

Cross-lingual Entity Alignment with Incidental Supervision

Muhao Chen^{1,2*}, Weijia Shi^{3*}, Ben Zhou¹, Dan Roth¹

¹Department of Computer and Information Science, UPenn

²Viterbi School of Engineering, USC

³Department of Computer Science, UCLA

{muhao, xyzhou, danroth}@seas.upenn.edu; swj0419@cs.ucla.edu

Abstract

Much research effort has been put to multilingual knowledge graph (KG) embedding methods to address the entity alignment task, which seeks to match entities in different language-specific KGs that refer to the same real-world object. Such methods are often hindered by the insufficiency of seed alignment provided between KGs. Therefore, we propose an incidentally supervised model, JEANS [\[1\]](#), which jointly represents multilingual KGs and text corpora in a shared embedding scheme, and seeks to improve entity alignment with incidental supervision signals from text. JEANS first deploys an entity grounding process to combine each KG with the monolingual text corpus. Then, two learning processes are conducted: (i) an *embedding learning process* to encode the KG and text of each language in one embedding space, and (ii) a self-learning based *alignment learning process* to iteratively induce the matching of entities and that of lexemes between embeddings. Experiments on benchmark datasets show that JEANS leads to promising improvement on entity alignment with incidental supervision, and significantly outperforms state-of-the-art methods that solely rely on internal information of KGs.¹

1 Introduction

A multilingual knowledge base (KB) such as DBpedia ([Lehmann et al., 2015](#)), ConceptNet ([Speer et al., 2017](#)) and Yago ([Mahdisoltani et al., 2015](#)) stores multiple language-specific knowledge graphs (KGs) that express relations of many concepts and real-world entities. As each KG thereof is either extracted independently from monolingual corpora ([Lehmann et al., 2015](#); [Mahdisoltani et al., 2015](#)) or contributed by speakers of the language ([Speer et al., 2017](#);

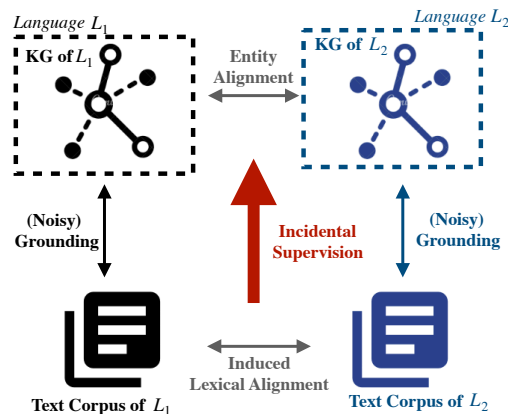


Figure 1: The learning framework of JEANS.

[Mitchell et al., 2018](#)), it is common for different KGs to constitute complementary knowledge ([Bleiholder and Naumann, 2009](#); [Bryl and Bizer, 2014](#)). Hence, aligning and synchronizing language-specific KGs support AI systems with more comprehensive commonsense reasoning ([Lin et al., 2019](#); [Li et al., 2019b](#); [Yeo et al., 2018](#)), and benefit various knowledge-driven NLP tasks, including machine translation ([Moussallem et al., 2018](#)), narrative prediction ([Chen et al., 2019](#)) and dialogue agents ([Sun et al., 2019a](#)).

Learning to align multilingual KGs is a non-trivial task, as KGs with distinct surface forms, heterogeneous schemata and inconsistent structures easily cause traditional symbolic methods to fall short ([Suchanek et al., 2011](#); [Wijaya et al., 2013](#); [Jiménez-Ruiz et al., 2012](#)). Recently, much attention has been paid to methods based on multilingual KG embeddings ([Chen et al., 2017a,b, 2018](#); [Sun et al., 2017, 2018, 2019b](#); [Zhang et al., 2019](#)). Those methods seek to separately encode the structure of each language-specific KG in an embedding space. Then, based on some seed entity alignment, the entity counterparts in different KGs can be easily matched via distances or transformations of embedding vectors. The principle is

* Indicating equal contributions.

¹Software and resources are available at http://cogcomp.org/page/publication_view/929.

that entities with relevant neighborhood information can be characterized with similar embedding representations. Such representations particularly are tolerant to the aforementioned heterogeneity of surface forms and schemata in language-specific KGs (Chen et al., 2017a; Sun et al., 2018, 2020a).

While multilingual KG embeddings provide a general and tractable way to align KGs, it still remains challenging for related methods to precisely infer the cross-lingual correspondence of entities. The challenge is that the seed entity alignment, which serves as the essential training data to learn the connection between language-specific KG embeddings, is often limitedly provided in KBs (Chen et al., 2018; Sun et al., 2018). Hence, the lack of supervision often hinders the precision of inferred entity counterparts, and affects even more significantly when KGs scale up and become inconsistent in contents and density (Pujara et al., 2017). Several methods also gain auxiliary supervision from profile information of entities, including descriptions (Chen et al., 2018; Yang et al., 2019; Zhang et al., 2019) and numerical attributes (Sun et al., 2017; Trsedya et al., 2019). However, such profile information is not available in many KGs (Speer et al., 2017; Mitchell et al., 2018; Bond and Foster, 2013), therefore causing these methods to be not generally applicable.

Unlike existing models that rely on internal information of KGs, we seek to create embeddings that incorporate both KGs and freely available text corpora, and exploit incidental supervision signals (Roth, 2017) from text corpora to enhance the alignment learning on KGs (Figure 1). In this paper, we propose a novel embedding model JEANS (Joint Embedding Based Entity Alignment with INcidental Supervision; ¶). Particularly, JEANS first performs a *grounding process* (Gupta et al., 2017; Upadhyay et al., 2018) to link entity mentions in each monolingual text corpus to the KG of the same language. Based on the KGs and grounded text in a pair of languages, JEANS conducts two learning processes, i.e. *embedding learning* and *alignment learning*. The embedding learning process distributes entities, relations and lexemes of each language in its embedding space, in which a KG embedding model and a language model for that language are jointly trained. This process seeks to leverage text contexts to help capture the proximity of entities. On top of that, alignment learning

captures cross-lingual correspondence for entities and lexemes in a self-learning manner (Artetxe et al., 2018). Starting from a small amount of seed entity alignment, this process iteratively induces a transformation between language-specific embedding spaces, and infers more alignment of entity and lexemes at each iteration to improve the learning at the next one. Moreover, we also employ the closed-form Procrustes solution (Conneau et al., 2018) to strengthen alignment induction within each iteration. Experimental results on two benchmarks confirm the effectiveness of JEANS in leveraging incidental supervision, leading to significant improvement to entity alignment and drastically outperforming existing methods.

