

Coordinated Reasoning for Cross-Lingual Knowledge Graph Alignment

Kun Xu¹, Linfeng Song¹, Yansong Feng², Yan Song³, Dong Yu¹

¹Tencent AI Lab, Seattle, US

²Peking University, Beijing, China

³Sinovation Ventures

Abstract

Existing entity alignment methods mainly vary on the choices of encoding the knowledge graph, but they typically use the same decoding method, which independently chooses the local optimal match for each source entity. This decoding method may not only cause the “many-to-one” problem but also neglect the coordinated nature of this task, that is, each alignment decision may highly correlate to the other decisions. In this paper, we introduce two coordinated reasoning methods, i.e., the Easy-to-Hard decoding strategy and joint entity alignment algorithm. Specifically, the Easy-to-Hard strategy first retrieves the model-confident alignments from the predicted results and then incorporates them as additional knowledge to resolve the remaining model-uncertain alignments. To achieve this, we further propose an enhanced alignment model that is built on the current state-of-the-art baseline. In addition, to address the many-to-one problem, we propose to jointly predict entity alignments so that the one-to-one constraint can be naturally incorporated into the alignment prediction. Experimental results show that our model achieves the state-of-the-art performance and our reasoning methods can also significantly improve existing baselines.

Introduction

Knowledge graphs (KGs), such as Freebase (Bollacker et al., 2008) and DBpedia (Auer et al., 2007), represent world-level factoid information of entities and their relations in a graph-based format. They have been successfully used in many natural language processing applications, such as question answering (Berant et al., 2013; Bao et al., 2014; Yih et al., 2015; Xu et al., 2016; Das et al., 2017) and relation extraction (Mintz et al., 2009; Hoffmann et al., 2011; Min et al., 2013; Zeng et al., 2015). To date, there have been many KGs in different languages, with each being created in one language (Franco-Salvador, Rosso, and Montesy Gómez, 2016). They share lots of the same facts, and each also provides rich additional information that the others do not cover. Thus, it is very beneficial to establish the cross-lingual alignments between KGs, so that the combined KG can provide richer knowledge for downstream tasks. There-




Figure 1: A challenging entity matching example.

fore, the cross-lingual KG alignment task, which automatically matches entities between multilingual KGs, is proposed to address this problem.

Most recently, several approaches based on cross-lingual entity embeddings (Hao et al., 2016; Chen et al., 2017; Sun, Hu, and Li, 2017) or graph neural networks (Wang et al., 2018; Xu et al., 2019; Wu et al., 2019) have been proposed for this task. In particular, Xu et al. (2019) introduces the topic entity graph to capture the local context information of an entity within the KG, and further tackles this task as a graph matching problem by proposing a graph matching network. This work significantly advanced the state-of-the-art accuracies across several datasets.

Despite the excitingly progressive results that have been shown, all previous works fail to consider the coordinated nature of this task, that is, each alignment decision may highly correlate to the other decisions. For example, all existing models independently align each source entity, which may result in the many-to-one mapping, i.e., more than one source entities are aligned to the same target entity. In particular, we analyze the results of Xu et al. (2019) and find that nearly 8% of the alignments are many-to-one mappings. One intuitive solution is to align these entities in a greedy fashion, that is, assign one alignment at each time with a constraint that all alignments are one-to-one mappings. However, this may introduce the error propagation, since each decision error may propagate to the future decisions. On the other hand, given the fact that the KGs are large, it is also impractical to jointly assign all alignments, due to the massive search space.

We analyze the results of existing alignment baselines and find the second type of errors are caused by the existence of *adversarial* entities that have similar surface strings and KG neighbors with the ground truth. It is challenging for exist-

ing approaches to disambiguate these entities since previous methods mainly rely on the embeddings that are derived by encoding the surface strings and KG neighbors. Figure 1 gives such an example, where it is ambivalent for a model to align 乔治·布什 (George Bush) to “George W. Bush” or “George H. W. Bush”, because both candidates have similar surface strings and share several common neighbors (such as “Republican Party” and “U.S. president”).

In this paper, we propose to alleviate these two types of errors using two coordinated reasoning methods, i.e., the Easy-to-Hard strategy and joint entity alignment algorithm. Specifically, the Easy-to-Hard strategy leverages an iterative approach, where the most model-confident (**easy**) alignments predicted in the previous iteration are provided as additional inputs to the current iteration for resolving the remaining model-uncertain (**hard**) alignments. This idea is motivated by our observation that the model-confident alignments are mostly correct, and thus they can provide reliable clues for other decisions with less model confidence.

