the one-hot encoding of entity e). From the perspective of structure, the target entity with the highest similarity to a source entity is returned as its alignment result.

4.1 Name representation learning

Recall that structural information is of limited use for aligning long-tail entities. Distinct from existing efforts that try to exploit structures, we take another angle by seeking some signal that is beneficial and generally available to long-tail entities.

To this end, we propose to incorporate the textual names of entities, a signal that has been largely overlooked by current embedding-based EA methods. In particular, the choice is alluring due to at least the following considerations: (1) entity name can normally identify an entity, and given two entities, comparing their names may be the most intuitive way to judge the equivalence preliminarily; (2) most real-life entities possess a name, the ratio of which is much higher than that of other textual information (i.e., descriptions and attributes), which tends to be lacking for long-tail entities.

Despite that there are many classic approaches for measuring the *string similarity* between entity names, we go for *semantic similarity*

since it can still work when the vocabularies of KGs differ, especially for the *cross-lingual* scenario. Specifically, we choose a general form of power mean embeddings [15], which encompasses many well-known means such as the arithmetic mean, the geometric mean and the harmonic mean. Given a sequence of word embeddings, $\mathbf{w}_1, \ldots, \mathbf{w}_l \in \mathbb{R}^d$, the power mean operation is formalized as

$$(\frac{w_{1i}^p + \cdots + w_{li}^p}{l})^{1/p}, \quad \forall i = 1, \ldots, d, \quad p \in \mathbb{R} \cup \pm\infty, \quad (1)$$

where l is the number of words and d denotes the dimension of embeddings. It can be seen that setting p to 1 results in the arithmetic mean, to 0 the geometric mean, to -1 the harmonic mean, to $+\infty$ the maximum operation and to $-\infty$ the minimum operation [16].

Given a word embedding space $\mathbb{E}^i$, the embeddings of the words in the name of entity s can be represented as $\mathbf{W}^i = [\mathbf{w}_1^i, \ldots, \mathbf{w}_l^i] \in \mathbb{R}^{l \times d^i}$. Correspondingly, $H_p(\mathbf{W}^i) \in \mathbb{R}^{d^i}$ denotes the power mean embedding vector after feeding $\mathbf{w}_1^i, \ldots, \mathbf{w}_l^i$ to Equation (1). To obtain summary statistics of entity s, we compute K power means of s and concatenate them to get the entity name representation $\mathbf{s}^i \in \mathbb{R}^{d^i \cdot K}$, i.e.,

$$\mathbf{s}^i = H_{p_1}(\mathbf{W}^i) \oplus \cdots \oplus H_{p_K}(\mathbf{W}^i), \quad (2)$$

where $\oplus$ represents concatenation along rows, and $p_1, \ldots, p_K$ are K different power mean values [16].

To get further representational power from different word embeddings, we generate the final entity name representation $\mathbf{n}_s$ by concatenating $\mathbf{s}^i$ obtained from different embedding spaces $\mathbb{E}^i$:

$$\mathbf{n}_s = \bigoplus_i \mathbf{s}^i. \quad (3)$$

Note that the dimensionality of this representation is $d_n = \sum_i d^i \cdot K$. The name embeddings of all entities can be denoted in matrix form as $\mathbf{N} \in \mathbb{R}^{n \times d_n}$.

Like word embeddings, semantically similar entity names will be placed adjacently in the entity name representation space. From the perspective of textual name, given two entities $e_1 \in G_1$ and $e_2 \in G_2$, their similarity $Sim_t(e_1, e_2)$ is the cosine similarity between $\mathbf{N}(e_1)$ and $\mathbf{N}(e_2)$, and the target entity with the highest similarity to a source entity is returned as its alignment result.

Discussion. The concatenated power mean word embedding [16] is a better solution for representing entity name than averaged word embedding in that it can better summarize the useful signals

in entity name [3]. The average of word embeddings discards large amount of information since variant names can be represented by similar averaged embedding. In contrast, the concatenation of different power means yields a more precise summary by reducing uncertainty about the semantic variation within a sentence. This is further verified through empirical results in Section 5.3.

Note that for cross-lingual EA, we utilize pre-trained multilingual word embeddings [7] that have already aligned words from different languages into the same embedding space. This ensures that entity names from various language sources also exist in the same semantic space, which avoids designing an additional mapping function for aligning multilingual embedding spaces.

Among others, the aforementioned method can be extended to support other textual information like attributes, without loss of generality. An immediate solution is to concatenate attributes and entity name to form a "sentence" that describes the entity, which is then encoded by concatenated power mean word embeddings. Nevertheless, the extension to incorporate additional information and more advanced adaptations are beyond the scope of this paper.

