

# From Alignment to Assignment: Frustratingly Simple Unsupervised Entity Alignment

Xin Mao<sup>1</sup>, Wenting Wang<sup>2</sup>, Yuanbin Wu<sup>1</sup>, Man Lan<sup>1\*</sup>

<sup>1</sup> School of Computer Science and Technology, East China Normal University

Lazada, Alibaba Group

xmao@stu.ecnu.edu.cn, wenting.wang@lazada.com

{ybwu,mlan}@cs.ecnu.edu.cn

## Abstract

Cross-lingual entity alignment (EA) aims to find the equivalent entities between cross-lingual KGs (Knowledge Graphs), which is a crucial step for integrating KGs. Recently, many GNN-based EA methods are proposed and show decent performance improvements on several public datasets. However, existing GNN-based EA methods inevitably inherit poor interpretability and low efficiency from neural networks. Motivated by the isomorphic assumption of GNN-based methods, we successfully transform the cross-lingual EA problem into an assignment problem. Based on this re-definition, we propose a frustratingly Simple but Effective Unsupervised entity alignment method (SEU) without neural networks. Extensive experiments have been conducted to show that our proposed unsupervised approach even beats advanced supervised methods across all public datasets while having high efficiency, interpretability, and stability.

## 1 Introduction

The knowledge graph (KG) represents a collection of interlinked descriptions of real-world objects and events, or abstract concepts (e.g., documents), which has facilitated many downstream applications, such as recommendation systems (Cao et al., 2019b; Wang et al., 2019) and question-answering (Zhao et al., 2020; Qiu et al., 2020). Over recent years, a large number of KGs are constructed from different domains and languages by different organizations. These cross-lingual KGs usually hold unique information individually but also share some overlappings. Integrating these cross-lingual KGs could provide a broader view for users, especially for the minority language users who usually suffer from lacking language resources. Therefore, how to fuse the knowledge from cross-lingual KGs has attracted increasing attentions.

As shown in Figure 1, cross-lingual entity alignment (EA) aims to find the equivalent entities




Figure 1: An example of cross-lingual knowledge graph entity alignment.

across multi-lingual KGs, which is a crucial step for integrating KGs. Conventional methods (Suchanek et al., 2011; Jiménez-Ruiz and Grau, 2011) usually solely rely on lexical matching and probability reasoning, which requires machine translation systems to solve cross-lingual tasks. However, existing machine translation systems are not able to achieve high accuracy with limited contextual information, especially for language pairs that are not alike, such as Chinese-English and Japanese-English.

Recently, Graph Convolutional Network (GCN) (Kipf and Welling, 2017) and subsequent Graph Neural Network (GNN) variants have achieved state-of-the-art results in various graph application. Intuitively, GNN is better in capturing structural information of KGs to compensate for the shortcoming of conventional methods. Specifically, several GNN-based EA methods (Xu et al., 2019; Wu et al., 2019a; Wang et al., 2020) indeed demonstrate decent performance improvements on public datasets. All these GNN-based EA methods are built upon a core premise, i.e., entities and their counterparts have similar neighborhood structures. However, better performance is not the only outcome of using GNN. Existing GNN-based methods inevitably inherit the following inborn defects from neural networks:

(1) **Poor Interpretability:** Recently, many researchers view GNN (Xu et al., 2019; Wu et al., 2019a) as a black box, focusing on improving performance metrics. The tight coupling between non-

linear operations and massive parameters makes GNN hard to be interpreted thoroughly. As a result, it is hard to judge whether the new designs are universal or just over-fitting on a specific dataset. A recent summary (Zhang et al., 2020) notes that several "advanced" EA methods are even beaten by the conventional methods on several public datasets.

(2) **Low Efficiency:** To further increase the performance, newly proposed EA methods try to stack novel techniques, e.g., Graph Attention Networks (Wu et al., 2019a), Graph Matching Networks (Xu et al., 2019), and Joint Learning (Cao et al., 2019a). Consequently, the overall architectures become more and more unnecessarily complex, resulting in their time-space complexities also dramatically increase. Zhao et al. (2020) present that the running time of complex methods (e.g., RDGCN (Wu et al., 2019a)) is  $10\times$  more than that of vanilla GCN (Wang et al., 2018).