To address the many-to-one problem, we propose a joint entity alignment algorithm that finds the global optimal entity alignments that satisfy the one-to-one constraint. This problem is essentially a fundamental combinatorial optimization problem whose exact solution can be found by the Hungarian algorithm (Kuhn, 1955). However, since this algorithm takes a high time complexity of $O(N^4)$ for KGs of N nodes, it is impractical to apply this algorithm in our framework directly. To address this, we propose a simple yet effective solution that breaks down the whole search space into small isolated pieces, so that each piece could be efficiently solved with the Hungarian algorithm. Experiments on the benchmark datasets show that our proposed coordinated reasoning methods can not only improve the current state-of-the-art performance but also significantly boost the performance of previous approaches.

Related Work

Our work is mainly related to two lines of research: network embedding and knowledge graph alignment.

Graph Convolutional Networks

Recently, there has been an increasing interest in extending neural networks to deal with graphs. Defferrard, Bresson, and Vandergheynst (2016) proposed a spectral graph theoretical formulation of CNNs on graphs and a convolutional network extending the conventional CNNs to non-Euclidean space. Kipf and Welling (2017) further extended this idea and proposed graph convolutional neural networks (GCNs) to integrate the connectivity patterns and feature attributes of graph-structured data, and achieved decent results in semi-supervised classification. Thereafter, a series of improvements and extensions were proposed based on GCN. GAT (Veličković et al., 2017) employs the attention mechanism to GCNs, in which each node gets an importance score based on its neighborhood, thus providing more expressive representations for nodes. Furthermore, the R-GCNs (Schlichtkrull et al., 2018) are proposed to model relational data and have been successfully exploited in link

prediction and entity classification. Inspired by the capability of GCNs on learning node representations, we employ the GCN to build our entity alignment framework.

Entity Alignment

Earliest approaches of entity alignment usually require expensive expert efforts to design model features (Mahdisoltani, Biega, and Suchanek, 2013). Recently, embedding based methods have been proposed to address this issue. MTransE (Chen et al., 2017) employs TransE (Bordes et al., 2013) to embed entities and relations of each knowledge graph in a separate space, and then provides five different variants of transformation functions to project the embedded vectors from one subspace to another. The candidate set of one entity’s correspondence in the other knowledge graph can be obtained by ranking the distance between them in the transformed space. ITransE (Zhu et al., 2017) utilizes TransE to learn one common low-dimensional subspace for all knowledge graphs, with the constraint that the observed anchor seeds from different knowledge graphs share the same vector representation in the subspace. AlignE (Sun, Hu, and Li, 2017) also adopts TransE to learn network embeddings, and applies parameter swapping to encode network into a unified space. NTAM (Li et al., 2018) utilizes a probabilistic model for the alignment task. Instead of using TransE to derive entity embeddings from the knowledge graph, various GCN based methods (Wang et al., 2018; Ye et al., 2019; Wu et al., 2019) that use the conventional GCN to encode the entities and relations have been proposed to perform the alignment. Different with those methods that still follow previous works that rely on learned entity embeddings to rank alignments, Xu et al. (2019) views this task as a graph matching problem and further proposes a graph matching neural network that additionally considers the matching information of an entity’s neighborhood to perform the prediction.

Despite these approaches achieve progressive results, all current works focus on encoding the entities and relations, while neglecting the fact that the decoding strategy may have a considerable impact over the final performance. In this paper, we explore the coordinated nature of this task and propose two types of reasoning methods to improve the performance of these baselines.

Problem Formulation

Formally, a KG is represented as $G = (E, R, T)$, where E , R , T are the sets of entities, relations, and triples, respectively. Let $G_1 = (E_1, R_1, T_1)$ and $G_2 = (E_2, R_2, T_2)$ be two heterogeneous KGs to be aligned. That is, an entity in G_1 (source entity) may have its counterpart in G_2 (target entity) in a different language or different surface names. As a starting point, we can collect a small number of equivalent entity pairs between G_1 and G_2 as the alignment seeds. We define the entity alignment task as automatically finding more equivalent entities using the alignment seeds as training data.

Coordinated Reasoning

All existing works follow the conventional framework that first encodes the context information of the source entity within the KG into a distributional representation and then ranks the candidate target entities according to the representation similarities. These works may differ in the choice of the encoder, such as TransE or GCN, but all of them utilize the same decoding method, which simply picks the local optimal candidate for each source entity without considering the global alignment coherence. For example, more than one source entities may be aligned to the same target entity, causing the *many-to-one* problem. This simple decoding strategy also neglects the coordinated nature of this task, that is, previously predicted alignments are also helpful to future predictions.

Motivated by these observations, we propose two types of coordinated reasoning methods. First, to address the many-to-one problem, we jointly predict alignments by explicitly incorporating the one-to-one constraint into the decoding. Second, we propose a new Easy-to-Hard decoding strategy that first resolves the most model-confident alignments and then uses them as additional evidence to better handle the model-uncertain alignments.