2 Related Work

We discuss relevant works in four topics. Each of them has a large body of work which we can only provide as a highly selected summary.

Entity alignment. Entity alignment in KBs has been a long-standing problem (Shvaiko and Euzenat, 2011). Aside from earlier approaches based on symbolic or schematic similarity of entities (Suchanek et al., 2011; Wijaya et al., 2013; Jiménez-Ruiz et al., 2012), more recent research addresses this task with multilingual KG embeddings. A representative method of such is MTransE (Chen et al., 2017a). MTransE jointly learns two model components. There are a translational embedding model (Bordes et al., 2013) that distributes the facts in language-specific KGs into separate embeddings, and a transformation-based alignment model that maps between entity counterparts across embedding spaces.

Following the general principle of MTransE, later approaches are developed through the following three lines. One is to incorporate various *embedding learning techniques* for KGs. Besides translational techniques, some models employ alternative *relation modeling* techniques to encode relation facts, such as circular correlation (Nickel et al., 2016), Hadamard product (Hao et al., 2019) and recurrent skipping networks (Guo et al., 2019). Others encode entities with *neighborhood aggregation* techniques, including GCN (Wang et al., 2018; Yang et al., 2019; Cao et al., 2019; Xu et al., 2019; Wu et al., 2019b), RGCN (Wu et al., 2019a) and GAT (Zhu et al., 2019). Their benefits are mainly to produce entity representations

capturing high-order proximity, so as to better suit the alignment task. A few works follow the second line to enhance the *alignment learning* with semi-supervised learning techniques. Representative ones include co-training (Chen et al., 2018), optimal transport (Pei et al., 2019b) and bootstrapping (Sun et al., 2018), which improve the preciseness of alignment captured with limited supervision. The third line of research seeks to obtain additional supervision from entity profiles, including descriptions (Chen et al., 2018; Yang et al., 2019), attributes (Sun et al., 2017; Trsedya et al., 2019; Pei et al., 2019a; Yang et al., 2020) and KG schemata (Zhang et al., 2019). While those alternative views of entities can effectively bridge the embeddings, the limitation of such methods lies in the unavailability of those views in many KGs (Speer et al., 2017; Mitchell et al., 2018; Bond and Foster, 2013). A survey on the entity alignment problem by Sun et al. (2020b) has provided a more comprehensive summarization of recent advances in these lines.

Our method is mainly related to the third line of research. While instead of leveraging specific intra-KB information, our method introduces supervision signals from text contexts that are freely accessible to almost any KBs with the aid of grounding techniques. Meanwhile, our paper also follows the second line to improve alignment learning techniques, and couples two mainstream techniques for embedding learning.

Joint embeddings of entities and text. Fewer efforts have been put to jointly characterize entities and text as embeddings. Wang et al. (2014b) propose to connect a translational embedding of Freebase (Bollacker et al., 2008) to a English word embedding based on Wikipedia anchors, therefore providing a joint embedding to enhance link prediction in the KG. Zhong et al. (2015) generalize the approach by Wang et al. (2014b) with distant supervision based on entity descriptions and text corpora. Toutanova et al. (2015) extract dependency paths from sentences and jointly embed them with a KG using DistMult (Yang et al., 2015) to support the relation extraction task. Several other approaches focus on jointly embedding words, entities (Yamada et al., 2017; Newman-Griffis et al., 2018; Cao et al., 2017; Almasian et al., 2019) and entity types (Gupta et al., 2017) appearing in the same textual contexts without considering relational structure of a KG. These ap-

proaches are employed in monolingual NLP tasks including entity linking (Gupta et al., 2017; Cao et al., 2017), entity abstraction (Newman-Griffis et al., 2018) and factoid QA (Yamada et al., 2017). As they focus on a monolingual and supervised scenario, they are essentially different from our goal to help cross-lingual KG alignment with incidental supervision from unparalleled corpora.

Multilingual word embeddings. Our model component of alignment induction from text is closely connected to multilingual word embeddings. Earlier approaches in this line, regardless of being supervised or weakly supervised, based on seed lexicon (Zou et al., 2013) or parallel corpora (Gouws et al., 2015), are systematically summarized in a recent survey (2017). While a number of methods in this line can be employed in our model to gain addition supervision for entity alignment, we choose to use a combination of Procrustes solution (Conneau et al., 2018) with self-learning to offer precise inference of cross-lingual alignment based on limited seed alignment. Note that recent contextualized embeddings such as M-BERT (Pires et al., 2019) and XLM-R (Conneau et al., 2020) do not directly suit our problem setting, since contextualization could cause ambiguity to entity representations, therefore impairing the alignment of entities across embedding spaces.

Incidental supervision. Incidental supervision is a recently introduced learning strategy (Roth, 2017), which seeks to retrieve supervision signals from data that are not labeled for the target task. This strategy has been applied to tasks including SRL (He et al., 2020), controversy prediction (Rethmeier et al., 2018) and dataless classification (Song and Roth, 2015). To the best of our knowledge, the proposed method here is the first of its kind that incorporates incidental supervision in embedding learning or alignment.

3 Method

We hereby begin introducing our method with the formalization of learning resources.

In a KB, \mathcal{L} denotes the set of languages, and \mathcal{L}^2 unordered language pairs. G_L is the language-specific KG of language $L \in \mathcal{L}$. E_L and R_L respectively denote the corresponding vocabularies of entities and relations. $T = (h, r, t)$ denotes a triple in G_L such that $h, t \in E_L$ and $r \in R_L$. Boldfaced \mathbf{h} , \mathbf{r} , \mathbf{t} represent the embedding vectors

of head h , relation r , and tail t respectively. For a language pair $(L_1, L_2) \in \mathcal{L}^2$, $I_E(L_1, L_2)$ denotes a set of entity alignments between L_1 and L_2 , such that $e_1 \in E_{L_1}$ and $e_2 \in E_{L_2}$ for each entity pair $(e_1, e_2) \in I_E(L_1, L_2)$. Following the convention of previous work (Chen et al., 2018; Sun et al., 2018; Yang et al., 2019), we assume the entity pairs to have a 1-to-1 mapping and it is specified in $I_E(L_1, L_2)$. This assumption is congruent to the design of mainstream KBs (Lehmann et al., 2015; Mahdisoltani et al., 2015) where disambiguation of entities is granted. Besides the definition of multilingual KGs, we use D_L to denote the text corpus of language L . D_L is a set of documents $\{d_L\}$, where each document $d_L = [w_1, w_2, \dots, w_l]$ is a sequence of tokens from the monolingual lexicon W_L . Each token w_i thereof is originally a lexeme, but may also be an entity surface form after the ground process, and we also use bold-faced \mathbf{w}_i to denote its vector. $I_W(L_1, L_2)$ denotes the seed lexicon between (L_1, L_2) , such that $w_1 \in W_{L_1}$ and $w_2 \in W_{L_2}$ for each lexeme pair $(w_1, w_2) \in I_W(L_1, L_2)$. Note that I_W only include the alignment between lexemes, and may optionally serve as external supervision data. To be consistent with previous problem settings of entity alignment (Chen et al., 2017a; Sun et al., 2018; Yang et al., 2019), I_W is not necessarily provided to training, but is defined to be compatible with the scenarios where it is available.