4.2 Degree-aware co-attention feature fusion

Different types of features characterize entity identities from different aspects. Thereby, a feature fusion module is of great significance to the effective combination of various signals. Some propose to combine different embeddings into a unified representation space [27], whereas it requires additional training for unifying irrelevant features. A more preferable strategy would be computing the similarity matrix within each feature-specific space first, and then combine the feature-specific similarity scores [12, 23]. Nevertheless, for entities with different degrees, the contributions of different features vary. Regarding long-tail entities that possess little structural information, entity name representation should be credited more; conversely, structural representation of popular entities is relatively more useful than entity name information. To capture this dynamic change, inspired by the bi-attention mechanism [17], we devise a degree-aware co-attention network, shown in Figure 4.

Formally, given the structural embedding matrix $\mathbf{Z}$, name embedding matrix $\mathbf{N}$, for every entity pair, $(e_1, e_2), e_1 \in G_1, e_2 \in G_2$, the similarity score between e_1 and e_2 is calculated, which is then utilized for determining the alignment result. In order to compute the overall similarity, we first calculate the feature-specific similarity scores between e_1 and e_2, i.e., $Sim_s(e_1, e_2)$ and $Sim_t(e_1, e_2)$, as introduced in previous subsections. Then our degree-aware co-attention network aims to determine the weights for $Sim_s(e_1, e_2)$ and $Sim_t(e_1, e_2)$ by integrating degree information. This network comprises three stages—feature matrix construction, co-attention similarity matrix calculation, and weight assignment.

Feature matrix construction. In addition to entity name and structural information, we incorporate entity degree information to construct feature matrix for each entity. In specific, entity degrees are represented as one-hot vectors of all possible degree values, and then forwarded to a fully-connected layer to generate the continuous degree vector. For instance, the degree vector of e_1 is
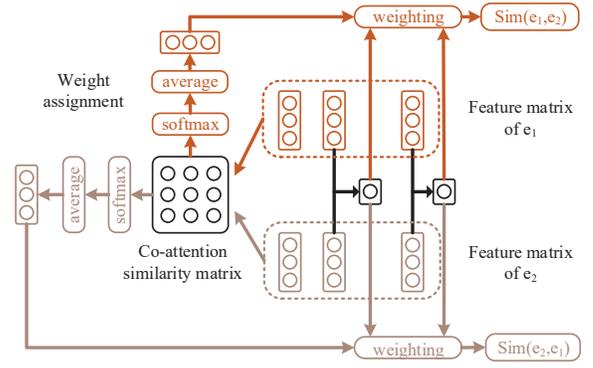

Figure 4: Degree-aware co-attention feature fusion

$\mathbf{g}_{e_1} = \mathbf{M} \cdot \mathbf{h}_{e_1} \in \mathbb{R}^{d_g}$, where $\mathbf{h}_{e_1}$ is its one-hot representation, $\mathbf{M}$ is the weight matrix in the fully-connected layer and d_g represents the dimension of degree vector. This continuous degree vector, along with structural and entity name representations, is stacked to form an entity's feature matrix. For entity e_1,

$$\mathbf{F}_{e_1} = [\mathbf{N}(e_1); \mathbf{Z}(e_1); \mathbf{g}_{e_1}] \in \mathbb{R}^{3 \times d_m}, \quad (4)$$

where ; denotes the concatenation along *columns*, $d_m = \max\{d_n, d_s, d_g\}$ and we pad the missing values with 0s.

Co-attention similarity matrix calculation. To model the interaction between $\mathbf{F}_{e_1}$ and $\mathbf{F}_{e_2}$, as well as highlight important features, we build a co-attention matrix $\mathbf{S} \in \mathbb{R}^{3 \times 3}$, where the similarity between the i-th feature of e_1 and the j-th feature of e_2 is computed by

$$\mathbf{S}_{ij} = \alpha(\mathbf{F}_{e_1}^{i:}, \mathbf{F}_{e_2}^{j:}) \in \mathbb{R}, \quad (5)$$

where $\mathbf{F}_{e_1}^{i:}$ is the i-th row vector and $\mathbf{F}_{e_2}^{j:}$ is the j-th row vector, $i = 1, 2, 3; j = 1, 2, 3$. $\alpha(\mathbf{u}, \mathbf{v}) = \mathbf{w}^\top(\mathbf{u} \oplus \mathbf{v} \oplus (\mathbf{u} \circ \mathbf{v}))$ is a trainable scalar function that encodes the similarity, where $\mathbf{w} \in \mathbb{R}^{3d_m}$ is a trainable weight vector, $\circ$ is the element-wise multiplication. Note that the implicit multiplication is a matrix multiplication.