Easy-to-Hard Decoding

All existing models independently predict alignments for source entities while neglecting the fact that the decoding strategy may have a significant impact over the performance. Figure 1 illustrates such an example where the goal is to align 乔治·布什 (George Bush) from the Chinese KG into the English KG. Given its two candidates, i.e., *George W. Bush* and *George H. W. Bush*, it is challenging for previous methods to find the correct alignment (*George W. Bush*) since these candidates have almost the same neighbors, except that George W. Bush graduated from Harvard University while George H. W. Bush not. On the other hand, we can see that the Chinese KG includes a fact, \langle 乔治·布什 graduated from 哈佛大学 (Harvard University) \rangle , which is strong evidence for aligning 乔治·布什 to George W. Bush. Intuitively, if a model could first align 哈佛大学 to the Harvard University and introduce this as additional knowledge, it could be more easy for the model to find the correct alignment for 乔治·布什. Compared to the alignment for 乔治·布什, which is **Hard** to resolve, the alignment for 哈佛大学 is relatively **Easier**.

Inspired by the above observation, in this paper, we propose a new decoding method, namely Easy-to-Hard strategy, which first attempts to resolve “easy” alignments in the test set and then incorporates them as additional knowledge into the model to better tackle the remaining “hard” alignments. There are two main challenges here. First of all, it is difficult to determine whether an alignment is easy or hard to resolve. Second, existing dominant models are mainly built on the graph neural networks, and it is unclear how to integrate such additional knowledge into their models.

We analyze the alignment results of three baseline methods, i.e., Wang et al. (2018), Xu et al. (2019) and Wu et al. (2019). Interestingly, we find that all these baselines could

achieve at least 99.5% accuracy for those alignments with normalized probabilities over 0.9. This result is coherent with our expectation since a higher probability typically suggests that the model is more confident about the prediction and also indicates that this alignment is easier for the model to resolve. Therefore, we apply the following steps to decode the test set iteratively.

Step	Description
1	Employ an alignment model to predict alignments for all source entities in the test set.
2	Use a predefined probability threshold α to refine those alignments. In particular, assignments with probabilities <i>higher than</i> α are regarded as easy alignments while the others are viewed as hard alignments.
3	If more than K easy alignments are found in Step 2, take these easy alignments as additional knowledge and incorporate them into the alignment model to establish alignments for the remaining entities (go to Step 1); otherwise, return all alignments.

After establishing easy assignments in each decoding step, we need to incorporate them as additional knowledge into the alignment model for the next round decoding. This design heavily depends on alignment model architecture. In this paper, we use the state-of-the-art alignment model (Xu et al., 2019) as our baseline method and propose two ways to enhance this model by incorporating easy assignment information.

Alignment Model Baseline. Xu et al. (2019) utilized a graph (namely *topic graph*) to capture the context information of an entity (namely *topic entity*) within the KG. For instance, Figure 2 gives the topic graphs of George Bush in both the Chinese and English KG. The entity alignment task is then viewed as a graph matching problem, whose goal is to calculate the similarity of these two topic graphs, say G_1 and G_2 . To achieve this, they further propose a neural graph matching model that includes the following four layers:

- **Input Representation Layer.** The goal of this layer is to learn embeddings for entities that occurred in topic entity graphs by using a graph convolution neural network (GCN) (Kipf and Welling, 2017).
- **Node-Level Matching Layer.** This layer is designed to capture local matching information by comparing each entity embedding of one topic entity graph against all entity embeddings of the other graph in both ways (from G_1 to G_2 and G_2 to G_1).
- **Graph-Level Matching Layer.** In this layer, the model applies another GCN to propagate the local matching information throughout the graph. The motivation behind it is that this GCN layer can encode the global matching state between the pairs of whole graphs. The model then feeds these matching representations to a fully-connected neural network and applies the element-wise *max* and *mean* pooling method to generate a fixed-length graph matching representation.