JEANS addresses entity alignment in three consecutive processes. (i) A *grounding process* first link entities of each KG to possible mentions of them in the corresponding monolingual corpus, therefore connecting entities and text tokens of the same language into a shared vocabulary. (ii) An *embedding learning* process characterizes the KG and text of each language into a separate embedding space. In this process, we couple both the translational technique (Bordes et al., 2013; Chen et al., 2017a, 2018) and the neighborhood aggregation technique (Wang et al., 2018; Yang et al., 2019), which are two representative techniques to characterize a KG. Simultaneously, the monolingual text tokens are encoded with a skip-gram language model (Mikolov et al., 2013). (iii) On top of the embeddings, starting from a small amount of seed entity alignment and *optional* seed lexicon, the *alignment learning* process iteratively infers more alignment both on KGs and text using self-learning and Procrustes solution (Schönemann,

1966). The processes of JEANS’s learning is consistent to Figure 1. The rest of this section introduces the technical details of each process.

3.1 (Noisy) Entity Grounding

The goal of the grounding process is to combine vocabularies of the KG and the text corpus in each language. This serves as the premise for the embedding learning process to produce a shared representation scheme for entities, relations and lexemes, therefore allowing the alignment learning process to leverage supervision signals for both entities and lexemes. It is noteworthy that, the purpose of entity grounding here is to combine the two data modalities. Hence, we only expect this process to discover enough entity contexts and offer a higher coverage on entity vocabularies, while being tolerant to possible noise in entity recognition and linking. Particularly, we consider two grounding techniques, one using a pre-trained entity discovery and linking (EDL) model, the other based on simple surface form matching (SFM).

Pre-trained EDL model. One technique is to use off-the-shelf EDL models (Khashabi et al., 2018; Manning et al., 2014). A typical model of such sequentially handles the steps of NER to detect entity mentions, and link each mention to candidate entities from the KG based on symbolic and contextual similarity. Many EDL models are easily trainable on large text corpora with anchors, and offer promising performance of grounding and disambiguation on multiple languages (Sil et al., 2018). In this paper, we do not go into details to the design of EDL models. Interested readers are referred to the aforementioned literature.

Surface form matching. Suppose a pre-trained EDL model is not available, then a simpler way of combining data is to match KG surface forms with text. This can be efficiently done by building a Completion Trie (Hsu and Ottaviano, 2013) for multi-token surface forms, and conducting a longest prefix matching (Dharmapurikar et al., 2006) between surface forms and sub-sequences of text tokens. While this simple technique does not necessarily disambiguate entity mentions, experiments find it sufficient to combine the two modalities and allow supervision signals from induced lexical alignment to propagate to entities.

Once the entity vocabulary E_L and the lexicon W_L of a language are combined, we assume that

entity mentions in D_L are properly tokenized as grounded surface forms in $E_L \cap W_L$. Specifically, we now use x to denote a token in the grounded corpus D_L that can either be an entity e or a lexeme w . Given the combined learning resources for each language, we next describe the processes of embedding learning and alignment learning.

3.2 Embedding Learning

The embedding learning process is responsible for capturing the combined KG and text corpus of each language in a shared embedding space \mathbb{R}^k . In this process, JEANS jointly learns two model components to respectively encode units of the KG and the text, among which the overlaps $E_L \cap W_L$ use shared representations. We hereby describe these two model components in detail.

3.2.1 KG Embedding

As discussed in §2, previous approaches respectively leverage two forms of embedding learning techniques: (i) *relation modeling* (Chen et al., 2017a; Sun et al., 2018) such as vector translations, circular correlation and Hadamard product seeks to capture relations as an arithmetic operation in the vector space; (ii) *neighborhood aggregation* (Wang et al., 2018; Yang et al., 2019; Cao et al., 2019) employs graph neural networks (GNN) to encode neighborhood contexts for better seizing the proximity of entities.

The KG embedding model proposed in this work couples both forms of techniques. This aims at seizing both relations and entity proximity, two factors that are both beneficial to produce transferable entity embeddings. To achieve this goal, the encoder first stacks n layers of GCN (Kipf and Welling, 2016) on the KG. Formally, the l -th layer representation $\mathbf{E}^{(l)}$ is computed as

$$\mathbf{E}^{(l)} = \phi \left(D^{-\frac{1}{2}} \tilde{A} D^{-\frac{1}{2}} \mathbf{E}^{(l-1)} \mathbf{M}^{(l-1)} \right),$$

where D is the diagonal degree matrix D of the KG, $\tilde{A} = A + I$ is the sum of the adjacency matrix A and an identity I , and $\mathbf{M}^{(l-1)}$ is a trainable weight matrix. The raw features of entities $\mathbf{E}^{(0)}$ can be either entity attributes or randomly initialized. The last layer outputs are regarded as entity embedding representations, i.e. $\mathbf{E} = \mathbf{E}^{(n)}$.

We use \mathbf{E}_L to denote the entity representations of language L , then the following log-softmax loss is optimized to perform relational modeling with translation vectors in the embedding space of L :

$$S_L^K = - \sum_{T \in G_L} \log \frac{\exp(b - f_r(h, t))}{\sum_{\hat{T} \notin G_L} \exp(b - f_r(\hat{h}, \hat{t}))},$$

where $f_r(h, t) = \|\mathbf{h} + \mathbf{r} - \mathbf{t}\|$ is the plausibility measure of a triple (Bordes et al., 2013), $\hat{T} = (\hat{h}, r, \hat{t})$ is a Bernoulli negative-sampled triple (Wang et al., 2014a) created by substituting either head or tail entities h or t in $T = (h, r, t)$. b is a positive bias to adjust the scale of the plausibility measure. All the entity representations optimized in S_L^K are from \mathbf{E}_L . Note that the reason for us to choose the translational technique over other relation modeling techniques is due to this technique being more robust in cases where KG structures are sparser (Pujara et al., 2017).