Weight assignment. The co-attention similarity matrix $\mathbf{S}$ is then utilized to obtain attention vectors, i.e., $\mathbf{att_1}$ and $\mathbf{att_2}$, in both directions. $\mathbf{att_1}$ signifies which feature vectors of e_1 are most relevant to feature vectors of e_2, while $\mathbf{att_2}$ signifies which feature vectors of e_2 are most relevant to feature vectors of e_1. In specific, we feed $\mathbf{S}$ to a softmax layer. The resulting matrix is then compressed by an average layer to generate attention vectors. Note that column-wise operation in the softmax layer and row-wise operation in the average layer produce $\mathbf{att_1}$, while row-wise operation in the softmax layer and column-wise operation in the average layer produce $\mathbf{att_2}$.

Eventually, we multiply the feature-specific similarity scores with the attention values to obtain the final similarity score,

$$Sim(e_1, e_2) = Sim_s(e_1, e_2) \cdot \mathbf{att_1}^s + Sim_t(e_1, e_2) \cdot \mathbf{att_1}^t, \quad (6)$$

where $\mathbf{att_1}^s$ and $\mathbf{att_1}^t$ are the corresponding weight values for structural and name similarity scores, respectively. Note that $Sim(e_1, e_2) \neq Sim(e_2, e_1)$ as they may have different attention weight vectors.

The co-attention feature fusion model is of low model complexity, with only parameters $\mathbf{M}$ and $\mathbf{w}$. In addition, it can also be easily adapted for including more features.

Training. The training objective is to maximize the similarity scores of the training entity pairs, which can be converted to minimizing the following loss function:

$$L = \sum_{(e_1, e_2) \in S} [-Sim(e_1, e_2) + \gamma]_+ + [-Sim(e_2, e_1) + \gamma]_+, \quad (7)$$

where $[x]_+ = max\{0, x\}$, and γ is a constant number.

Discussion. There could be other ways to implement degree-aware weighting, e.g., by using $\texttt{sigmoid}(\mathbf{W} \cdot [\mathbf{N}(e), \mathbf{Z}(e), \mathbf{g}_e])$, where $\mathbf{W}$ is the parameter. Here we exploit a co-attention mechanism to consolidate different channels of signals with degree-aware weights, in order to demonstrate the effectiveness of leveraging degrees for EA in tail. More in-depth comparison with additional implementations is left for future work.

4.3 Iterative KG completion

Iterative self-training is an interesting idea to explore, which is shown to be effective [20, 28]. However, existing research overlooked the opportunity to enrich structure information during the iterative process. In fact, we found that although long-tail entities may lack structural information in the source KG, the target KG may possess such information in a complementary manner. Intuitively, if we can replenish the original KG with facts from its counterpart by mining confident EA results as pseudo matching pairs anchoring the subgraphs, the KGs' structural sparsity could be mitigated. This can substantially improve the coverage of KGs and reduce the amount of long-tail entities. From the amplified KGs, the structural learning model can generate increasingly better structural embedding, which in turn leads to more accurate EA results in the subsequent rounds, and it is naturally iterative.

First, we detail the inclusion of EA pairs with high confidence, with an emphasis on avoiding false pairs that might adversely affect the model. In particular, we devise a novel strategy for selecting EA pairs. For every given entity $e_1 \in E_1 - S_1$ (in G_1 but not in the training set), suppose its most similar entity in G_2 is e_2, its second most similar entity is e_2' and the difference between the similarity scores is $\Delta_1 \triangleq Sim(e_1, e_2) - Sim(e_1, e_2')$, if for e_2, its most similar entity in G_1 is exactly e_1, its second most similar entity is e_1', the difference between the similarity scores is $\Delta_2 \triangleq Sim(e_2, e_1) - Sim(e_2, e_1')$, and Δ_1, Δ_2 are both above a given threshold θ, (e_1, e_2) would be considered as a correct pair. This is a relatively strong constraint, as it requires that (1) the similarity between the two entities is the highest from both sides, respectively, and (2) there is a margin between the top-2 candidates.

After adding the highly confident EA results to seed entity pairs, we use these entities (S_a) to bridge two KGs and enrich the KGs with facts from the other side. For instance, for a triple $t_1 \in T_1$, if both its head and tail entities correspond to certain entries in S_a, we replace entities in t_1 with the corresponding entities in E_2 and add it to T_2. Although this is a very intuitive and simple approach, it effectively reduces the number of long-tail entities and improves the coverage of KGs in practice. Finally, we utilize the augmented KGs for learning better structural representation, which in turn

Table 2: Statistics of *SRPRS*

Dataset	KGs	#Triples	#Entities
EN-FR	DBpedia (English)	36,508	15,000
	DBpedia (French)	33,532	15,000
EN-DE	DBpedia (English)	38,281	15,000
	DBpedia (German)	37,069	15,000
DBP-WD	DBpedia	38,421	15,000
	Wikidata	40,159	15,000
DBP-YG	DBpedia	33,571	15,000
	YAGO3	34,660	15,000

contributes to the improvement of EA performance. This iterative completion process lasts for ζ rounds.