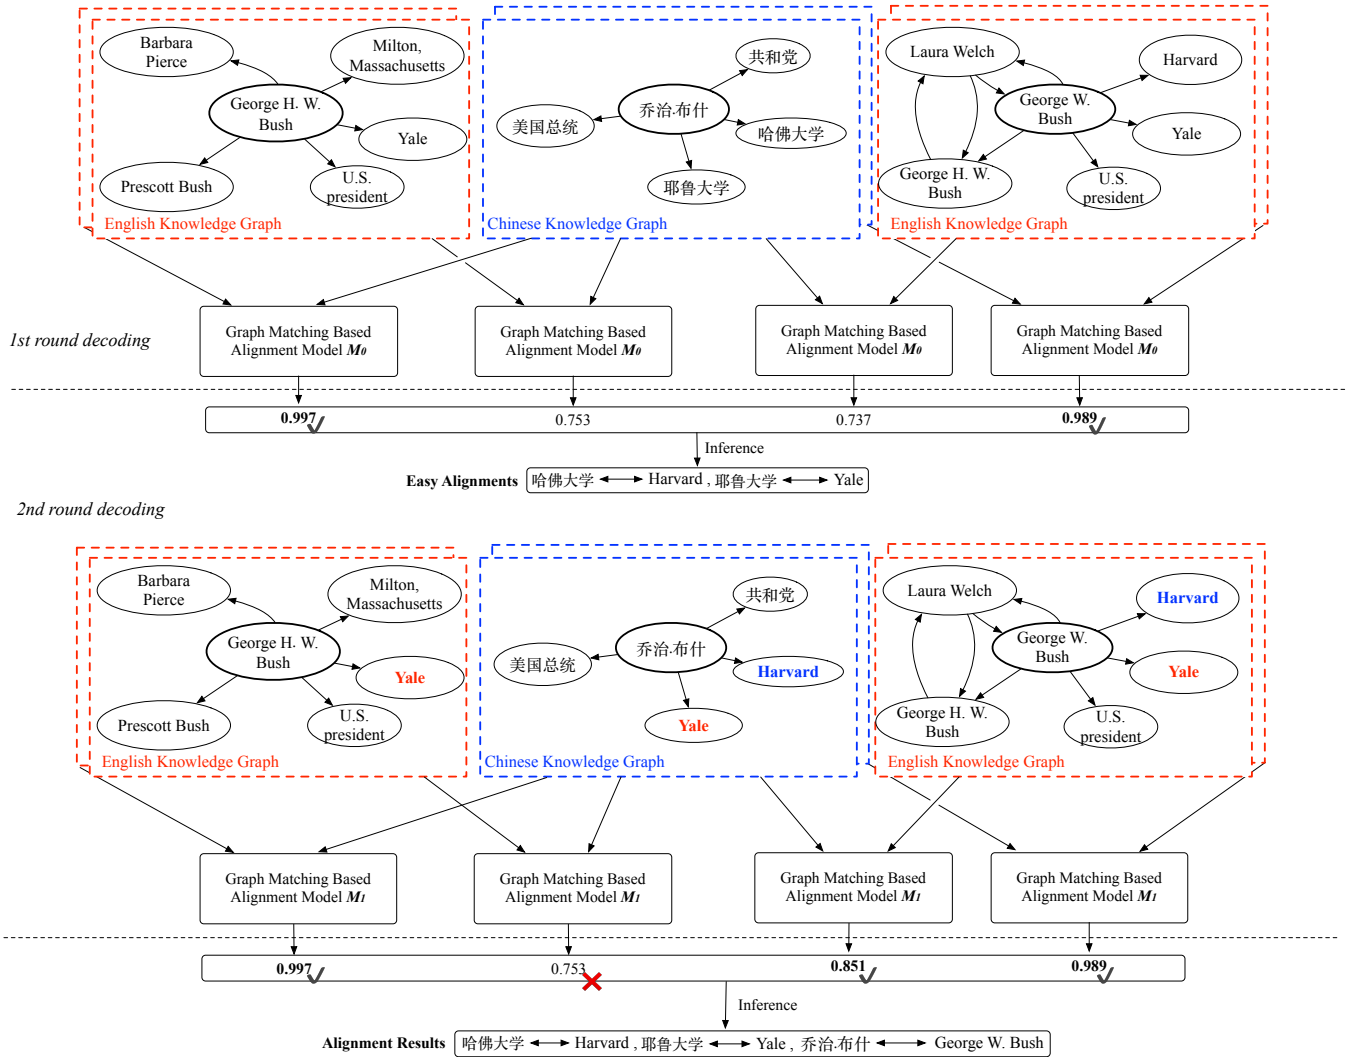


Figure 2: A running example of our Easy-to-Hard decoding strategy for aligning *George Bush* in the English and Chinese knowledge graph. After the first round of decoding, the baseline model aligns 哈佛大学 to Harvard and 耶鲁大学 to Yale, because their probabilities predicted by M_0 is higher than α . After introducing these information, our enhanced model M_1 increased the probability of aligning 乔治·布什 to George W. Bush while decreasing the probability of its alignment to George H. W. Bush.

- **Prediction Layer.** The model finally uses a two-layer feed-forward neural network to consume the fixed-length graph matching representation and applies the *softmax* function in the output layer.