3.2.2 Text Embedding

In addition to the KG embedding, the text embedding seeks to leverage the contextual information of free text to help the embedding better capture the proximity of entities. This model employs the continuous skip-gram language model, which is inline with a number of word embedding methods (Mikolov et al., 2013; Bojanowski et al., 2017; Conneau et al., 2018), and is realized by optimizing the following log-softmax loss:

$$S_L^T = - \sum_{x \in E_L \cup W_L} \sum_{x_c \in C_{x, D_L}} \log \frac{\exp(d(x, x_c))}{\sum_{x_n} \exp(d(x, x_n))}.$$

The text context C_{x, D_L} thereof is the set of tokens that surround a token x in the entity-grounded corpus D_L , d denotes the l_2 distance, and x_n denotes a randomly sampled token in $E_L \cup W_L$.

3.2.3 Embedding Learning Objective

For each language $L \in \mathcal{L}$, the goal of embedding learning is to optimize the joint loss

$$S_L^E = S_L^K + S_L^T.$$

As mentioned, the grounded entity surface forms in $E_L \cap W_L$ use shared representations in both model components, hence are optimized with both S_L^K and S_L^T . The rest lexeme, relation and entity representations are optimized alternately by either component. In both model components, the number of negative samples of triples and tokens are both adjustable hyperparameters.

It is noteworthy that, both model components may choose alternative techniques, including other

KG encoders such as GAT (Veličković et al., 2018), multi-channel GCN (Cao et al., 2019) and gated GNN (Sun et al., 2020a), and text embeddings such as GloVe (Pennington et al., 2014). As experimenting with different embedding techniques is not a main contribution of this work, we leave them as future work. Specifically, contextualized text representations (Peters et al., 2018; Devlin et al., 2019) cannot directly apply, as contextualization will cause ambiguity to token representations that hinder the match of embeddings.

3.3 Alignment Learning

Once the KG and text units of each language are captured in a shared embedding, the alignment learning process therefore bridges the alignment between each pair of embeddings. This process seeks to exploit additional alignment labels from text embeddings, and use those to help the alignment of entities. Different from the majority of methods in §2 that jointly learn embeddings and alignment, the alignment learning process in JEANS is a retrofitting process (Shi et al., 2019; Faruqui et al., 2015). Hence, the embedding of each language is fixed and does not require duplicate training for different language pairs (Chen et al., 2017a; Sun et al., 2017).

Given a pair of languages $(L_i, L_j) \in \mathcal{L}^2$, the objective of alignment learning is to induce a transformation $\mathbf{M}_{ij} \in \mathbb{R}^{k \times k}$ between the two embedding spaces. The following loss is minimized

$$S_{L_i, L_j}^A = \sum_{(x_i, x_j) \in I(L_i, L_j)} \|\mathbf{M}_{ij}\mathbf{x}_i - \mathbf{x}_j\|_2,$$

in which $I(L_i, L_j) = I_E(L_i, L_j) \cup I_W(L_i, L_j)$, and the word seed lexicon I_W is considered additional supervision data that are optionally provided. Each \mathbf{x}_i (\mathbf{x}_j) denotes a fixed representation of either an entity or a lexeme of L_i (L_j).

Starting from a small amount of seed alignment in $I(L_i, L_j)$, JEANS conducts an iterative self-learning process to exploit more alignment labels for both entities and lexemes to improve the learning of \mathbf{M}_{ij} . In each iteration, we follow Conneau et al. (2018) to induce a Procrustes solution for \mathbf{M}_{ij} . To propose new alignment labels, the self-learning technique in JEANS deploys a mutual nearest neighbor (NN) constraint, which requires a suggested pair of matched items to appear in the NN of each other. More specifically, define

$\mathcal{N}_{L_i}^K(\mathbf{x})$ as the K -NN of vector \mathbf{x} in the embedding space of L_i , this constraint requires a proposed match (x_i, x_j) to be inserted into I only if $\mathbf{M}_{ij}\mathbf{x}_i$ is in $\mathcal{N}_{L_j}^1(\mathbf{x}_j)$, and \mathbf{x}_j mutually appears in $\mathcal{N}_{L_i}^1(\mathbf{M}_{ij}\mathbf{x}_i)$. Besides, we also require (x_i, x_j) to be of the same type, i.e. both being entities or being lexemes. Particularly, we only select entities that have not been aligned in I to form the newly-proposed (x_i, x_j) . This respects the 1-to-1 matching constraint of entities being defined at the beginning of this section, and effectively reduces the candidate space after each iteration of self-learning. Meanwhile, 1-to-1 matching is not required for lexemes. To mitigate hubness, we also follow Conneau et al. (2018) to employ the Cross-domain Similarity Local Scaling (CSLS) measure.

After the iteration, the newly proposed alignment labels are inserted to I to enhance the learning at the next iteration. The iterative self-learning is stopped once the number of proposed entity alignment in an iteration is below certain quantity (e.g. 1% of $|E_{L_i}|$). With more and more matched entities and lexemes being exploited within each iteration, a better \mathbf{M}_{ij} is induced, whereas the lexical alignment naturally serve as incidental supervision signals for entity alignment.

After the alignment learning process, given a query $(e_i, ?e_j)$ to find the counterpart entity of $e_i \in E_{L_i}$ from E_{L_j} , the answer e_j is predicted as the 1-NN entity after applying \mathbf{M}_{ij} to transform e_i , denoted $\{e_j\} = \mathcal{N}_{E_{L_j}}^1(\mathbf{M}_{ij}e_i)$. The inference phase by default also adopts CSLS as the distance measure, which is consistent with the default setting of recent works (Sun et al., 2019b, 2020a).

4 Experiment

In this section, we evaluate JEANS on two benchmark datasets for cross-lingual entity alignment, and compare against a wide selection of recent baseline methods. We also provide detailed ablation study on model components of JEANS.

4.1 Experimental Settings

Datasets. Experiments are conducted on DBP15k (Sun et al., 2017) and WK3160k (Chen et al., 2018) that are widely used benchmarks on the studied task. DBP15k contains four language-specific KGs that are respectively extracted from English (En), Chinese (Zh), French (Fr) and Japanese (Ja) DBpedia (Lehmann et al.,

