JarKA: Modeling Attribute Interactions for Cross-lingual Knowledge Alignment

Bo Chen^{1,2}, Jing Zhang^{1,2(✉)}, Xiaobin Tang^{1,2}, Hong Chen^{1,2},
and Cuiping Li^{1,2}

¹ Key Laboratory of Data Engineering and Knowledge Engineering
of Ministry of Education, Renmin University of China, Beijing, China

² Information School, Renmin University of China, Beijing, China
{bochen,zhang-jing,txb,chong,licuiping}@ruc.edu.cn

Abstract. Cross-lingual knowledge alignment is the cornerstone in building a comprehensive knowledge graph (KG), which can benefit various knowledge-driven applications. As the structures of KGs are usually sparse, attributes of entities may play an important role in aligning the entities. However, the heterogeneity of the attributes across KGs prevents from accurately embedding and comparing entities. To deal with the issue, we propose to model the interactions between attributes, instead of globally embedding an entity with all the attributes. We further propose a joint framework to merge the alignments inferred from the attributes and the structures. Experimental results show that the proposed model outperforms the state-of-art baselines by up to 38.48% HitRatio@1. The results also demonstrate that our model can infer the alignments between attributes, relationships and values, in addition to entities.

1 Introduction

DBpedia, Freebase, YAGO and so on have been published as noteworthy large and freely available knowledge graphs (KGs), which can benefit many knowledge-driven applications. However, the knowledge embedded in different languages is extremely unbalanced. For example, DBpedia contains about 2.6 billion triplets in English, but only 889 million and 278 million triplets in French and Chinese respectively. Creating the linkages between cross-lingual KGs can reduce the gap of acquiring knowledge across multiple languages and benefit many applications such as machine translation, cross-lingual QA and cross-lingual IR.

Recently, much attention has been paid to leveraging the embedding techniques to align entities between two KGs. Some of them only leverage the structures of the KGs, i.e., the relationship triplets in the form of $\langle \text{entity}, \text{relationship}, \text{entity} \rangle$ to learn the structure embeddings of entities [3, 6, 10]. However, the structures of some KGs are sparse, making it difficult to learn the structure embeddings accurately. Other efforts are made to incorporate the attribute triplets in the form of $\langle \text{entity}, \text{attribute}, \text{value} \rangle$ to learn the attribute embeddings of entities [9, 11, 12, 15]. For example, JAPE [9] embeds attributes via attributes'

Fig. 1. Illustration of different attributes of same entities in two cross-lingual knowledge graphs from wikipedia.

concurrence. Wang et al. [12] adopt GCNs to embed entities with the one-hot representations of the attributes. Trsedya et al. [11] and MultiKE [15] embed the literal values of the attributes. Despite the existing studies on incorporating the attribute triplets to align entities, there are still unsolved challenges.

Challenge 1: Heterogeneity of Attributes. Different KGs may hold heterogeneous attributes, resulting in the difficulty of aligning entities. For example, in Fig. 1, two entities from cross-lingual KGs named “Audi RSQ” are the same entity. Although the attributes “Manufacturer” and “Body style” and their values in English correspond to certain attribute triplets in Chinese, there are still many attribute such as “Designer” and “Engine” in English that cannot find any counterpart in Chinese. However, if we embed an entity by all its attribute triplets and then compare two entities by their attribute embeddings [11, 15], the effects of the same attribute triplets will be diluted by other different ones.

Challenge 2: Multi-view Combination. To combine the effects from attributes and structures, existing works usually learn a combined embedding for each entity, based on which they infer the alignments. For example, JAPE [9] and AttrE [11] refine the structure embeddings by the closeness of the corresponding attribute embeddings. MultiKE [15] map the attribute and structure embeddings into a unified space. However, the issue of the missing attributes or relationships triplets may result in the inaccurate attribute or structure embeddings, which will propagate the errors to the combined embeddings.

Besides the above two challenges, most of the existing works [3, 9, 15] only focus on aligning entities, or at most relationships, but ignore attributes and values. However, the alignment of different objects influence each other. A unified way to align all of these objects simultaneously is worth studying.