In this paper, we notice that existing GNN-based EA methods inherit considerable complexity from their neural network lineage. Naturally, we consider eliminating the redundant designs from existing EA methods to enhance interpretability and efficiency without losing accuracy. Leveraging the core premise of GNN-based EA methods, we restate the assumption that both structures and textual features of source and target KGs are isomorphic. With this assumption, we are able to successfully transform the cross-lingual EA problem into an assignment problem, which is a fundamental and well-studied combinatorial optimization problem. Afterward, the assignment problem could be easily solved by the Hungarian algorithm (Kuhn, 1955) or Sinkhorn operation (Cuturi, 2013).

Based on the above findings, we propose a frustratingly Simple but Effective Unsupervised EA method (SEU) without neural networks. Compared to existing GNN-based EA methods, SEU only retains the basic graph convolution operation for feature propagation while abandoning the complex neural networks, significantly improving efficiency and interpretability. Experimental results on the public datasets show that SEU could be completed in several seconds with the GPU or tens of seconds with the CPU. More startlingly, our unsupervised method even outperforms the state-of-the-art supervised approaches across all public datasets. Furthermore, we discuss the possible reasons behind the unsatisfactory performance of existing complex EA methods and the necessity of neural networks




Figure 2: The architecture of existing EA methods.

in cross-lingual EA. The main contributions are summarized as follows:

- By assuming that both structures and textual features of source and target KGs are isomorphic, we successfully transform the cross-lingual EA problem into an assignment problem. Based on this finding, we propose a frustratingly Simple but Effective Unsupervised entity alignment method (SEU).
- Extensive experiments on public datasets indicate that our unsupervised method outperforms all advanced supervised competitors while preserving high efficiency, interpretability, and stability.

## 2 Task Definition

KG stores the real-world knowledge in the form of triples  $(h, r, t)$ . A KG could be defined as  $G = (E, R, T)$ , where  $E$ ,  $R$ , and  $T$  represent the entity set, relation set, and triple set, respectively. Given a source graph  $G_s = (E_s, R_s, T_s)$  and a target graph  $G_t = (E_t, R_t, T_t)$ , EA aims to find the entity correspondences  $P$  between KGs.

## 3 Related Work

### 3.1 Cross-lingual Entity Alignment

Existing cross-lingual EA methods are based on the premise that equivalent entities in different KGs have similar neighboring structures. Following this idea, most of them can be summarized into two steps (as shown in Figure 2): (1) Using KG embedding methods (e.g., TransE (Bordes et al., 2013) and GCN (Kipf and Welling, 2016)) to generate low-dimensional embeddings for entities and relations in each KGs. (2) Mapping these embeddings into a unified vector space through contrastive losses (Hadsell et al., 2006; Schroff et al., 2015) and pre-aligned entity pairs.




Figure 3: An example of isomorphic graph.

Based on the vanilla GCN, many EA methods design task-specific modules for improving the performance of EA. Cao et al. (2019a) propose a multi-channel GCN to learn multi-aspect information from KGs. Wu et al. (2019a) use a relation-aware dual-graph network to incorporate relation information with structural information. Moreover, due to the lack of labeled data, some methods (Sun et al., 2018; Mao et al., 2020) apply iterative strategies to generate semi-supervised data. In order to provide a multi-aspect view from both structure and semantic, some methods (Wu et al., 2019b; Yang et al., 2019) use word vectors of translated entity names as the input features of GNNs.

### 3.2 Assignment Problem

The assignment problem is a fundamental and well-studied combinatorial optimization problem. An intuitive instance is to assign  $N$  jobs for  $N$  workers. Assuming that each worker can do each job at a term, though with varying degrees of efficiency, let  $x_{ij}$  be the profit if the  $i$ -th worker is assigned to the  $j$ -th job. Then the problem is to find the best assignment plan (which job should be assigned to which person in one-to-one basis) so that the total profit of performing all jobs is maximum. Formally, it is equivalent to maximizing the following equation:

$$\arg \max_{\mathbf{P} \in \mathbb{P}_N} \langle \mathbf{P}, \mathbf{X} \rangle_F \quad (1)$$

$\mathbf{X} \in \mathbb{R}^{N \times N}$  is the profit matrix.  $\mathbf{P}$  is a permutation matrix denoting the assignment plan. There are exactly one entry of 1 in each row and each column in  $\mathbf{P}$  while 0s elsewhere.  $\mathbb{P}_N$  represents the set of all  $N$ -dimensional permutation matrices. Here,  $\langle \cdot \rangle_F$  represents the Frobenius inner product. In this paper, we adopt the Hungarian algorithm (Kuhn, 1955) and the Sinkhorn operation (Cuturi, 2013) to solve the assignment problem.