Our Model. In contrast to Xu et al. (2019) that only takes two topic graphs as input, we can utilize additional information such as easy assignments found in previous decoding steps to resolve hard assignments. In particular, we introduce two ways to enhance this baseline model by explicitly integrating the easy assignment information into two layers of Xu et al. (2019):

- **Enhanced Input Representation Layer.** In this layer, Xu et al. (2019) utilizes a GCN to learn entity embeddings from the topic graph, where the entity surface form has been proved to be a key feature in deriving their embeddings. Therefore, we require that the aligned entities

found in the easy alignments should have the same surface forms so that they could share the common embeddings. For example, in Figure 2, after the first round of decoding, 哈佛大学 (Harvard University) is aligned to Harvard, we then change the surface form of “哈佛大学” to “Harvard” in the second decoding step.

- **Enhanced Node-Level Matching Layer.** As concluded in Xu et al. (2019), the node-level matching layer has a significant impact on the matching performance, since it captures the local entity matching information. In the baseline model, the entity similarities are calculated based on the entity embeddings derived from the first GCN layer. Although in the enhanced input representation layer the aligned entities have the same surface forms, it can still not guarantee that their embeddings are close. It is because the first GCN layer is supposed to encode not only

the surface form but also the structural information into their representations. Therefore, we explicitly incorporate the easy alignment information into this layer by enforcing that the normalized similarities between the aligned entities to be 1.0. Then, we feed the revised entity similarities to the graph-level matching layer and the final prediction layer.

Notice that, in practice, there are two possible options to build the enhanced alignment model in our framework. First, we can directly use the pre-trained baseline but replace its first two layers with our proposed enhanced layers. Because we do not modify the model architecture, no more parameters are needed to be learned. The second way is to train a new enhanced alignment model with randomly sampled alignments as simulated easy alignments. The motivation behind is that given more easy alignments, the model could more focus on learning to disambiguate hard alignments. Experimental results show that the latter achieves much better performance. We will discuss these two options in the experiment section.

Joint Entity Alignment

As shown in Figure 3(a), our model typically outputs a 2-dimensional matrix of probabilities after decoding, where each cell item (such as $p(e_t|e_s)$) represents the likelihood of aligning source entity e_s to target entity e_t . The goal of the entity alignment task is then equal to find the best solution (a set of one-to-one alignments) with the highest probability:

$$\arg \max_{\mathbf{A}} \prod_{(e_s, e_t) \in \mathbf{A}} p(e_t|e_s) \quad (1)$$

where \mathbf{A} represents one solution. Since knowledge graphs are usually huge, this problem cannot be solved by naive enumeration, which takes $O(N!)$ time for KGs with N entities. Existing works choose the optimal local match for each source entity while neglecting the *one-to-one* nature, and as a result, multiple source entities may be mapped to one target entity.

Here, for the first time, we propose to explicitly incorporate this one-to-one constraint into the alignment prediction. To achieve this, we first reformat the goal from maximizing the product of probabilities (Equation 1) to minimizing the sum of negative log-likelihoods.

$$\arg \min_{\mathbf{A}} \sum_{(e_s, e_t) \in \mathbf{A}} -\log p(e_t|e_s) \quad (2)$$

As a result, the entity alignment problem is equivalently converted to the well-studied “task assignment” problem¹, where each agent/task is assigned to exactly one task/agent, and each agent-task assignment has a fixed cost that does not depend on the other assignments.

The Hungarian algorithm (Kuhn, 1955) has been proven to be efficient for finding the best solution for this problem. It takes a cost matrix as input, which can be easily achieved by padding rows or columns of a constant value for the non-square matrix. For a brief introduction, the algorithm takes

¹https://en.wikipedia.org/wiki/Assignment_problem




Figure 3: (a) The original alignment results between entities $\{A, B, C\} \in G_1$ and $\{1, 2, 3\} \in G_2$, where the thickness of a line represents its alignment probability and very weak alignments are shown as dotted lines; (b) The sub-spaces after separation.

the following four main steps for the cost matrix with $N \times N$ elements, where the last two steps repeat until a solution is found.² It is guaranteed that a solution could be found within $O(N^4)$ time.

Step	Description
1	Find the lowest item for each row and subtract it from the others in that row.
2	Similarly, find the lowest item for each column and subtract it from the others in that column.
3	Cover all zeros in the resulting matrix using a minimum number of horizontal and vertical lines. If <i>less than</i> N lines are required, go to Step 4; otherwise, a solution is found.
4	Find the smallest item v not covered by any line in Step 3. Subtract v from all uncovered items, and add v to all items covered by two lines. Go to Step 3.

One can see that naively applying Hungarian is impractical, as it still takes $O(N^4)$ computation time for matching two KGs of N nodes. To further decrease the time consumption, we break the whole search space into many isolated sub-spaces, where each sub-space contains only a subset of source and target entities for making alignments. Specifically, we discard the candidate alignments with a probability lower than a predefined threshold τ from the original search space. Based on this, we define two source entities being connected only if they share common candidates in the target. Doing in this way fits the intuition where a large KG usually contains many domains, such as politics, sports and science, and only the entities within each domain have densely interacted. Our experiments show that τ has little effect on performance, while it dramatically reduces the search time.