Discussion. Some EA methods also employ bootstrapping or iterative training strategies; but they merely aim to enlarge the training signals to update the embeddings, without touching the actual structure of KGs. Besides, their procedures to select confident EA pairs are single-sided and time-consuming [20] or apt to introduce incorrect instances [28]. In contrast, we substantially augment the KGs, and judiciously design the selection—pairing two entities only if they treat each other as the first priority, the superiority of which is witnessed in Section 5.5.

5 EXPERIMENTS

This section reports the experiments with in-depth analysis [4].

5.1 Experiment setting

Dataset. We adopt *SRPRS* [10], since the KG pairs thereof follow real-life degree distributions. This dataset was constructed by using inter-language links and reference links in DBpedia, and each entity has an equivalent entity in the other KG. The detailed information is shown in Table 2. 30% of entity pairs are harnessed for training.

Parameter settings. For the *Structural representation learning module*, we follow the settings in [10], except for assigning d_s to 300. Regarding *Name representation learning module*, we set $\mathbf{p} = [p_1, \ldots, p_K]$ to $[1, \min, \max]$. For mono-lingual datasets, we merely use the fastText embeddings [2] as the word embedding (i.e., only one embedding space in Equation 3). For cross-lingual datasets, the multilingual word embeddings are obtained from MUSE [5]. Two word embedding spaces (from two languages) are used in Equation 3. As for *Degree-aware fusion module*, we set d_g to 300, γ to 0.8, batch size to 32. Stochastic gradient descent is harnessed to minimize the loss function, with learning rate set to 0.1, and we use early stopping to prevent over-fitting. In *KG completion module*, θ is set to 0.05 and ζ is set to 3.

Evaluation metric. We utilize Hits@k (k=1, 10) and mean reciprocal rank (MRR) as evaluation metrics. For each source entity, entities in the other KG are ranked according to their similarity scores Sim to the source entity in descending order. Hits@k reflects the percentage of correctly aligned entities in top-k similar entities to source entities. In particular, Hit@1 represents the accuracy of

---

[4]The source code is available at https://github.com/DexterZeng/DAT
[5]https://github.com/facebookresearch/MUSE

Table 3: Overall results of entity alignment

Methods	EN-FR			EN-DE			DBP-WD			DBP-YG		
	Hits@1	Hits@10	MRR	Hits@1	Hits@10	MRR	Hits@1	Hits@10	MRR	Hits@1	Hits@10	MRR
MTransE	25.1	55.1	0.35	31.2	58.6	0.40	22.3	50.1	0.32	24.6	54.0	0.34
IPTransE	25.5	55.7	0.36	31.3	59.2	0.41	23.1	51.7	0.33	22.7	50.0	0.32
BootEA	31.3	62.9	0.42	44.2	70.1	0.53	32.3	63.1	0.42	31.3	62.5	0.42
RSNs	34.8	63.7	0.45	49.7	73.3	0.58	39.9	66.8	0.49	40.2	68.9	0.50
MuGNN	13.1	34.2	0.20	24.5	43.1	0.31	15.1	36.6	0.22	17.5	38.1	0.24
KECG	29.8	61.6	0.40	44.4	70.7	0.54	32.3	64.6	0.43	35.0	65.1	0.45
TransEdge	40.0	67.5	0.49	55.6	75.3	0.63	46.1	73.8	0.56	44.3	69.9	0.53
GCN-Align	15.5	34.5	0.22	25.3	46.4	0.33	17.7	37.8	0.25	19.3	41.5	0.27
JAPE	25.6	56.2	0.36	32.0	59.9	0.41	21.9	50.1	0.31	23.3	52.7	0.33
RDGCN	67.5	76.9	0.71	78.3	88.4	0.82	83.4	90.7	0.86	85.8	93.8	0.89
HGCN	67.0	77.0	0.71	76.3	86.3	0.80	82.3	88.7	0.85	82.2	88.8	0.85
GM-Align [1]	62.7	-	-	67.7	-	-	81.5	-	-	82.8	-	-
DAT	**75.8**	**89.9**	**0.81**	**87.6**	**95.5**	**0.90**	**92.6**	**97.7**	**0.94**	**94.0**	**98.5**	**0.96**

[1] When running GM-Align, it is noted that entities without valid name embeddings are excluded from evaluation, and hence we consider that GM-Align fails to align these entities without specifying rankings, which leads to the lack of Hits@10 and MRR values.