## 4 The Proposed Method

### 4.1 From Alignment to Assignment

The inputs of our proposed SEU are four matrices:  $\mathbf{A}_s \in \mathbb{R}^{|E_s| \times |E_s|}$  and  $\mathbf{A}_t \in \mathbb{R}^{|E_t| \times |E_t|}$  represent the adjacent matrices of the source graph  $G_s$  and the target graph  $G_t$ .  $\mathbf{H}_s \in \mathbb{R}^{|E_s| \times d}$  and  $\mathbf{H}_t \in \mathbb{R}^{|E_t| \times d}$  represent the textual features of entities that have been pre-mapped into a unified semantic space through machine translation systems or cross-lingual word embeddings.

Similar to the assignment plan, aligned entity pairs in EA also needs to satisfy the one-to-one constraint. Let a permutation matrix  $\mathbf{P} \in \mathbb{P}_{|E|}$  represent the entity correspondences between  $G_s$  and  $G_t$ .  $P_{ij} = 1$  indicates that  $e_i \in G_s$  and  $e_j \in G_t$  are an equivalent entity pair. The goal of SEU is to solve  $\mathbf{P}$  according to  $\{\mathbf{A}_s, \mathbf{A}_t, \mathbf{H}_s, \mathbf{H}_t\}$ . Consider the following ideal situation:

(1)  $\mathbf{A}_s$  and  $\mathbf{A}_t$  are isomorphic, i.e.,  $\mathbf{A}_s$  could be transformed into  $\mathbf{A}_t$  by reordering the entity node indices according to  $\mathbf{P}$  (as shown in Figure 3):

$$\mathbf{P} \mathbf{A}_s \mathbf{P}^{-1} = \mathbf{A}_t \quad (2)$$

(2) The textual features of equivalent entity pairs are mapped perfectly by the translation system. Therefore,  $\mathbf{H}_s$  and  $\mathbf{H}_t$  could also be aligned according to the entity correspondences  $\mathbf{P}$ :

$$\mathbf{P} \mathbf{H}_s = \mathbf{H}_t \quad (3)$$

By combining Equation (2) and (3), the connection between the 5-tuple  $\{\mathbf{A}_s, \mathbf{A}_t, \mathbf{H}_s, \mathbf{H}_t, \mathbf{P}\}$  could be described as follows:

$$\begin{aligned} (\mathbf{P} \mathbf{A}_s \mathbf{P}^{-1})^l \mathbf{P} \mathbf{H}_s &= \mathbf{A}_t^l \mathbf{H}_t \quad \forall l \in \mathbb{N} \\ \Rightarrow \mathbf{P} \mathbf{A}_s^l \mathbf{H}_s &= \mathbf{A}_t^l \mathbf{H}_t \quad \forall l \in \mathbb{N} \end{aligned} \quad (4)$$

Based on Equation (4),  $\mathbf{P}$  could be solved by minimizing the Frobenius norm  $\|\mathbf{P} \mathbf{A}_s^l \mathbf{H}_s - \mathbf{A}_t^l \mathbf{H}_t\|_F^2$  under the one-to-one constraint  $\mathbf{P} \in \mathbb{P}_{|E|}$ . Theoretically, for arbitrarily depth  $l \in \mathbb{N}$ , the solution of  $\mathbf{P}$  should be the same. However, the above inference is based on the ideal isomorphic situation. In practice,  $G_s$  and  $G_t$  are not strictly isomorphic and the translation system cannot perfectly map the textual features into a unified semantic space either. In order to reduce the impact of noise existing in practice,  $\mathbf{P}$  should be fit for various depths  $l$ . Therefore, we propose the following equation to solve the cross-lingual EA problem:

$$\arg \min_{\mathbf{P} \in \mathbb{P}_{|E|}} \sum_{l=0}^L \left\| \mathbf{P} \mathbf{A}_s^l \mathbf{H}_s - \mathbf{A}_t^l \mathbf{H}_t \right\|_F^2 \quad (5)$$

**Theorem 1** Equation (5) is equivalent to solving the following assignment problem:

$$\arg \max_{\mathbf{P} \in \mathbb{P}_{|E|}} \left\langle \mathbf{P}, \sum_{l=0}^L \mathbf{A}_t^l \mathbf{H}_t \left( \mathbf{A}_s^l \mathbf{H}_s \right)^T \right\rangle_F \quad (6)$$