Solution. To deal with the above challenges, we propose a joint model—JarKA to Jointly model the atttributes interactions and relationships for cross-lingual Knowledge Alignment. The two views are carefully merged to reinforce the training performance iteratively. The contributions can be summarized as:

- We comprehensively formalize cross-lingual knowledge alignment as linking entities, relationships, attributes and values across cross-lingual KGs.
- To tackle the first challenge, we propose an interaction-based attribute model to capture the attribute-level interactions between two entities instead of

globally representing the two entities. A matrix-based strategy is further proposed to accelerate the similarity estimation.

- To deal with the second challenge, we propose a joint framework to combine the alignments inferred by the attribute model and relationship model respectively instead of learning a combined embedding. Three different merge strategies are proposed to solve the conflicting alignments.
- Experimental results on several datasets of cross-lingual KGs demonstrate that JarKA significantly outperforms state-of-the-art comparison methods (improving 2.35–38.48% in terms of Hit Ratio@1).

2 Problem Definition

Definition 1. Knowledge Graph: We denote the KG as union of the **relationship triplets** and the **attribute triplets**, i.e., $G = \{(h, r, t)\} \cup \{(h, a, v)\}$, where (h, r, t) is a relationship triplet consisting of a head entity h , a relationship r , and a tail entity t , and (h, a, v) is an attribute triplet consisting of a head entity h , an attribute a and its value v . We also use e to denote entity.

We distinguish the two kinds of triplets as they are independent views that can take different effects on alignment.

Problem 1. Cross-lingual Knowledge Alignment: Given two cross-lingual KGs G and G' , and the seed set I of the aligned entities, relationships, attributes, and values, i.e., $I = \{(e \sim e')\} \cup \{(r \sim r')\} \cup \{(a \sim a')\} \cup \{(v \sim v')\}$ ¹, the goal is to augment I by the inferred new alignments between G and G' .

3 JarKA Model

We propose an interaction-based attribute model to leverage the (h, a, v) triplets, an embedding-based relationship model to leverage the $\{(h, r, t)\}$ triplets, and then incorporate the two models by a carefully designed joint framework.

3.1 Interaction-Based Attribute Model

Existing methods represent an entity globally by all its associated (h, a, v) and then compare the entity embedding between entities [11, 15]. However, as shown in Fig. 1, two entities from cross-lingual KGs may have heterogeneous attribute. The irrelevant attribute triplets between two entities may dilute the effects of their similar attribute triplets if globally embedding the entities.

To deal with the above issue, we propose an interaction-based attribute model to directly estimate the similarity of two entities by capturing the interactions between their attributes and values. The model mimics the process that humans solve the problem. The humans usually align two entities if they have many

¹ Please refer to Sect. 4 for how to obtain I .

same attributes with same values. Following this, we firstly find all the aligned attribute pairs of two entities, and then compare its values. Since the number of the attributes are far smaller than that of the values in KGs, we initialize the aligned attributes by the attribute seed pairs and gradually extend them by our joint framework, which will be introduced in the following section. To compare the large number of cross-lingual values, we train a machine translation model and use it to estimate the BLEU score [8] of two cross-lingual values as their similarity. Unfortunately, following the idea, we need to enumerate and invoke the translation model for maximal M attribute pairs for each entity pair, resulting in $O(N \times N' \times M)$ time complexity when there are N and N' entities in G and G' respectively, which is too inefficient to finish within available time. To accelerate the similarity estimation, we represent each knowledge graph as a 3-dimension value embedding matrix and then perform an efficient matrix-based strategy to calculate entity similarities. Figure 2 illustrates the whole process of the proposed attribute model. In the following part, we will explain the details.

Embed Cross-lingual Attribute Values. We build an neural machine translation model (NMT) [2] to capture semantic similarities between cross-lingual values. We pretrain NMT based on the value seeds². Since the seeds are limited, we will update NMT by the newly discovered value seeds iteratively.