Table 4: Hits@1 results by degrees

Degree	EN-FR			EN-DE			DBP-WD			DBP-YG		
	RSNs	RDGCN	DAT	RSNs	RDGCN	DAT	RSNs	RDGCN	DAT	RSNs	RDGCN	DAT
1	14.3	56.5	**57.4**	22.9	75.8	**83.1**	20.6	80.3	**86.7**	16.2	84.2	**89.5**
2	27.5	64.4	**72.4**	40.1	78.6	**84.3**	26.9	84.0	**90.8**	31.5	80.9	**92.6**
3	36.1	77.3	**82.9**	54.2	77.4	**88.4**	33.5	75.2	**88.2**	49.3	89.2	**97.0**
≥4	57.8	75.8	**91.6**	69.9	80.0	**92.3**	62.6	88.6	**98.8**	71.9	90.2	**99.4**
All	34.8	67.5	**75.8**	49.7	78.3	**87.6**	39.9	83.4	**92.6**	40.2	85.8	**94.0**

alignment results. MRR, on the other hand, denotes the average of reciprocal ranks of ground truth results. Note that higher Hits@k and MRR indicate better performance. Unless otherwise specified, the results of Hits@k are represented in percentages. We denote the best performance in **bold** in the tables.

Competitors. Overall 13 state-of-the-art methods are involved in comparison. The group that solely utilizes structural feature includes (1) MTransE [6], which proposes to utilize TransE for EA; (2) IPTransE [28], which uses an iterative training process to improve the alignment results; (3) BootEA [20], which devises an alignment-oriented KG embedding framework and a bootstrapping strategy; (4) RSNs [10], which integrates recurrent neural networks with residual learning; (5) MuGNN [3], which puts forward a multi-channel graph neural network to learn alignment-oriented KG embeddings; (6) KECG [11], which proposes to jointly learn knowledge embeddings that encode inner-graph relationships, and a cross-graph model that enhances entity embeddings with their neighbors' information; and (7) TransEdge [21], which presents a novel edge-centric embedding model that contextualizes relation representations in terms of specific head-tail entity pairs.

Methods incorporating other types of information include (1) JAPE [19], where attributes of entities are harnessed to refine the structural information; (2) GCN-Align [23], which generates entity embeddings and attribute embeddings to align entities in different KGs; (3) GM-Align [26], where a local subgraph of an entity

is constructed to represent entity, and entity name information is harnessed for initializing the framework; (4) MultiKE [27], which offers a novel framework that unifies the views of entity names, relations and attributes at *representation-level for mono-lingual* EA; (5) RDGCN [24], which proposes a relation-aware dual-graph convolutional network to incorporate relation information via attentive interactions between KG and its dual relation counterpart; and (6) HGCN [25], where a learning framework is built to jointly learn entity and relation representations for EA.

5.2 Results

Table 3 summarizes the results. Solutions in the first group merely harness structural information for aligning. BootEA and KECG achieve better performance than MTransE and IPTransE due to the carefully designed alignment-oriented KG embedding framework and attention-based graph embedding model, respectively. RSNs further advances the results by taking into account long-term relational dependencies between entities, which can capture more structural signals for alignment. TransEdge attains the best performance due to its edge-centric KG embedding and the bootstrapping strategy. MuGNN fails to yield promising results as there are no aligned relations on *SRPRS*, where rule transferring cannot work, and the number of detected rules is rather limited. Noteworthily,

Table 5: Experiment results on dense datasets

Methods	DBP-WD-100K			DBP-YG-100K		
	Hits@1	Hits@10	MRR	Hits@1	Hits@10	MRR
MultiKE	91.9	96.3	0.94	88.0	95.3	0.91
DAT	**97.4**	**99.6**	**0.98**	**94.3**	**98.6**	**0.96**

Hits@1 values on most datasets are below 50%, showing the inefficiency of solely relying on KG structure, especially when long-tail entities account for the majority.

Regarding the second group, both GCN-Align and JAPE exploit attribute information to complement structural signals. However, they fail to outperform the leading method in the first group, which can be attributed to the limited effect of attributive information. The rest four methods all employ the generally available entity name information. The significant performance gain compared with the first group validates the usefulness of this feature. Among others, DAT outperforms GM-Align, RDGCN and HGCN by approximately 10% on Hit@1 over all datasets, verifying that our framework can best exploit entity name information. The underlying reason is that GM-Align, RDGCN and HGCN fuse features on representation level, which potentially causes the loss of information as the resulting unified feature representation cannot keep the characteristics of original features. DAT, however, utilizes a co-attention network to determine feature weights and fuses features at outcome level, namely, feature-specific similarity scores.

Evaluation by degree. To demonstrate that DAT is effective at aligning long-tail entities, we show the results by degree in Table 4. Note that for DAT, the degree refers to the initial degree distribution (as the completion process alters entity degrees).

From Table 4, it reads that for entities with degree 1, the Hits@1 values of DAT are two or even three times higher than those of RSNs, validating the effectiveness of DAT for dealing with long-tail issue. Although for popular entities, the performance also increases, the improvement over RSNs narrows greatly in comparison to the entities in tail. Moreover, DAT outperforms RDGCN across all degree groups in four datasets, despite that both of them utilize entity name information as an external signal for EA.