Settings	DBP15k _{En-Fr}			DBP15k _{En-Zh}			DBP15k _{En-Ja}			WK3160k _{En-Fr}			WK3160k _{En-De}		
	<i>H@1</i>	<i>H@10</i>	<i>MRR</i>	<i>H@1</i>	<i>H@10</i>	<i>MRR</i>	<i>H@1</i>	<i>H@10</i>	<i>MRR</i>	<i>H@1</i>	<i>H@5</i>	<i>MRR</i>	<i>H@1</i>	<i>H@5</i>	<i>MRR</i>
MTransE (Chen et al., 2017a) ^{†‡}	0.224	0.556	0.335	0.308	0.614	0.364	0.279	0.575	0.349	0.140	0.203	0.177	0.034	0.101	0.072
GCN-Align (Wang et al., 2018) [‡]	0.373	0.745	0.532	0.413	0.744	0.549	0.399	0.745	0.546	0.215	0.378	0.293	0.138	0.246	0.190
AlignE (Sun et al., 2018) [†]	0.481	0.824	0.599	0.472	0.792	0.581	0.448	0.789	0.563	---	---	---	---	---	---
GCN-JE (Wu et al., 2019b)	0.483	0.778	---	0.459	0.729	---	0.466	0.746	---	---	---	---	---	---	---
RotatE (Sun et al., 2019c) [†]	0.345	0.738	0.476	0.485	0.788	0.589	0.442	0.761	0.550	---	---	---	---	---	---
KECG (Li et al., 2019a)	0.486	0.851	0.610	0.478	0.835	0.598	0.490	0.844	0.610	---	---	---	---	---	---
MuGCN (Cao et al., 2019) [†]	0.495	0.870	0.621	0.494	0.844	0.611	0.501	0.857	0.621	---	---	---	---	---	---
RSN (Guo et al., 2019) [†]	0.516	0.768	0.605	0.508	0.745	0.591	0.507	0.737	0.590	---	---	---	---	---	---
GMN (Xu et al., 2019)	0.596	0.876	0.679	0.433	0.681	0.479	0.465	0.728	0.580	---	---	---	---	---	---
AliNet (Sun et al., 2020a) [†]	0.552	0.852	0.657	0.539	0.826	0.628	0.549	0.831	0.645	---	---	---	---	---	---
JAPE (Sun et al., 2017) ^{†‡}	0.324	0.667	0.430	0.412	0.745	0.490	0.363	0.685	0.476	0.169	0.354	0.271	0.147	0.239	0.192
SEA (Pei et al., 2019a) [†]	0.400	0.797	0.533	0.424	0.796	0.548	0.385	0.783	0.518	---	---	---	---	---	---
HMAN (Yang et al., 2019)	0.543	0.867	---	0.537	0.834	---	0.565	0.866	---	---	---	---	---	---	---
BootEA (Sun et al., 2018) ^{†‡}	0.653	0.874	0.731	0.629	0.847	0.703	0.622	0.854	0.701	0.333	0.511	0.425	0.233	0.393	0.316
KDCoE (Chen et al., 2018)	---	---	---	---	---	---	---	---	---	0.483	0.569	0.496	0.335	0.380	0.339
MMR (Shi and Xiao, 2019)	0.635	0.878	---	0.647	0.858	---	0.623	0.847	---	---	---	---	---	---	---
NAEA (Zhu et al., 2019)	0.673	0.894	0.752	0.650	0.867	0.720	0.641	0.873	0.718	---	---	---	---	---	---
OTEA (Pei et al., 2019b) [‡]	---	---	---	---	---	---	---	---	---	0.361	0.541	0.447	0.270	0.440	0.352
JEANS-SFM	0.766	0.939	0.814	0.713	0.885	0.773	0.723	0.913	0.793	0.463	0.558	0.538	0.337	0.450	0.412
JEANS-EDL	0.769	0.940	0.827	0.719	0.895	0.791	0.737	0.914	0.798	0.451	0.544	0.529	0.312	0.431	0.390

Table 1: Entity alignment results. Baselines are separated in accord with the three groups described in Section 4.1. [†] indicates results obtained from (Sun et al., 2020a), and [‡] indicates those from (Pei et al., 2019b). Results of KECG, GCN-JE, MMR, HMAN, KDCoE and NAEA are from original papers. Hyphens denote not available. MRR were not reported by GCN-JE, MMR and HMAN. Top results (incl. w/ and w/o seed lexicon) are boldfaced. Note that results by GCN-JE, GMN and HMAN are reported only for the versions where the extra cross-lingual alignment information (such as machine translation) is removed, so as to conduct fair comparison with all the rest models that are trained using only the alignment labels in the benchmark training sets.

2015), each of which contains around 65k-106k entities. Three sets of 15k alignment labels are constructed to align entities between each of the other three languages and English. WK3160k contains larger KGs with around 57k to 65k entities in En, Fr and German (De) KGs, and around 55k reference entity alignment for En-Fr and En-De settings. Dataset statistics are given in Appendix §A.2 (Chen et al., 2021).

We also use the text of Wikipedia dumps (dated Jan 01, 2019) in the five participating languages in training. For Chinese and Japanese corpora thereof, we obtain the segmented versions respectively from PKUSEG (Luo et al., 2019) and MeCab (Kudo, 2006).

Baseline methods. We compare with a wide selection of recent approaches for entity alignment on multilingual KGs. The baseline methods include (i) those employing different structure embedding techniques, namely MTransE (Chen et al., 2017a), GCN-Align (Wang et al., 2018), AlignE (Sun et al., 2018), GCN-JE (Wu et al., 2019b), KECG (Li et al., 2019a), MuGCN (Cao et al., 2019), RotatE (Sun et al., 2019c), RSN (Guo et al., 2019) and AliNet (Sun et al., 2020a); (ii) methods that incorporate auxiliary information of entities, namely JAPE (Sun et al., 2017), SEA (Pei et al., 2019a), GMN (Xu et al., 2019) and HMAN (Yang et al., 2019); (iii) semi-

supervised alignment learning methods, including BootEA (Sun et al., 2018), KDCoE (Chen et al., 2018), MMR (Shi and Xiao, 2019), NAEA (Zhu et al., 2019) and OTEA (Pei et al., 2019b). Descriptions of these methods are given in Appendix §A.1 (Chen et al., 2021).

Note that some works have allowed to incorporate extra cross-lingual signals such as machine translation in training, or using pre-aligned word embeddings to delimit candidate spaces (Wu et al., 2019a,b; Xu et al., 2019). For example, Wu et al. (2019a,b) used Google Translate to translate surface forms of entities in all other languages to English, and initialize the entity embeddings in their model with pre-trained word embedding of translated entity names. Results for these models are reported for the versions where the extra cross-lingual alignment information is removed so as to conduct fair comparison with all the rest models that are trained from scratch and using the same alignment labels in the benchmark datasets. This also necessarily prevents potential leakage of testing data to training (Liu et al., 2020), considering that training a comprehensive NMT system may have subsumed many of the testing data in the entity alignment benchmarks.

Evaluation protocols. The use of the datasets are consistent with previous studies of the baseline methods. On each language pair in DBP15k,

around 30% of seed alignment is used for training, the rest for testing. On WK3160k, 20% of seed alignment on En-Fr and En-De settings is respectively used for training. Following the convention, we calculate several ranking metrics on test cases, including the accuracy $H@1$, the proportion of cases that are ranked no larger than p $H@p$, and mean reciprocal rank MRR . Note that to align with the results in previous studies (Sun et al., 2020a; Pei et al., 2019b), p is set to 10 on DBP15k and 5 on WK3160k. All metrics are preferred higher to indicate better performance.