Then we use NMT to project cross-lingual values into the same vector space. Specifically, for each attribute of each entity in G , we first invoke NMT to predict the translated value given its original value, and then look up the word embedding for each word in the translated value. While for each attribute of each entity in G' , we directly look up the word embedding for each word in its original value. With the help of NMT, the embeddings of the cross-lingual values can be unified in the same space. Then we average all the word embeddings in the value as its value embedding. The dimension is denoted as D_v .

Estimate Entity Similarities by Matrix-Based Strategy. We construct a 3-dimension value embedding matrix $\mathbf{V} \in \mathbb{R}^{N \times M \times D_v}$ for G and a similar matrix $\mathbf{V}' \in \mathbb{R}^{N' \times M \times D_v}$ for G' , where each element $\mathbf{V}_{mi}$ indicates the i -th value embedding of the m -th entity. Then we use the einsum operation

$$\text{einsum}(NMD_v, N'MD_v \rightarrow NN'MM), \quad (1)$$

i.e., Einstein summation convention [1], to make a multi-dimensional matrix product of $\mathbf{V}$ and $\mathbf{V}'$ to obtain the value similarity matrix $\hat{\mathbf{V}} \in \mathbb{R}^{N \times N' \times M \times M}$.

What's more, it is unnecessary to compare the values of different attributes. For example, although the attributes "birthplace" and "deathplace" have the same value "New York", they cannot reflect the similarity of two entities. So, we build an attribute mask matrix $\hat{\mathbf{A}}$ to limit the computation within the values of the aligned attributes. Specifically, we prepare a 3-dimension attribute identification matrix $\mathbf{A} \in \mathbb{R}^{N \times M \times K}$ for G and $\mathbf{A}' \in \mathbb{R}^{N' \times M \times K}$ for G' , where K denotes the number of the united frequent attributes in G and G' . Each row in $\mathbf{A}$ or $\mathbf{A}'$ is an one-hot vector, with an element $A_{mik} = 1$ if the i -th value of the

² In the future, external cross-lingual corpus can be easily used to pre-train the model.

Fig. 2. Illustration of the proposed attribute model. V and V' are the value embedding matrices, and A and A' are the attribute identification matrices for G and G' respectively. The figure can be read from left to right and top to bottom.

m -th entity belongs to the k -th attribute, and $A_{mik} = 0$ otherwise. Note the one-hot identification vector depends on the existing aligned attributes, which will be gradually extended with the joint model iteratively. Whenever two attributes are discovered to be aligned, we will unify their identification. For example, when the k -th attribute in G and the t -th attribute in G' are aligned, we replace the identification k with t , i.e., any row with $A_{mik} = 1$ will be changed to $A_{mit} = 1$. Then we multiply A and A' in the same way as Eq. (1) to obtain an attribute mask matrix $\tilde{A} \in \mathbb{R}^{N \times N' \times M \times M}$, where each element $\tilde{A}_{mnij} = 1$ if the i -th value of the m -th entity in G corresponds to the same attribute of the j -th value of the n -th entity in G' , and $\tilde{A}_{mnij} = 0$ otherwise. Then we calculate the element-wise product of $\tilde{V}$ and $\tilde{A}$, i.e., $\tilde{V} \odot \tilde{A}$, to get the masked value similarity matrix. Finally, we summarize the similarities of all the M^2 attribute pairs for each entity pair to obtain an entity similarity matrix:

$$S_e^A = \sum_i^M \sum_j^M \tilde{V}_{\cdot, i, j} \odot \tilde{A}_{\cdot, i, j}, \quad (2)$$

$\mathbf{S}_e^A \in \mathbb{R}^{N \times N'}$. The superscript A indicates the entity similarities are estimated by the attribute model. The above matrix computation is quite efficient, as the most expense comes from the construction of the value embedding matrices, which only requires invoking the translation model $O(N \times M)$ times.

3.2 Embedding-Based Relationship Model

Due to the success of the existing works on modeling the structures of the graph comprised by $\{(h, r, t)\}$ [10, 15], we adopt TransE algorithm to maximize the energy (possibility) that h can be translated to t in the KG, i.e., $E(h, r, t) = \|\mathbf{h} + \mathbf{r} - \mathbf{t}\|$, where $\mathbf{h}$, $\mathbf{r}$ and $\mathbf{t}$ represent the structure embeddings.