*Proof:* According to the property of Frobenius norm  $\|\mathbf{A} - \mathbf{B}\|_F^2 = \|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2 - 2\langle \mathbf{A}, \mathbf{B} \rangle_F$ , Equation (5) could be derived into following:

$$\begin{aligned} & \arg \min_{\mathbf{P} \in \mathbb{P}_{|E|}} \sum_{l=0}^L \left\| \mathbf{P} \mathbf{A}_s^l \mathbf{H}_s - \mathbf{A}_t^l \mathbf{H}_t \right\|_F^2 \\ &= \arg \min_{\mathbf{P} \in \mathbb{P}_{|E|}} \sum_{l=0}^L \left\| \mathbf{P} \mathbf{A}_s^l \mathbf{H}_s \right\|_F^2 + \left\| \mathbf{A}_t^l \mathbf{H}_t \right\|_F^2 \\ & \quad - 2 \left\langle \mathbf{P} \mathbf{A}_s^l \mathbf{H}_s, \mathbf{A}_t^l \mathbf{H}_t \right\rangle_F \end{aligned} \quad (7)$$

Here, the permutation matrix  $\mathbf{P}$  must be orthogonal, so both  $\|\mathbf{P} \mathbf{A}_s^l \mathbf{H}_s\|_F^2$  and  $\|\mathbf{A}_t^l \mathbf{H}_t\|_F^2$  are constants. Then, Equation (7) is equivalent to maximizing as below:

$$\arg \max_{\mathbf{P} \in \mathbb{P}_{|E|}} \sum_{l=0}^L \left\langle \mathbf{P} \mathbf{A}_s^l \mathbf{H}_s, \mathbf{A}_t^l \mathbf{H}_t \right\rangle_F \quad (8)$$

For arbitrarily real matrices  $\mathbf{A}$  and  $\mathbf{B}$ , these two equations always hold:  $\langle \mathbf{A}, \mathbf{B} \rangle_F = \text{Tr}(\mathbf{A}\mathbf{B}^T)$  and  $\langle \mathbf{A}, \mathbf{B} + \mathbf{C} \rangle_F = \langle \mathbf{A}, \mathbf{B} \rangle_F + \langle \mathbf{A}, \mathbf{C} \rangle_F$ , where  $\text{Tr}(\mathbf{X})$  represents the trace of matrix  $\mathbf{X}$ . Therefore, Theorem 1 is proved:

$$\begin{aligned} & \arg \max_{\mathbf{P} \in \mathbb{P}_{|E|}} \sum_{l=0}^L \left\langle \mathbf{P} \mathbf{A}_s^l \mathbf{H}_s, \mathbf{A}_t^l \mathbf{H}_t \right\rangle_F \\ &= \arg \max_{\mathbf{P} \in \mathbb{P}_{|E|}} \sum_{l=0}^L \text{Tr} \left( \mathbf{P} \mathbf{A}_s^l \mathbf{H}_s (\mathbf{A}_t^l \mathbf{H}_t)^T \right) \\ &= \arg \max_{\mathbf{P} \in \mathbb{P}_{|E|}} \sum_{l=0}^L \left\langle \mathbf{P}, \mathbf{A}_t^l \mathbf{H}_t (\mathbf{A}_s^l \mathbf{H}_s)^T \right\rangle_F \\ &= \arg \max_{\mathbf{P} \in \mathbb{P}_{|E|}} \left\langle \mathbf{P}, \sum_{l=0}^L \mathbf{A}_t^l \mathbf{H}_t (\mathbf{A}_s^l \mathbf{H}_s)^T \right\rangle_F \end{aligned} \quad (9)$$

□

By Theorem 1, we successfully transform the EA problem into the assignment problem. Compared to GNN-based EA methods, our proposed method retains the basic graph convolution operation for feature propagation but replaces the complex neural networks with the well-studied assignment problem. Note that the entity scales  $|E_s|$  and  $|E_t|$  are usually inconsistent in practice, resulting in the profit matrix not being a square matrix. This kind of unbalanced assignment problem could be reduced to the balanced assignment problem easily. Assuming that  $|E_s| > |E_t|$ , a naive reduction is to

pad the profit matrix with zeros such that its shape becomes  $\mathbb{R}^{|E_s| \times |E_s|}$ . This naive reduction is suitable for the dataset with a small gap between  $|E_s|$  and  $|E_t|$ . For the dataset with a large entity scale gap, there is a more efficient reduction algorithm available (Ramshaw and Tarjan, 2012).