Comparison with MultiKE on dense datasets. The results of MultiKE on *SRPRS* are not provided since it can merely cope with mono-lingual dataset and demands semantics of relations to be known in advance. To better appreciate DAT, we supply the experimental results of DAT on the *dense* datasets where MultiKE was evaluated. Concretely, the dense datasets, DBP-WD-100K and DBP-YG-100K, are similar to DBP-WD and DBP-YG, but some with larger scale (100K entities on each side) and higher density [20].

On dense datasets, DAT yields better results, with Hits values over 90% and MRR over 0.95 (shown in Table 5), demonstrating that DAT better exploits name information, which can be attributed to our degree-aware feature fusion module and the choice of calculating scores within each view first, instead of learning a combined representation that might cause information loss.

5.3 Ablation study

We report an ablation study on EN-FR dataset in Table 6.

Table 6: Experiment results of ablation

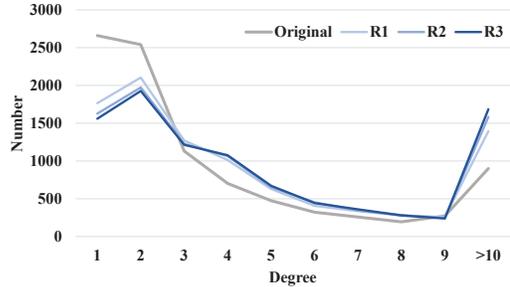| Methods	EN-FR		
	Hits@1	Hits@10	MRR
DAT	**75.8**	**89.9**	**0.81**
DAT w/o IKGC	72.1	85.4	0.77
DAT w/o KGC	73.9	88.6	0.79
DAT w/o ATT	73.1	88.5	0.79
DAT w/o CPM	75.3	89.7	0.80

Figure 5: Distribution of entity degree in EN-FR

Iterative KG completion. By removing the entire module, EA performance drops by 3.7% on Hits@1 (DAT vs. DAT w/o IKGC). If we merely eliminate the KG completion module while keep the iterative process (similar to [28]), Hits@1 also declines by 1.9% (DAT vs. DAT w/o KGC), validating the significance of KG completion. We further show the dynamic change of degree distribution after each round, i.e., Original, R1, R2, R3, in Figure 5, which indicates that the embedded KG completion enhances KG coverage and reduces the number of long-tail entities.

Degree-aware co-attention feature fusion. Table 6 shows that replacing *Degree-Aware Fusion Module* with fixed equal weights results in a 2.7% drop on Hits@1 (DAT vs. DAT w/o ATT). This validates that dynamically adjusting the weights of features according to the degree leads to better feature integration and hence more
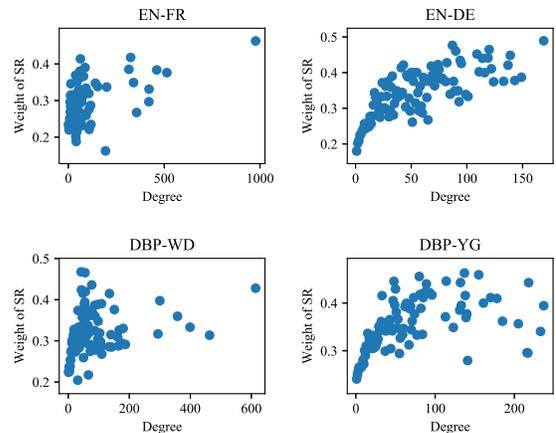

Figure 6: Weight distribution of structural representation

(a) Number of pairs selected   (b) Accuracy of pairs selected   (c) Running time consumption (s)
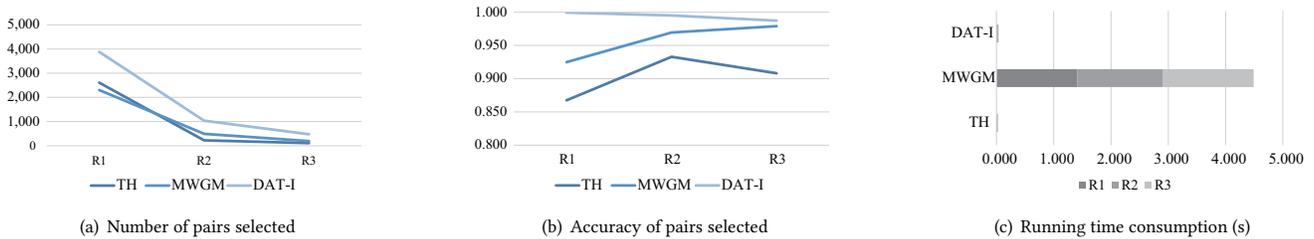
Figure 7: Comparison results of iterative training strategies

accurate alignment results. We further display the weight of structural representation generated by our degree-aware fusion model across different degrees (in 1st round) in Figure 6. It reveals that, by and large, the significance of structure information increases as the degree of entities rises, which is in accordance with our anticipation.