Model Configurations. We use AMSGrad (Reddi et al., 2018) to optimize the training losses of the embedding learning process, for which we set the learning rate α to 0.001, the exponential decay rates β_1 and β_2 to 0.9 and 0.999, and batch sizes to 512 for both S_L^K and S_L^T . Trainable parameters are initialized using Xavier initialization (Glorot and Bengio, 2010). The dimension k is set to 300, which is often used for bilingual word embedding models trained on Wikipedia corpora (Conneau et al., 2018; Gouws et al., 2015), considering that the vocabulary sizes and training data density here are relatively close to those models. The number of GCN layers is set to 2. We set negative sample sizes of triples and text contexts to 5, the text context width to be 10 and the bias b in S_L^K to be 2. More implementation details are in Appendix §A.3 (Chen et al., 2021). Specifically, we evaluate variants of JEANS by adjusting two technical details. First, for the grounding process, aside from the simple surface form matching (marked with SFM), we also explore with the off-the-shelf Wikification-based EDL model (Upadhyay et al., 2018, marked with EDL). A grounding performance estimation is given in §4.4. In addition, we consider both CSLS and l_2 in inference.

4.2 Results

We report the entity alignment results in Table 1.

Considering the baseline results on DBP15k, we can see that the simplest variant of JEANS using SFM-based grounding has consistently outperformed all baselines on three cross-lingual settings. Particularly, it leads to 17.0-17.4% of absolute improvement in $H@1$ over the best structure-based baseline, 14.0-22.3% over the best entity profile based one, and 6.30-9.30% over the best semi-supervised one. This shows that while JEANS preserves the key merit of a semi-

Setting	DBP15k _{En-Fr}			DBP15k _{En-Ja}		
Metrics	$H@1$	$H@10$	MRR	$H@1$	$H@10$	MRR
JEANS-SFM	0.766	0.939	0.814	0.723	0.913	0.793
—w/o Self-learning	0.628	0.845	0.720	0.622	0.835	0.728
—w/o GCN	0.742	0.913	0.809	0.709	0.905	0.789
—w/o Text	0.725	0.891	0.786	0.681	0.857	0.761
—w/o KG	0.699	0.872	0.771	0.635	0.819	0.706
—w/o CSLS	0.697	0.905	0.762	0.687	0.893	0.768
—w/ seed lexicons	0.788	0.947	0.848	0.738	0.931	0.803

Table 2: The ablation study results for components of JEANS based on DBP15k_{En-Fr} and DBP15k_{En-Ja}. Note that the additional seed lexicon is not used in the main experiment, and is not obligatory. The last row in this table is only to show the effectiveness of leveraging available supervision data on lexemes.

Language	SFM		EDL	
Estimation	Coverage	Avg match	Coverage	Avg match
En	0.982	1,268	0.933	1,367
Fr	0.987	295	0.926	929
Zh	0.855	141	0.774	348
Ja	0.982	159	0.797	881
De	0.981	297	0.951	1,092

Table 3: Estimated vocabulary coverage and average match per entity on each of the five language-specific Wikipedia corpora.

supervised entity alignment method, and effectively enhances the alignment of KGs by exploiting incidental supervision signals from unaligned text corpora. Considering different grounding techniques, we observe that SFM variants often perform closely to EDL ones. This indicates that simple SFM is enough to combine KG and text corpora for JEANS’s embedding learning without EDL-related resources. The results on Wk3160k generally exhibit similar observations. In comparison to KDCoE that leverages strong but expensive supervision data of entity descriptions in co-training, JEANS offers comparable performance based on very accessible resources.

In general, the experiments here show that JEANS promisingly improves SOTA performance for entity alignment, with only the need for unparalleled free text and no need for additional labels.

4.3 Ablation Study

In Table 2 we report an ablation study for JEANS-SFM based on DBP15k, so as to understand the importance of each incorporated technique.

From the results, we observe that self-learning is the most important factor. The removal of it can lead to a drop of 10.1-13.8% in $H@1$, as well as drastic drop of other metrics. This also

explains why semi-supervised baselines (group 3) typically perform better than others. However, even with self-learning, the removal of text can lead to $H@1$ drop of 2.4% on En-Fr and 4.2% on En-Ja. This shows that context information JEANS retrieves from free text effectively infers the match of entities. On the other hand, the structure encoding of KGs is more important than textual contexts, as it causes higher performance drops of 6.7-8.8% in $H@1$ by removing KGs. Note that the model without KG learns entity embeddings solely based on free text. Its results show that context information from text alone can provide a strong starting point from which incorporating KGs can further enhance its performance. Employing GCN leads to relatively slight performance gain, as joint learning the relation model and the language model can satisfyingly capture entity proximity. Changing the distance metric to l_2 also leads to 3.6-6.9% of decrease in $H@1$. This shows CSLS’s ability to handle hubness and isolation is also important for similarity inference in the dense embedding space for the metric words and entities. Hence, this metric is also recommended by recent work (Sun et al., 2020a, 2019b; Zhang et al., 2019). In addition, if we introduce additional 5k seed lexicon (with only word alignment information, not including any entity alignment) provided by Conneau et al. (2018) for each language pair, it leads to additional improvement of 1.5-2.2% in $H@1$. This shows that JEANS effectively leverages available supervision data on lexemes to further enhance entity alignment, although it is not obligatory.

4.4 Grounding Performance Estimation

Due to the lack of ground truths on unlabeled text, it is hard to estimate the precision of entity grounding by the two types of (noisy) grounding techniques. However, as the requirement of the grounding process is to simply connect two data modalities for training the embeddings, we may encourage a technique that handles enough entity mentions and offer a higher coverage on entity vocabularies. Accordingly, the estimations of these two factors for the two techniques are reported in Table 3. As we can observe that, without considering disambiguation, SFM can overall cover higher proportions of the entity vocabularies, while pre-trained EDL generally discovers more entity mentions for each entity. However, both techniques are

sufficient to support the noisy grounding process and combine two data modalities for embedding learning and alignment induction.

5 Conclusion

This paper introduces JEANS for entity alignment. Different from previous methods that leverage only internal information of KGs, JEANS extends the learning on any text corpora that may contain the KG entities. For each language, a noisy grounding process first connects both data modalities, followed by an embedding learning process coupling GCN with relational modeling, and a self-learning based alignment process. Without introducing additional labeled data, JEANS offers significantly improved performance over SOTA models on benchmarks. Hence, it shows the effectiveness and feasibility of exploiting incidental supervision from free text for entity alignment.