To preserve the cross-lingual relations of entities and relationships included in the existing alignments, we swap the entities or relationships in each alignment ($e \sim e'$) or ($r \sim r'$) to generate new relationship triplets [10]. Then a margin-based loss function is optimized on the all relationship triplets to obtain entity embeddings $\mathbf{H}_e \in \mathbb{R}^{N \times D_e}$, $\mathbf{H}'_e \in \mathbb{R}^{N' \times D_e}$ and relationship embeddings $\mathbf{H}_r \in \mathbb{R}^{L \times D_r}$ and $\mathbf{H}'_r \in \mathbb{R}^{L' \times D_r}$ for both G and G' , where L and L' are the number of relationships and D_e and D_r are the embedding sizes with $D_e = D_r$.

3.3 Jointly Modeling the Attribute and Relationship Model

Different from existing works that learn combined embeddings by the attribute and the relationship model [9, 11, 15], we propose a joint framework to firstly infer the confident alignments by the two models and then combine their inferences by three different merge strategies. Algorithm 1 illustrates the whole process. At each iteration, for modeling the attribute triplets, we first train the translation model based on the seed set of the aligned values (Line 3). Then we construct the value embedding matrices by the translation model (Line 4) and meanwhile construct the mask matrices by the existing aligned attributes (Line 5), based on which we perform an efficient matrix-based strategy to calculate the entity similarities (Line 6) and finally infer the new alignments of entities, attributes, and values based on the estimated similarities and existing alignments (Line 7). For modeling the relationship triplets, we train the entity and relationship embeddings based on the swapped relationship triplets between two graphs (Line 8), then we infer the new alignments of entities and relationships (Line 9). Finally, we merge the new aligned entity seeds from the attribute and the relationship model (Line 10) and augment the seed set by all the new alignments (Line 11 and 12). The framework bootstraps the two models iteratively by the extended alignments. Note we remove the new alignments from the candidate pairs at each iteration to avoid duplicate inference (Line 13).

Infer Alignments by the Attribute Model. We select an entity pair ($e_m \sim e'_n$) from all the $N \times N'$ candidate entity pairs into the new aligned set of entities I_e^A if their similarity $S_e^A[m, n]$ is larger than a threshold τ_e^A :

$$I_e^A \leftarrow (e_m \sim e'_n), \text{ if } S_e^A[m, n] > \tau_e^A. \quad (3)$$

Algorithm 1. JarKA

Input: G , G' and the seed alignments
 $I = \{(e \sim e')\} \cup \{(r \sim r')\} \cup \{(a \sim a')\} \cup \{(v \sim v')\}$.

Output: The augmented alignments I .

```

1  $C =$  all the candidate pairs except the seed pairs in  $I$ ;
2 repeat
3   /* Attribute model */
4   Train a translation model on  $\{(v \sim v')\}$ ;
5   Construct the value embedding matrices  $\mathbf{V}$  and  $\mathbf{V}'$  by the translation model;
6   Construct the mask matrices  $\mathbf{A}$  and  $\mathbf{A}'$  by  $\{(a \sim a')\}$ ;
7   Calculate entity similarities  $\mathbf{S}_e^A$  by Eq. (2);
8   Infer new alignments  $I_e^A$ ,  $I_a$ ,  $I_v$  from  $C$  by Eq. (3),(4),(5);
9   /* Relationship model */
10  Train the relationship model on  $\{(h, r, t)\}$ ,  $\{(h', r', t')\}$ ,  $\{(e \sim e')\}$  and  $\{(r \sim r')\}$  to
11  obtain  $\mathbf{H}_e$ ,  $\mathbf{H}'_e$ ,  $\mathbf{H}_r$  and  $\mathbf{H}'_r$ ;
12  Infer new alignments  $I_e^R$  and  $I_r$  from  $C$  by Eq. (6);
13  /* Merge new alignments */
14   $I_e \leftarrow \text{merge}(I_e^A, I_e^R)$ ;
15   $\Delta I \leftarrow (I_e, I_r, I_a, I_v)$ ;
16   $I = I + \Delta I$ ;
17   $C = C - \Delta I$ ;
18 until  $\Delta I = \emptyset$ ;
```