Figure 3 illustrates the search space separation process, where thin and dotted lines correspond to low-confident alignments. After dropping out these alignments with low model scores, the whole search space is split into two independent sub-spaces, as shown in Figure 3(b). Here (A) and (B) are in the same sub-space, as they share the same target can-

²<http://www.hungarianalgorithm.com/> provides a detailed explanation and an online demo.

Datasets		Entities	Relations	Triples
DBP15K _{ZH-EN}	Chinese	66,469	2,830	153,929
	English	98,125	2,317	237,674
DBP15K _{JA-EN}	Japanese	65,744	2,043	164,373
	English	95,680	2,096	233,319
DBP15K _{FR-EN}	French	66,858	1,379	192,191
	English	105,889	2,209	278,590

Table 1: Dataset summary.

didate ①. Removed connections (such as ④ to ②) are considered as infinite cost. As the next step, each sub-space is solved with the Hungarian algorithm, before their results are combined to form our final outputs.

Experimental Setup

Datasets. We evaluate our approach on three large-scale cross-lingual datasets from DBP15K (Sun, Hu, and Li, 2017). These datasets are built upon Chinese, English, Japanese and French versions of DBpedia (Auer et al., 2007). Each dataset contains 15,000 inter-language links connecting equivalent entities in two KGs of different languages. We use the same training/testing split as previous works, 30% for training, and 70% for testing. Table 1 lists their statistical summaries.

Evaluation Metrics. Like previous works, we use *Hits@1* to evaluate our model, where a *Hits@1* score (higher is better) is computed by measuring the proportion of correctly aligned entities ranked in the top one.

Comparison Models. We compare our approach against existing alignment methods: JE (Hao et al., 2016), MTransE (Chen et al., 2017), JAPE (Sun, Hu, and Li, 2017), IPTransE (Zhu et al., 2017), BootEA (Sun, Hu, and Li, 2017), GCN (Wang et al., 2018), GM (Xu et al., 2019) and RDGCN (Wu et al., 2019).

Model Variants. To evaluate different reasoning methods, we provide three implementation variants of our model for ablation studies, including (1) *X*-EHD: the baseline model *X* that only uses our proposed **Easy-to-Hard Decoding** strategy; (2) *X*-JEA: the baseline model *X* that only uses our proposed **Joint Entity Alignment** method; (3) *X*-EHD-JEA: the baseline model *X* that uses both of these two reasoning methods.

Implementation details. For the configurations of the alignment model, we use the same settings as Xu et al. (2019). Specifically, we use the Adam optimizer (Kingma and Ba, 2014) to update parameters with mini-batch size 32. The learning rate is set to 0.001. The hop size of two GCN layers is set to 2 and 3, respectively. Following Wu et al. (2019), we use Google Translate to translate Chinese, Japanese, and French entity names into English, and then use Glove embeddings (Pennington, Socher, and Manning, 2014) to construct the initial entity representations in the model. For all datasets, we first use the baseline model to

Method	ZH-EN	JA-EN	FR-EN
JE	21.27	18.92	15.38
MTransE	30.83	27.86	24.41
JAPE	41.18	36.25	32.39
IPTransE	40.59	36.69	33.30
GCN	41.25	39.9	37.29
BootEA	62.94	62.23	65.30
GM	67.93	73.97	89.38
RDGCN	70.75	76.74	88.64
GCN-JEA	43.43	45.00	39.78
BootEA-JEA	64.56	64.17	69.31
RDGCN-JEA	72.03	77.56	90.49
GCN-EHD	44.37	41.72	39.09
BootEA-EHD	65.27	65.36	68.92
RDGCN-EHD	71.15	77.07	91.01
GM-EHD	70.31	77.92	90.49
GM-JEA	72.05	78.73	91.08
GM-EHD-JEA	73.58	79.15	92.43

Table 2: Evaluation results on the datasets.

retrieve the top 10 alignments, normalize their scores as probabilities and then perform the proposed coordinated reasoning methods over them. For the Easy-to-Hard decoding method, α is set to 0.75, and K is set to 20. For the joint entity alignment, τ is set to 0.10. For training the enhanced alignment model, for each topic graph pair, we randomly choose at most two gold alignments from the ground truth as the simulated easy alignments.