Concatenated power mean word embeddings. We compare concatenated power mean word embeddings and averaged word embeddings against aligning entities (DAT vs. DAT w/o CPM). The results suggest that stacking different power mean embeddings does capture more features for alignment.

5.4 Error analysis

We perform error analysis on EN-FR dataset to demonstrate the contribution of each module, as well as cases where DAT fails. On EN-FR, solely using structural information leads to 65.5% error rate on the Hits@1. 67.0% of the entities are long-tail (i.e., with degree ≤ 3), among which 65.1% are wrongly aligned. Introducing entity name information and dynamically fusing it with structural information reduces the overall Hits@1 error rate to 27.9%. For long-tail entities, the error rate is also reduced to 33.2%. On top of it, employing iterative KG completion to replenish structure and propagate the signals can further reduce the overall Hits@1 error rate to 24.2%. The percentage of long-tail entities is also reduced to 49.7%, of which merely 8.3% are wrongly aligned. It implies that long-tail entities account for the most errors initially, and then applying the proposed techniques reduces not only the error rate but also the portion of (contributing) long-tail entities.

As for the very difficult cases that DAT cannot resolve, we give following analysis with the focus on entity name information. Among the false cases (24.2% on EN-FR), 41% do not have an proper entity name embedding as all words in the name are OOVs, and 31% suffer from partial OOVs. Moreover, 15% could have achieved correct results by solely using name information but get misled by structural signals, and 13% fail to align due to the insufficiency of the entity name representation approach or the fact that the entities of the identical name denote different physical ones.

5.5 Further experiment

We further justify the effectiveness of our iterative training strategy by conducting the following experiments.

Our iterative process differs from existing methods in not only the embedded KG completion process but also the selection of confident pairs. To demonstrate the merit, we exclude the KG completion

module from DAT resulting DAT-I, to compare with the selection methods of [20, 28]. In [28], for each *non-aligned* source entity, it finds the most similar *non-aligned* target entity, and if the similarity between the two entities is above a given threshold, it is confident to consider these two as a pair. We refer to this as threshold-based method (TH). In [20], for each source entity, it computes alignment likelihood to every target entity, and only those with likelihood above a given threshold are involved in a maximum likelihood matching processing under 1-to-1 mapping constraint, which generates a solution containing confident EA pairs. We refer to this as maximum weight graph matching (MWGM). We re-implement the methods under our framework and tune the parameters according to their original papers. To measure the effectiveness of different iterative training methods, we adopt the number of selected confident EA pairs, the accuracy of these EA pairs and the running time of each round as the main metrics.

For fair comparison, we report the results of first three rounds in Figure 7. It is observed that, DAT-I outperforms the other two in terms of the number and the quality of selected pairs in comparatively less time. MWGM requires much more time as it needs to solve a global optimization problem, whereas compared with TH, it performs better in terms of the accuracy of selected pairs.

6 CONCLUSION

In this work, we offer a revisited framework DAT for entity alignment with emphasis on long-tail entities. Observing the limit of sole reliance on structural information, we propose to introduce entity name information in the form of concatenated power mean embedding in pre-alignment phase. In alignment, to consolidate different signals, we devise a co-attention feature fusion network that dynamically adjusts the weights of different features under the guide of degree. During post-alignment, we complete KG with knowledge from the other side using confident EA results as anchors in an iterative fashion, which amplifies the structural information and boosts the performance. DAT is evaluated on both cross-lingual and mono-lingual EA benchmarks and achieves superior results.

REFERENCES

[1] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. 2007. DBpedia: A Nucleus for a Web of Open Data. In *The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007, Busan, Korea, November 11-15, 2007.* 722–735.

[2] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146.

[3] Yixin Cao, Zhiyuan Liu, Chengjiang Li, Zhiyuan Liu, Juanzi Li, and Tat-Seng Chua. 2019. Multi-Channel Graph Neural Network for Entity Alignment. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers.* 1452–1461.

[4] Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., and Tom M. Mitchell. 2010. Toward an Architecture for Never-Ending Language Learning. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010.*

[5] Muhao Chen, Yingtao Tian, Kai-Wei Chang, Steven Skiena, and Carlo Zaniolo. 2018. Co-training Embeddings of Knowledge Graphs and Entity Descriptions for Cross-lingual Entity Alignment. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.* 3998–4004.

[6] Muhao Chen, Yingtao Tian, Mohan Yang, and Carlo Zaniolo. 2017. Multilingual Knowledge Graph Embeddings for Cross-lingual Knowledge Alignment. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017.* 1511–1517.

[7] Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word Translation Without Parallel Data. *CoRR* abs/1710.04087 (2017).