The candidates of the aligned attributes and values depend on the aligned entities. Specifically, for each aligned entity pair $(e_m \sim e'_n)$, if the similarity of a value pair $\tilde{V}_{m,n,i,j}$ is larger than a threshold τ_v , we select their corresponding attribute pair $(a_i \sim a'_j)$ into the new aligned attribute set I_a :

$$I_a \leftarrow (a_i \sim a'_j), \forall (e_m \sim e'_n) \in I, \text{ if } \tilde{V}_{m,n,i,j} > \tau_v. \quad (4)$$

Then for each pair of attribute triplets $(e_m, a_i, v_i) \in G$ and $(e'_n, a'_j, v'_j) \in G'$, if the entities and the attributes are both aligned, we select their corresponding value pairs $(v_i \sim v'_j)$ into the new aligned value set I_v :

$$I_v \leftarrow (v_i \sim v'_j), \forall (e_m, a_i, v_i) \in G \ \& \ (e'_n, a'_j, v'_j) \in G', \quad (5)$$

$$\text{if } (e_m \sim e'_n) \in I \ \& \ (a_i \sim a'_j) \in I.$$

Infer Alignments by the Relationship Model. We calculate the similarity matrix $\mathbf{S}_e^R$ as the dot product of the entity embeddings where the superscript R indicates the entity similarities are estimated by the relationship model. Then we select an entity pair $(e_m \sim e'_n)$ into the new aligned entity set I_e^R if their similarity $S_e^R[m, n]$ is larger than a threshold τ_e^R :

$$I_e^R \leftarrow (e_m \sim e'_n), \text{ if } S_e^R[m, n] > \tau_e^R. \quad (6)$$

The new aligned relationships are inferred in the same way but with a different threshold τ_r .

Merge Alignments of the Two models. We propose three strategies to merge I_e^A and I_e^R into I_e .

Standard Multi-view Merge Strategy. Following the standard co-training algorithm, we firstly infer I_e^A from candidate entity pairs C_e by Eq. (3). Then we remove I_e^A from C_e , and then infer I_e^R from the remaining candidates $C_e - I_e^A$.

Score-Based Merge Strategy. Due to the missing attributes and relationships, the labels inferred from the two views may have conflicts. For all the conflicting counterparts of an entity e_m , i.e., $C_m = \{e'_n | (e_m \sim e'_n) \in I_e^A \cup I_e^R\}$, we select the counterpart with the maximal score $S_e[m, n] = S_e^A[m, n] + S_e^R[m, n]$ into the final new alignments. The strategy assumes that the alignments discovered by more views will be more confident:

$$I_e \leftarrow (e_m \sim e'_n), e'_n = \operatorname{argmax}_{e'_n \in C_m} S_e[m, n]. \quad (7)$$

Rank-Based Merge Strategy. Directly comparing the similarities estimated by the two models may suffer from the different scales of scores. Thus, we compare the normalize ranking indexes of the conflicting alignments. Specifically, for all the conflicting counterparts C_m of e_m , we select the counterpart with the minimal ranking ratio $R[m, n]$ into the final new alignments:

$$I_e \leftarrow (e_m \sim e'_n), e'_n = \operatorname{argmin}_{e'_n \in C_m} R[m, n],$$

$$R[m, n] = \begin{cases} r^A[m, n] / |I_e^A|, & \text{if } (e_m \sim e'_n) \in I_e^A; \\ r^R[m, n] / |I_e^R|, & \text{if } (e_m \sim e'_n) \in I_e^R. \end{cases}$$

where $r^A[m, n]$, $r^R[m, n]$ denote the ranking index of the alignment $(e_m \sim e'_n)$ in I_e^A and I_e^R respectively, and $R[m, n]$ denotes the normalized ranking index.

Table 1. Data statistics. Notation #Rt denotes the number of the relationship triplets, #Ar denotes the number of the attribute triplets.