For future work, aside from experimenting with other embedding learning techniques for KGs and text, we plan to extend JEANS to learn associations on KGs with different specificity (Hao et al., 2019). We also seek to extend the representation scheme in hyperbolic spaces (Nickel and Kiela, 2017; Chen and Quirk, 2019) along with the incorporation of hyperbolic lexical embedding techniques (Tifrea et al., 2018), aiming at better capturing the associations for hierarchical ontologies.

Acknowledgement

We appreciate the anonymous reviewers for their insightful comments and suggestions. Also, we would like thank Jennifer Sheffield, Yi Zhang and other members of the UPenn Cognitive Computation Group for giving suggestions that improved the manuscript.

This research is supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA Contract No. 2019-19051600006 under the BETTER Program, and by Contract FA8750-19-2-1004 with the US Defense Advanced Research Projects Agency (DARPA). The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

- Satya Almasian, Andreas Spitz, and Michael Gertz. 2019. Word embeddings for entity-annotated texts. In *Proceedings of the European Conference on Information Retrieval (ECIR)*.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Jens Bleilholder and Felix Naumann. 2009. Data fusion. *ACM Computing Surveys*, 41(1):1–41.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics (TACL)*, 5.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of ACM SIGMOD International Conference on Management of Data (SIGMOD)*.
- Francis Bond and Ryan Foster. 2013. Linking and extending an open multilingual Wordnet. In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems (NIPS)*.
- Volha Bryl and Christian Bizer. 2014. Learning conflict resolution strategies for cross-language wikipedia data fusion. In *Proceedings of the World Wide Web Conference (WWW)*, pages 1129–1134.
- Yixin Cao, Lifu Huang, Heng Ji, Xu Chen, and Juanzi Li. 2017. Bridge text and knowledge by learning multi-prototype entity mention embedding. In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Yixin Cao, Zhiyuan Liu, Chengjiang Li, Juanzi Li, and Tat-Seng Chua. 2019. Multi-channel graph neural network for entity alignment. In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*, pages 1452–1461.
- Jiaao Chen, Jianshu Chen, and Zhou Yu. 2019. Incorporating structured commonsense knowledge in story completion. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.
- Muhao Chen and Chris Quirk. 2019. Embedding edge-attributed relational hierarchies. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*.
- Muhao Chen, Weijia Shi, Ben Zhou, and Dan Roth. 2021. Cross-lingual entity alignment with incidental supervision. *arXiv preprint arXiv:2005.00171v2 (with Appendices)*.
- 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 International Joint Conference on Artificial Intelligence (IJCAI)*.
- Muhao Chen, Yingtao Tian, Mohan Yang, and Carlo Zaniolo. 2017a. Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Muhao Chen, Tao Zhou, et al. 2017b. Multi-graph affinity embeddings for multilingual knowledge graphs. In *Automated Knowledge Base Construction (AKBC)*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *International Conference on Learning Representations (ICLR)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Sarang Dharmapurikar, Praveen Krishnamurthy, and David E Taylor. 2006. Longest prefix matching using bloom filters. *IEEE/ACM Transactions on Networking*, 14(2):397–409.
- Manaal Faruqui, Jesse Dodge, Sujay Kumar Jauhar, Chris Dyer, Eduard Hovy, and Noah A Smith. 2015. Retrofitting word vectors to semantic lexicons. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*.

- Stephan Gouws, Yoshua Bengio, et al. 2015. Bilbowa: Fast bilingual distributed representations without word alignments. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Lingbing Guo, Zequn Sun, and Wei Hu. 2019. Learning to exploit long-term relational dependencies in knowledge graphs. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Nitish Gupta, Sameer Singh, and Dan Roth. 2017. Entity linking via joint encoding of types, descriptions, and context. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Junheng Hao, Muhao Chen, Wenchao Yu, Yizhou Sun, and Wei Wang. 2019. Universal representation learning of knowledge bases by jointly embedding instances and ontological concepts. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*.
- Hangfeng He, Qiang Ning, and Dan Roth. 2020. Quase: Question-answer driven sentence encoding. In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Bo-June Paul Hsu and Giuseppe Ottaviano. 2013. Space-efficient data structures for top-k completion. In *Proceedings of the World Wide Web Conference (WWW)*.
- Ernesto Jiménez-Ruiz, Bernardo Cuenca Grau, Yujiao Zhou, and Ian Horrocks. 2012. Large-scale interactive ontology matching: Algorithms and implementation. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*.
- Daniel Khashabi, Mark Sammons, Ben Zhou, Tom Redman, Christos Christodoulopoulos, Vivek Srikumar, Nickolas Rizzolo, Lev Ratinov, Guanheng Luo, Quang Do, et al. 2018. Cogcompnlp: Your swiss army knife for nlp. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.
- Taku Kudo. 2006. Mecab: Yet another part-of-speech and morphological analyzer. <http://mecab.sourceforge.jp>.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, et al. 2015. Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6(2):167–195.
- Chengjiang Li, Yixin Cao, Lei Hou, Jiaxin Shi, Juanzi Li, and Tat-Seng Chua. 2019a. 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)*.
- Shiyang Li, Jianshu Chen, and Dian Yu. 2019b. Teaching pretrained models with commonsense reasoning: A preliminary kb-based approach. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. 2019. Kagnet: Knowledge-aware graph networks for commonsense reasoning. 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)*.
- Zhiyuan Liu, Yixin Cao, Liangming Pan, Juanzi Li, and Tat-Seng Chua. 2020. Exploring and evaluating attributes, values, and structure for entity alignment. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6355–6364.
- Ruixuan Luo, Jingjing Xu, Yi Zhang, Xuancheng Ren, and Xu Sun. 2019. Pkuseg: A toolkit for multi-domain chinese word segmentation. *CoRR*, abs/1906.11455.
- Farzaneh Mahdisoltani, Joanna Biega, et al. 2015. Yago3: A knowledge base from multilingual Wikipedias. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Tomas Mikolov, Kai Chen, et al. 2013. Efficient estimation of word representations in vector space. *International Conference on Learning Representations (ICLR)*.
- Tom Mitchell, William Cohen, Estevam Hruschka, Partha Talukdar, B Yang, J Betteridge, A Carlson, B Dalvi, M Gardner, B Kisiel, et al. 2018. Never-ending learning. *Communications of the ACM*.
- Diego Moussallem, Matthias Wauer, and Axel-Cyrille Ngonga Ngomo. 2018. Machine translation using semantic web technologies: A survey. *Journal of Web Semantics*, 51:1–19.
- Denis Newman-Griffis, Albert M Lai, and Eric Fosler-Lussier. 2018. Jointly embedding entities and text with distant supervision. In *Proceedings of the Workshop on Representation Learning for NLP*.
- Maximilian Nickel, Lorenzo Rosasco, Tomaso A Poggio, et al. 2016. Holographic embeddings of knowledge graphs. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.