## 4.2 Two Algorithms for Solving the Assignment Problem

The first polynomial time-complexity algorithm for the assignment problem is the Hungarian algorithm (Kuhn, 1955), which is based on improving a matching along the augmenting paths. The time complexity of the original Hungarian algorithm is  $O(n^4)$ . Later, Jonker and Volgenant (1987) improve the algorithm to achieve  $O(n^3)$  running time, which is one of the most popular variants.

Besides the Hungarian algorithm, the assignment problem could also be regarded as a special case of the optimal transport problem. In the optimal transport problem, the assignment plan  $\mathbf{P}$  could be any doubly stochastic matrix instead of a permutation matrix. Based on the Sinkhorn operation (Sinkhorn, 1964; Adams and Zemel, 2011), Cuturi (2013) proposes a fast and completely parallelizable algorithm for the optimal transport problem:

$$\begin{aligned} S^0(\mathbf{X}) &= \exp(\mathbf{X}), \\ S^k(\mathbf{X}) &= \mathcal{N}_c(\mathcal{N}_r(S^{k-1}(\mathbf{X}))), \end{aligned} \quad (10)$$

$$\text{Sinkhorn}(\mathbf{X}) = \lim_{k \rightarrow \infty} S^k(\mathbf{X}).$$

where  $\mathcal{N}_r(\mathbf{X}) = \mathbf{X} \oslash (\mathbf{X} \mathbf{1}_N \mathbf{1}_N^T)$  and  $\mathcal{N}_c = \mathbf{X} \oslash (\mathbf{1}_N \mathbf{1}_N^T \mathbf{X})$  are the row and column-wise normalization operators of a matrix,  $\oslash$  represents the element-wise division, and  $\mathbf{1}_N$  is a column vector of ones. Then, Mena et al. (2018) further prove that the assignment problem could also be solved by the Sinkhorn operation as a special case of the optimal transport problem:

$$\begin{aligned} & \arg \max_{\mathbf{P} \in \mathbb{P}_N} \langle \mathbf{P}, \mathbf{X} \rangle_F \\ &= \lim_{\tau \rightarrow 0^+} \text{Sinkhorn}(\mathbf{X}/\tau) \end{aligned} \quad (11)$$

In general, the time complexity of the Sinkhorn operation is  $O(kn^2)$ . Because the number of iteration  $k$  is limited, the Sinkhorn operation can only obtain an approximate solution in practice. But according to our experimental results, a very small  $k$  is enough to achieve decent performance in entity alignment. Therefore, compared to the Hungarian algorithm, the Sinkhorn operation is  $n$  times more efficient, i.e.,  $O(n^2)$ .



### 4.3 Implementation Details

The above two sections introduce how to transform the cross-lingual EA problem into the assignment problem and how to solve the assignment problem. This section will clarify two important implementation details of our proposed method SEU.

#### 4.3.1 Textual Features $H$

The input features of SEU include two aspects:

**Word-Level.** In previous cross-lingual EA methods (Xu et al., 2019; Wu et al., 2019a), the most commonly used textual features are word-level entity name vectors. Specifically, these methods first use machine translation systems or cross-lingual word embeddings to map entity names into a unified semantic space and then average the pre-trained entity name vectors to construct the initial features. To make fair comparisons, we adopt the same entity name translations and word vectors provided by Xu et al. (2019).

**Char-Level.** Because of the contradiction between the extensive existence of proper nouns (e.g., person and city name) and the limited size of word vocabulary, the word-level EA methods suffer from a serious out of vocabulary (OOV) issue. Therefore, many EA methods explore the char-level features, using char-CNN (Wang et al., 2020) or name-BERT (Liu et al., 2020) to extract the char/sub-word features of entities. In order to keep the simplicity and consistency of our proposed method, we adopt the character bigrams of translated entity names as the char-level input textual features instead of complex neural networks.

In addition to these text-based methods, we notice that some structure-based EA methods (Wang et al., 2018; Guo et al., 2019) do not require any textual information at all, where the entity features are randomly initialized. Section 5.6 will discuss the connection between text-based and structure-based methods and challenge the necessity of neural networks in cross-lingual EA.

#### 4.3.2 Adjacent Matrix $A$

In Section 4.1, all deductions are built upon the assertion that the adjacency matrices  $A_s$  and  $A_t$  are isomorphic. Obviously, let  $D$  be the degree matrix of adjacency matrix  $A_{s/t}$ , the equal probability random walk matrix  $A_r = D^{-1}A_{s/t}$  and the symmetric normalized Laplacian matrix  $A_L = I - D^{-1/2}A_{s/t}D^{-1/2}$  of  $A_s$  and  $A_t$  are also isomorphic too. Therefore, if  $A_{s/t}$  is replaced by  $A_r$  or  $A_L$ , our method still holds.