Dataset	#Ent.	#Rel.	#Attr.	#Rt	#At
ZH-EN	164,594	5,147	15,286	391,603	947,439
JA-EN	161,424	4,139	11,948	397,692	851,849
FR-EN	172,747	3,588	10,969	470,781	1,105,208

4 Experiments

4.1 Experimental Settings

Dataset. We evaluate the proposed model on DBP15K³, a well-known public dataset for KG alignment. DBP15K contains 3 pairs of cross-lingual KGs, each of which contains 15,000 inter-lingual links (ILLs). The proportion of the ILLs for training, validating and testing is 4:1:10. Table 1 shows the data statistics.

³ <https://github.com/nju-websoft/JAPE>.

Baseline Methods. We compare several existing methods:

MuGNN [3]: Learns the structure embeddings by a multi-channel GNNs.

BootEA [10]: Is a bootstrap method that finds new alignments by performing a maximal matching between the structure embeddings of the entities.

JAPE [9]: Leverages the attributes and the type of values to refine the structure embeddings.

GCNs [12]: Learns the structure embeddings by GCNs and use the one-hot representation of all the attributes as the initial input of an entity.

MultiKE [15]: Learns a global attribute embedding for each entity and combines it with the structure embedding. Since it is to solve monolingual entity alignment, for a fair comparison, we translate all the words into English by Google’s translator and then applies MultiKE.

JarKA: Is our model. The variant JarKA-r removes the relationship model and JarKA-a removes the attribute model, the bootstrap strategy is still adopted.

As KDCoE [4] and Yang et al. [14] leverage the descriptions of entities, and Xu et al. [13] adopts external cross-lingual corpus to train embeddings, we do not compare with them and leave the studies with these resources in the future.

Evaluation Metrics. In the test set, for each entity in G , we rank all the entities in G' by either $\mathbf{S}_e^A$ or $\mathbf{S}_e^R$, and evaluate the ranking results by HitRatio@K (HRK), i.e., the percentage of entities with the rightly aligned entities ranked before top K, and Mean Reciprocal Rank (MRR), i.e., the average of the reciprocal ranks of the rightly aligned entities.

Implementation Details. In the attribute model, the value embedding size D_v is 100, the maximal number of attributes M is 20, and the frequent attributes are those occurred more than 50 times in $\{(h, a, v)\}$. In the relationship model, the entity/relationship embedding size D_e or D_r is 75, and $\gamma = 1.0$. The thresholds τ_e^A and τ_e^R for selecting the aligned entities are set as the values when the best HR1 is obtained on the validation set. τ_v for selecting the aligned attributes is 0.8, and τ_r for selecting the aligned relationships is 0.9.

Initial Seeds Construction. The existing ILLs can be viewed as the entity seed alignments. Some relationships or attributes in cross-lingual knowledge graphs are both represented in English. So we can easily treat a pair of relationships or attributes with the same name⁴ as a relationship or attribute seed alignment. Finally, the corresponding values of the aligned attributes for any aligned entity pairs are added into the seed set of the aligned values.

4.2 Experimental Results

Overall Alignment Performance. Table 2 shows the overall performance of entity alignment. MuGNN and BootEA only leverage relationship triplets.

⁴ Attributes and relationships are less ambiguous than entities. The sampled 500 pairs of attributes and relationships present that about 95% of them can be safely aligned only based on the same names.

Table 2. Overall performance of entity alignment (%).

Model	DBP15K _{ZH-EN}			DBP15K _{JA-EN}			DBP15K _{FR-EN}		
	HR1	HR10	MRR	HR1	HR10	MRR	HR1	HR10	MRR
MuGNN	49.40	84.40	61.10	50.10	85.70	62.10	49.50	87.00	62.10
BootEA	62.94	84.75	70.30	62.23	85.39	70.10	65.30	87.44	73.10
JAPE	41.18	74.46	49.00	36.25	68.50	47.60	32.39	66.68	43.00
GCNs	41.25	74.38	55.80	39.91	74.46	55.20	37.29	74.49	53.40
MultiKE	50.87	57.61	53.20	39.30	48.85	42.60	63.94	71.19	66.50
JarKA-r	57.18	70.44	61.80	50.63	60.36	54.30	53.92	60.40	56.30
JarKA-a	58.64	83.89	67.10	55.74	83.23	65.10	59.25	85.74	68.60
JarKA(M1)	68.59	86.56	74.90	62.65	82.79	69.70	68.43	87.86	75.10
JarKA(M2)	69.32	87.37	75.50	63.01	83.37	70.00	70.87	87.05	76.50
JarKA(M3)	70.58	87.81	76.60	64.58	85.50	70.80	70.41	88.81	76.80
JarKA-IT	66.39	87.29	73.40	60.08	84.45	68.20	68.31	88.33	75.40