Results and Discussion

Main Results

Table 2 shows the performance of all compared approaches on the evaluation datasets. We can see that both of the Easy-to-Hard decoding strategy (referred as EHD in Table 2) and the joint entity alignment method (referred as JEA in Table 2) could significantly improve the performance of GM. When these two methods are combined, the overall performance is further improved, outperforming previous works. We also investigate whether our proposed reasoning methods could also boost existing baselines. From Table 2, we can see also that the joint entity alignment method could also improve the performance of GCN, BootEA and RDGCN, indicating that our method is able to avoid the *many-to-one* problem effectively. Recall that, the Easy-to-Hard decoding method requires an enhanced alignment model that could integrate the easy alignment information. Since designing enhanced versions for these baselines is beyond our goal, here we only enforce that that the aligned entities found in the easy alignments have the same surface form. We find that this simplified strategy could still improve these baselines, which also suggests that our proposed decoding strategy is generally helpful to the alignment models.

Discussion

Let us first look at the impacts of alignment-dropping threshold τ to both the performance and running time for our joint entity alignment algorithm. From Table 3, we can see that

τ	Max sub-space	Time	FR-EN(hit@1)
0.05	5238	–	–
0.10	1562	24m34s	91.02
0.15	116	39s	90.90
0.20	<100	38s	90.78

Table 3: Performance and computation time for different τ values, where *Max sub-space* shows the number of source nodes in the largest sub-space. Baseline accuracy is 89.38.

α	Decoding Rounds	ZH-EN	JA-EN	FR-EN
–	–	66.29	72.31	88.07
0.95	2	69.05	74.33	88.25
0.85	5	71.71	75.10	90.31
0.75	10	72.09	76.62	91.18
0.65	20	67.12	72.15	88.60

Table 4: Hit@1 accuracies and decoding rounds on the development set for different α values. The first row lists the accuracies of the GM baseline.

decreasing τ can slightly improve the performance but with a huge cost of computation time. For example, when τ is changed from 0.15 to 0.10, the accuracy could increase by 0.12% but the computation time dramatically increases from 39s to almost 25 minutes. Moreover, if τ is set to 0.05, we cannot even get the results. As shown in Table 3, in order to better understand why the running time changes, we additionally analyze the size of the largest sub-space. We find that the size of the maximal sub-space under $\tau = 0.05$ is 3 times more than the size under $\tau = 0.10$, thus the running time under $\tau = 0.05$ is expected to be roughly 32 hours, which is 81 (3^4) times than the time under $\tau = 0.10$. The running time does not significantly change when increasing τ from 0.15 to 0.20, because the Hungarian algorithm does not take much time for this situation, and the most time consumption is data processing.

We also investigated the impact of the probability threshold α on the performance for our Easy-to-Hard decoding method. We experimented with different α values and evaluated our model on the development set of the DBP15K. Table 4 reports hit@1 accuracies on these datasets. We can see that our model could benefit from decreasing α until it reaches 0.75. It is expected to find that lower α may hurt the performance since it incorporates some incorrect predictions as easy (gold) alignments into the model. Recall that in our decoding algorithm, we continuously perform the inference until less than K new easy alignments are found in the previous round. As shown in Table 4, we observed that decreasing α not only achieves worse performance but also requires more converge rounds. To better understand why more converge rounds are required, we analyzed the intermediate established alignments during the inference. We find this is due to those incorrect alignments introduced by reducing α produce a chain reaction, which offers the model more confidence about some uncertain but incorrect alignments, resulting in more decoding rounds.

Recall that there are two options to build the enhanced

alignment model, where the first one directly replaces two layers of a pre-trained GM model with our proposed enhanced layers while keeping the parameters the same; the second one trains a new GM model with simulated easy alignments. We evaluate these two options on several datasets and observe that both of these two ways could improve the performance but the model could gain more performance improvement from the second way. We further manually analyze the predicted alignments of these two options and find that the new trained GM model could resolve more ambiguous (hard) alignments. We think this is due to that introducing the simulated easy alignments into the training phase could allow the model to learn how to properly utilize these additional evidence to disambiguate the hard alignments.

Here one natural question is how many simulated easy alignments are needed for training the new GM model. In experiments, we find that using two simulated easy alignments to train the model could achieve the best performance; introducing more easy alignments to train the model could not further improve the results. However, this observation is in conflict with our intuition, that is, more easy alignment information could better help the model to disambiguate those uncertain predictions. By analyzing the entities in the test set, we find this is due to that among these entities, at most three entities co-occur in the same topic graphs, and consequently, during the decoding, the model could only introduce at most two easy alignments. Motivated by this observation, we conducted an additional experiment that predicts alignments for all entities in the KGs except the training seeds. We find that our reasoning methods could achieve more performance improvement, and considering more than two easy alignments into the training also further improves the overall performance as we expected. Note that, although this experiment may consume almost 5 times more than the original decoding time, we believe that some optimization could be adopted to reduce the time complexity, which we leave for the future work.