- Maximillian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems (NIPS)*.
- Shichao Pei, Lu Yu, Robert Hoehndorf, et al. 2019a. semi-supervised entity alignment via knowledge graph embedding with awareness of degree difference. In *Proceedings of the Web Conference (WWW)*.
- Shichao Pei, Lu Yu, and Xiangliang Zhang. 2019b. Improving cross-lingual entity alignment via optimal transport. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Jeffrey Pennington, Richard Socher, et al. 2014. Glove: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual bert? In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Jay Pujara, Eriq Augustine, and Lise Getoor. 2017. Sparsity and noise: Where knowledge graph embeddings fall short. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. 2018. On the convergence of adam and beyond. In *International Conference on Learning Representations (ICLR)*.
- Nils Rethmeier, Marc Hübner, and Leonhard Hennig. 2018. Learning comment controversy prediction in web discussions using incidentally supervised multi-task cnns. In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*.
- Dan Roth. 2017. Incidental supervision: Moving beyond supervised learning. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2017. A survey of cross-lingual word embedding models. *Journal of Artificial Intelligence Research*.
- Peter H Schönemann. 1966. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10.
- Weijia Shi, Muhao Chen, Pei Zhou, and Kai-Wei Chang. 2019. Retrofitting contextualized word embeddings with paraphrases. 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)*.
- Xiaofei Shi and Yanghua Xiao. 2019. Modeling multi-mapping relations for precise cross-lingual 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)*.
- Pavel Shvaiko and Jérôme Euzenat. 2011. Ontology matching: state of the art and future challenges. *IEEE Transactions on Knowledge and Data Engineering*, 25(1):158–176.
- Avirup Sil, Heng Ji, Dan Roth, and Silviu Cucerzan. 2018. Multi-lingual entity discovery and linking. In *Proceedings of the Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Yangqiu Song and Dan Roth. 2015. Unsupervised sparse vector densification for short text similarity. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Robert Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.
- Fabian M Suchanek, Serge Abiteboul, et al. 2011. Paris: Probabilistic alignment of relations, instances, and schema. *Proceedings of the VLDB Endowment (PVLDB)*, 5(3).
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019a. Dream: A challenge data set and models for dialogue-based reading comprehension. *Transactions of the Association for Computational Linguistics (TACL)*, 7:217–231.
- Zequan Sun, Wei Hu, and Chengkai Li. 2017. Cross-lingual entity alignment via joint attribute-preserving embedding. In *Proceedings of the International Semantic Web Conference (ISWC)*.
- Zequan Sun, Wei Hu, Qingheng Zhang, and Yuzhong Qu. 2018. Bootstrapping entity alignment with knowledge graph embedding. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Zequan Sun, Chengming Wang, Wei Hu, Muhao Chen, Jian Dai, Wei Zhang, and Yuzhong Qu. 2020a. Knowledge graph alignment network with gated multi-hop neighborhood aggregation. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.

- Zequan Sun, Jiacheng Huang Wang, Wei Hu, Muhao Chen, and Yuzhong Qu. 2019b. Transedge: Translating relation-contextualized embeddings for knowledge graphs. In *Proceedings of the International Semantic Web Conference (ISWC)*.
- Zequan Sun, Qingheng Zhang, Wei Hu, Chengming Wang, Muhao Chen, Farahnaz Akrami, and Chengkai Li. 2020b. A benchmarking study of embedding-based entity alignment for knowledge graphs. *Proceedings of the VLDB Endowment (PVLDB)*, 13.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019c. Rotate: Knowledge graph embedding by relational rotation in complex space. In *International Conference on Learning Representations (ICLR)*.
- Alexandru Tifrea, Gary Becigneul, and Octavian-Eugen Ganea. 2018. Poincare glove: Hyperbolic word embeddings. In *International Conference on Learning Representations*.
- Kristina Toutanova, Danqi Chen, et al. 2015. Representing text for joint embedding of text and knowledge bases. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Bayu Distiawan Trsedya, Jianzhong Qi, and Rui Zhang. 2019. Entity alignment between knowledge graphs using attribute embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Shyam Upadhyay, Nitish Gupta, and Dan Roth. 2018. Joint multilingual supervision for cross-lingual entity linking. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations (ICLR)*.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014a. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhen Wang, Jianwen Zhang, et al. 2014b. Knowledge graph and text jointly embedding. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Zhichun Wang, Qingsong Lv, Xiaohan Lan, and Yu Zhang. 2018. Cross-lingual knowledge graph alignment via graph convolutional networks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Derry Wijaya, Partha Pratim Talukdar, et al. 2013. Pidgin: ontology alignment using web text as interlingua. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*.
- Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, Rui Yan, and Dongyan Zhao. 2019a. Relation-aware entity alignment for heterogeneous knowledge graphs. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, and Dongyan Zhao. 2019b. 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)*.
- 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 Annual Meeting of Associations for Computational Linguistics (ACL)*.
- Ikuya Yamada, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. 2017. Learning distributed representations of texts and entities from knowledge base. *Transactions of the Association for Computational Linguistics (TACL)*, 5:397–411.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding entities and relations for learning and inference in knowledge bases. *International Conference on Learning Representations (ICLR)*.
- Hsiu-Wei Yang, Yanyan Zou, Peng Shi, Wei Lu, Jimmy Lin, and Xu Sun. 2019. Aligning cross-lingual entities with multi-aspect information. 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)*.
- Kai Yang, Shaoqin Liu, Junfeng Zhao, Yasha Wang, and Bing Xie. 2020. Cotsae: Co-training of structure and attribute embeddings for entity alignment. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.
- Jinyoung Yeo, Geunyu Wang, Hyunsouk Cho, Seungtaek Choi, and Seung-won Hwang. 2018. Machine-translated knowledge transfer for commonsense causal reasoning. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.
- Qingheng Zhang, Zequan Sun, Wei Hu, Muhao Chen, Lingbing Guo, and Yuzhong Qu. 2019. Multi-view knowledge graph embedding for entity alignment. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Huaping Zhong, Jianwen Zhang, et al. 2015. Aligning knowledge and text embeddings by entity descriptions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Qiannan Zhu, Xiaofei Zhou, Jia Wu, Jianlong Tan, and Li Guo. 2019. Neighborhood-aware attentional representation for multilingual knowledge graphs. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Will Y Zou, Richard Socher, et al. 2013. Bilingual word embeddings for phrase-based machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.