BootEA bootstraps the alignments iteratively and performs better than MuGNN. Although JAPE and GCN additionally consider the attribute triplets, they perform much worse than BootEA, as they only leverage the attributes but ignore their corresponding values. MultiKE utilizes the values and performs better than JAPE and GCN. However, it learns and compares the global embeddings of entities, which may bring in additional noises by the irrelevant attribute triplets.

JarKA proposes an interaction-based attribute model which directly compares the values of the aligned attributes, thus it clearly performs better than others (+2.35-38.48% in HR1). JarKA also outperforms the variant JarKA-r and JarKA-a. Specifically, JarKA-r is comparable to JarKA-a in HR1 but underperforms in HR10 and MRR. Because in JarKA-r, we set a strict threshold (Cf. Fig. 3(a)) to obtain high-qualified alignments, which makes the translation model not easy to include the difficult alignments into the training data, i.e., the seemingly irrelevant value pairs which in fact indicate the same things.

The Effect of Different Merge Strategies. We compare the effects of the proposed three merge strategies and show the results of JarKA(M1), (M2) and (M3) in Table 2. We can see that the standard multi-view merge strategy (M1) performs worst, as it does not solve the conflicts from the two views. The score-based merge strategy (M2) and the rank-based merge strategy (M3) solve the conflicts, thus perform better than M1 (+1.26-2.43% in HR1). M3 avoids comparing the scores of different scales, thus performs better than M2 in most of the metrics. Later, JarKA indicates the proposed model with M3.

The Effect of Iteratively Update the Translation Model. We validate the effect of iteratively updating the translation model (IT) during the joint modeling process. Specifically, we compare JarKA with the translation model

being trained only once at the beginning, which is denoted as JarKA-IT. From Table 2, we can see that JarKA-IT performs worse JarKA by 2.56-4.50% in HR1, which indicates that the newly discovered value alignments by our model can boost the performance of the translation model.

The Effect of τ_v and τ_r . We verify how the new aligned attributes can benefit the entity alignment. Specifically, we vary the threshold τ_v from 0.6 to 1.0 with interval 0.1 and show the results of JarKA-r on DBP15K-1 FR-EN in Fig. 3(a). It is shown that when $\tau_v = 1.0$, i.e., #new aligned attributes is 0, the accuracy of entity alignment is significantly hurt. When $\tau_v < 1.0$, with the increase of #new aligned attributes, the accuracy improves and approaches the best when $\tau_v = 0.8$, as the quantity and the quality of the new aligned attributes are well balanced. The threshold τ_r for finding the new aligned relationships is set in the same way.

(a) Effect of τ_v

Relationship	夫-mate, équipe-team, コーチ-coach		
Attribute	面积排名 - arearank, année - year, 誕生日 - birthdate		
Value	塔吉克足球联赛 - tajik league, nordEst-northeast, ダマスカス-damascus		
	Relationship	Attribute	Value
#Init. align.	774	430	21,913
#Final align.	1580	1544	188,872
Accuracy	90%	89%	88%

(b) Case study

Fig. 3. Parameter analysis and case study.

Case Study. We present several cases of the new aligned relationships, attributes and values in different languages by JarKA on DBP15K in Fig. 3(b). We also show the number of initial and the finally discovered alignments on DBP15K ZH-EN. Most of the newly discovered alignments are high-frequent attributes or relationships. The low-frequent attributes or relationships are difficult to be aligned by the current method and will be studied in the future. We randomly sample 100 final alignments and manually evaluate the accuracy, as their ground truth is not available. The results demonstrate the effectiveness of our model. The whole alignments together with the codes are available online⁵.