Conclusion

Previous entity alignment methods mainly use the same decoding strategy that independently chooses the optimal local match for each source entity without considering the global alignment coherence, thereby may cause the many-to-one problem. To address this, we propose two reasoning method, including a new Easy-to-Hard decoding strategy and joint entity alignment method. Specifically, the Easy-to-Hard decoding method iteratively decodes the test set by taking the most model-confident alignments predicted in the previous iteration as additional inputs to the current iteration for resolving the model-uncertain alignments. The joint entity alignment method views the entity alignment as the task assignment problem and employs the Hungarian algorithm to guarantee the predicted alignments are one-to-one mappings. Experimental results on the DBP15K dataset show that our reasoning methods are general to these baselines and can significantly improve their performance.

References

- Auer, S.; Bizer, C.; Kobilarov, G.; Lehmann, J.; Cyganiak, R.; and Ives, Z. 2007. Dbpedia: A nucleus for a web of open data. In *The semantic web*. Springer, 722–735.
- Bao, J.; Duan, N.; Zhou, M.; and Zhao, T. 2014. Knowledge-based question answering as machine translation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 967–976.
- Berant, J.; Chou, A.; Frostig, R.; and Liang, P. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 1533–1544.
- Bollacker, K.; Evans, C.; Paritosh, P.; Sturge, T.; and Taylor, J. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, 1247–1250. AcM.
- Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, 2787–2795.
- Chen, M.; Tian, Y.; Yang, M.; and Zaniolo, C. 2017. Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 1511–1517.
- Das, R.; Zaheer, M.; Reddy, S.; and McCallum, A. 2017. Question answering on knowledge bases and text using universal schema and memory networks. *CoRR* abs/1704.08384.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, 3844–3852.
- Franco-Salvador, M.; Rosso, P.; and Montes-y Gómez, M. 2016. A systematic study of knowledge graph analysis for cross-language plagiarism detection. *Information Processing & Management* 52(4):550–570.
- Hao, Y.; Zhang, Y.; He, S.; Liu, K.; and Zhao, J. 2016. A joint embedding method for entity alignment of knowledge bases. In *China Conference on Knowledge Graph and Semantic Computing*, 3–14.
- Hoffmann, R.; Zhang, C.; Ling, X.; Zettlemoyer, L.; and Weld, D. S. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, 541–550. Association for Computational Linguistics.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *CoRR* abs/1412.6980.
- Kipf, T. N., and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Kuhn, H. W. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly* 2(1-2):83–97.
- Li, S.; Li, X.; Ye, R.; Wang, M.; Su, H.; and Ou, Y. 2018. Non-translational alignment for multi-relational networks. In *IJCAI*, 4180–4186.
- Mahdisoltani, F.; Biega, J.; and Suchanek, F. M. 2013. Yago3: A knowledge base from multilingual wikipedias.
- Min, B.; Grishman, R.; Wan, L.; Wang, C.; and Gondek, D. 2013. Distant supervision for relation extraction with an incomplete knowledge base. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 777–782.
- Mintz, M.; Bills, S.; Snow, R.; and Jurafsky, D. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, 1003–1011. Association for Computational Linguistics.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; Van Den Berg, R.; Titov, I.; and Welling, M. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, 593–607. Springer.
- Sun, Z.; Hu, W.; and Li, C. 2017. Cross-lingual entity alignment via joint attribute-preserving embedding. In *Proceedings of the 16th International Semantic Web Conference*.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Wang, Z.; Lv, Q.; Lan, X.; and Zhang, Y. 2018. Cross-lingual knowledge graph alignment via graph convolutional networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 349–357.
- Wu, Y.; Liu, X.; Feng, Y.; Wang, Z.; Yan, R.; and Zhao, D. 2019. Relation-aware entity alignment for heterogeneous knowledge graphs.
- Xu, K.; Reddy, S.; Feng, Y.; Huang, S.; and Zhao, D. 2016. Question answering on freebase via relation extraction and textual evidence. In *ACL 2016*.
- Xu, K.; Wang, L.; Yu, M.; Feng, Y.; Song, Y.; Wang, Z.; and Yu, D. 2019. Cross-lingual knowledge graph alignment via graph matching neural network. In *ACL 2019*, 3156–3161.
- Ye, R.; Li, X.; Fang, Y.; Zang, H.; and Wang, M. 2019. A vectorized relational graph convolutional network for multi-relational network alignment. 4135–4141.
- Yih, W.-t.; Chang, M.-W.; He, X.; and Gao, J. 2015. Semantic parsing via staged query graph generation: Question answering with knowledge base. In *ACL*, 1321–1331.
- Zeng, D.; Liu, K.; Chen, Y.; and Zhao, J. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *EMNLP*, 1753–1762.
- Zhu, H.; Xie, R.; Liu, Z.; and Sun, M. 2017. Iterative entity alignment via joint knowledge embeddings. In *IJCAI*, 4258–4264.