[8] Xin Dong, Evgeniy Gabrilovich, Geremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: a web-scale approach to probabilistic knowledge fusion. In *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014.* 601–610.

[9] Luis Galárraga, Simon Razniewski, Antoine Amarilli, and Fabian M. Suchanek. 2017. Predicting Completeness in Knowledge Bases. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM 2017, Cambridge, United Kingdom, February 6-10, 2017.* 375–383.

[10] Lingbing Guo, Zequn Sun, and Wei Hu. 2019. Learning to Exploit Long-term Relational Dependencies in Knowledge Graphs. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA.* 2505–2514.

[11] Chengjiang Li, Yixin Cao, Lei Hou, Jiaxin Shi, Juanzi Li, and Tat-Seng Chua. 2019. Semi-supervised Entity Alignment via Joint Knowledge Embedding Model and Cross-graph Model. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP).* 2723–2732.

[12] Ning Pang, Weixin Zeng, Jiuyang Tang, Zhen Tan, and Xiang Zhao. 2019. Iterative Entity Alignment with Improved Neural Attribute Embedding. In *Proceedings of the Workshop on Deep Learning for Knowledge Graphs (DL4KG2019) Co-located with the 16th Extended Semantic Web Conference 2019 (ESWC 2019), Portoroz, Slovenia, June 2, 2019.* 41–46.

[13] Heiko Paulheim. 2017. Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web* 8, 3 (2017), 489–508.

[14] Shichao Pei, Lu Yu, Robert Hoehndorf, and Xiangliang Zhang. 2019. Semi-Supervised Entity Alignment via Knowledge Graph Embedding with Awareness of Degree Difference. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019.* 3130–3136.

[15] George Polya, Godfrey Harold Hardy, and Littlewood. 1952. *Inequalities.* University Press.

[16] Andreas Rücklé, Steffen Eger, Maxime Peyrard, and Iryna Gurevych. 2018. Concatenated p-mean Word Embeddings as Universal Cross-Lingual Sentence Representations. *CoRR* abs/1803.01400 (2018). arXiv:1803.01400 http://arxiv.org/abs/1803.01400

[17] Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional Attention Flow for Machine Comprehension. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings.* https://openreview.net/forum?id=HJ0UKP9ge

[18] Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007.* 697–706.

[19] Zequn Sun, Wei Hu, and Chengkai Li. 2017. Cross-Lingual Entity Alignment via Joint Attribute-Preserving Embedding. In *The Semantic Web - ISWC 2017 - 16th International Semantic Web Conference, Vienna, Austria, October 21-25, 2017, Proceedings, Part I.* 628–644.

[20] Zequn Sun, Wei Hu, Qingheng Zhang, and Yuzhong Qu. 2018. Bootstrapping Entity Alignment with Knowledge Graph Embedding. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.* 4396–4402.

[21] Zequn Sun, JiaCheng Huang, Wei Hu, Muhao Chen, Lingbing Guo, and Yuzhong Qu. 2019. TransEdge: Translating Relation-Contextualized Embeddings for Knowledge Graphs. In *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference, Auckland, New Zealand, October 26-30, 2019, Proceedings, Part I.* 612–629.

[22] Bayu Distiawan Trisedya, Jianzhong Qi, and Rui Zhang. 2019. Entity Alignment between Knowledge Graphs Using Attribute Embeddings. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019.* 297–304.

[23] Zhichun Wang, Qingsong Lv, Xiaohan Lan, and Yu Zhang. 2018. Cross-lingual Knowledge Graph Alignment via Graph Convolutional Networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018.* 349–357.

[24] Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, Rui Yan, and Dongyan Zhao. 2019. Relation-Aware Entity Alignment for Heterogeneous Knowledge Graphs. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019.* 5278–5284.

[25] Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, and Dongyan Zhao. 2019. Jointly Learning Entity and Relation Representations for Entity Alignment. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP).* 240–249.

[26] Kun Xu, Liwei Wang, Mo Yu, Yansong Feng, Yan Song, Zhiguo Wang, and Dong Yu. 2019. Cross-lingual Knowledge Graph Alignment via Graph Matching Neural Network. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers.* 3156–3161.

[27] Qingheng Zhang, Zequn Sun, Wei Hu, Muhao Chen, Lingbing Guo, and Yuzhong Qu. 2019. Multi-view Knowledge Graph Embedding for Entity Alignment. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019.* 5429–5435.

[28] Hao Zhu, Ruobing Xie, Zhiyuan Liu, and Maosong Sun. 2017. Iterative Entity Alignment via Joint Knowledge Embeddings. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017.* 4258–4264.

[29] Qiannan Zhu, Xiaofei Zhou, Jia Wu, Jianlong Tan, and Li Guo. 2019. Neighborhood-Aware Attentional Representation for Multilingual Knowledge Graphs. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019.* 1943–1949.