5 Conclusions and Future Work

We present the first attempt to formalize the problem of cross-lingual knowledge alignment as comprehensively linking entities, relationships, attributes and values. We propose an interaction-based attribute model to compare the aligned

⁵ <https://github.com/BoChen-Daniel/PAKDD-20-JarKA>.

attributes of entities instead of globally embedding entities. A matrix-based strategy is adopted to accelerate the comparing process. Then we propose a joint framework together with three merge strategies to solve the conflicts of the alignments inferred from the attribute model and the relationship model. The experimental results demonstrate the effectiveness of the proposed model. In the future, we plan to incorporate the descriptions of entities and the pre-trained cross-lingual language model to enhance our model's performance.

Acknowledgements. This work is supported by National Key R&D Program of China (No. 2018YFB1004401) and NSFC under the grant No. 61532021, 61772537, 61772536, 61702522.

References

1. Ahlander, K.: Einstein summation for multidimensional arrays. *Comput. Math. Appl.* **44**, 1007–1017 (2002)
2. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. In: *ICLR 2015* (2015)
3. Cao, Y., Liu, Z., Li, C., Li, J., Chua, T.-S.: Multi-channel graph neural network for entity alignment. In: *ACL 2019*, pp. 1452–1461 (2019)
4. Chen, M., Tian, Y., Chang, K., Skiena, S., Zaniolo, C.: Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment. In: *IJCAI 2018*, pp. 3998–4004 (2018)
5. Chen, M., Tian, Y., Yang, M., Zaniolo, C.: Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. In: *IJCAI 2017*, pp. 1511–1517 (2017)
6. Hao, Y., Zhang, Y., He, S., Liu, K., Zhao, J.: A joint embedding method for entity alignment of knowledge bases. In: Chen, H., Ji, H., Sun, L., Wang, H., Qian, T., Ruan, T. (eds.) *CCKS 2016. CCIS*, vol. 650, pp. 3–14. Springer, Singapore (2016). https://doi.org/10.1007/978-981-10-3168-7_1
7. Li, S., Li, X., Ye, R., Wang, M., Su, H., Ou, Y.: Non-translational alignment for multi-relational networks. In: *IJCAI 2018* (2018)
8. Papineni, K., Roukos, S., Ward, T., Zhu, W.-J.: Bleu: a method for automatic evaluation of machine translation. In: *ACL 2002*, pp. 311–318 (2002)
9. Sun, Z., Hu, W., Li, C.: Cross-lingual entity alignment via joint attribute-preserving embedding. In: *ISWC 2017*, pp. 628–644 (2017)
10. Sun, Z., Hu, W., Zhang, Q., Qu, Y.: Bootstrapping entity alignment with knowledge graph embedding. In: *IJCAI 2018*, pp. 4396–4402 (2018)
11. Trsedya, B.D., Qi, J., Zhang, R.: Entity alignment between knowledge graphs using attribute embeddings. In: *AAAI 2019* (2019)
12. Wang, Z., Lv, Q., Lan, X., Zhang, Y.: Cross-lingual knowledge graph alignment via graph convolutional networks. In: *EMNLP 2018*, pp. 349–357 (2018)
13. Xu, K., et al.: Cross-lingual knowledge graph alignment via graph matching neural network. In: *ACL 2019*, pp. 1452–1461 (2019)
14. Yang, H.-W., Zou, Y., Shi, P., Lu, W., Lin, J., Xu, S.: Aligning cross-lingual entities with multi-aspect information. In: *EMNLP 2019*, pp. 4422–4432 (2019)
15. Zhang, Q., Sun, Z., Hu, W., Chen, M., Guo, L., Qu, Y.: Multi-view knowledge graph embedding for entity alignment. In: *AAAI 2019* (2019)
16. Zhu, H., Xie, R., Liu, Z., Sun, M.: Iterative entity alignment via joint knowledge embeddings. In: *IJCAI 2017*, pp. 4258–4264 (2017)