Datasets		$ E $	$ R $	$ T $
DBP <sub>ZH-EN</sub>	Chinese	19,388	1,701	70,414
	English	19,572	1,323	95,142
DBP <sub>JA-EN</sub>	Japanese	19,814	1,299	77,214
	English	19,780	1,153	93,484
DBP <sub>FR-EN</sub>	French	19,661	903	105,998
	English	19,993	1,208	115,722
SRPRS <sub>FR-EN</sub>	French	15,000	177	33,532
	English	15,000	221	36,508
SRPRS <sub>DE-EN</sub>	German	15,000	120	37,377
	English	15,000	222	38,363

Table 1: Statistical data of DBP15K and SRPRS.

However, the above matrices ignore the relation types in the KGs and treat all types of relations equally important. We believe the relations with less frequency should have higher weight because they represent more unique information. Following this intuition, we apply a simple strategy to generate the relational adjacency matrix  $A_{rel}$ , for  $a_{ij} \in A_{rel}$ :

$$a_{ij} = \frac{\sum_{r_j \in R_{i,j}} \ln(|T|/|T_{r_j}|)}{\sum_{k \in \mathcal{N}_i} \sum_{r_k \in R_{i,k}} \ln(|T|/|T_{r_k}|)} \quad (12)$$

where  $\mathcal{N}_i$  represents the neighboring set of entity  $e_i$ ,  $R_{i,j}$  is the relation set between  $e_i$  and  $e_j$ ,  $|T|$  and  $|T_r|$  represent the total number of all triples and the triples containing relation  $r$ , respectively.

## 5 Experiments

Our experiments are conducted on a workstation with a GeForce GTX Titan X GPU and a Ryzen ThreadRipper 3970X CPU. The code and datasets are available in [github.com/MaoXinn/SEU](https://github.com/MaoXinn/SEU).

### 5.1 Datasets

To make fair comparisons with previous EA methods, we experiment with two widely used public datasets: (1) **DBP15K** (Sun et al., 2017): This dataset consists of three cross-lingual subsets from multi-lingual DBpedia: DBP<sub>FR-EN</sub>, DBP<sub>ZH-EN</sub>, DBP<sub>JA-EN</sub>. Each subset contains 15,000 entity pairs. (2) **SRPRS**: Guo et al. (2019) propose this sparse dataset, including two cross-lingual subsets: SRPRS<sub>FR-EN</sub> and SRPRS<sub>DE-EN</sub>. Each subset also contains 15,000 entity pairs but with much fewer triples compared to DBP15K.

The statistics of these datasets are summarized in Table 1. Most of the previous studies (Wang et al., 2018; Cao et al., 2019a) randomly split 30% of the entity pairs for training and development, while using the remaining 70% for testing. Because our proposed method is unsupervised, all of the entity pairs could be used for testing.

Method	DBP <sub>ZH-EN</sub>			DBP <sub>JA-EN</sub>			DBP <sub>FR-EN</sub>			SRPRS <sub>FR-EN</sub>			SRPRS <sub>DE-EN</sub>		
	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR
GCN-Align	0.434	0.762	0.550	0.427	0.762	0.540	0.411	0.772	0.530	0.243	0.522	0.340	0.385	0.600	0.460
MuGNN	0.494	0.844	0.611	0.501	0.857	0.621	0.495	0.870	0.621	0.131	0.342	0.208	0.245	0.431	0.310
BootEA	0.629	0.847	0.703	0.622	0.853	0.701	0.653	0.874	0.731	0.365	0.649	0.460	0.503	0.732	0.580
MRAEA	0.757	0.930	0.827	0.758	0.934	0.826	0.781	0.948	0.849	0.460	0.768	0.559	0.594	0.818	0.666
JEANS	0.719	0.895	0.791	0.737	0.914	0.798	0.769	0.940	0.827	-	-	-	-	-	-
GM-Align	0.679	0.785	-	0.739	0.872	-	0.894	0.952	-	0.574	0.646	0.602	0.681	0.748	0.710
RDGCN	0.697	0.842	0.750	0.763	0.897	0.810	0.873	0.950	0.901	0.672	0.767	0.710	0.779	0.886	0.820
HGCN	0.720	0.857	0.760	0.766	0.897	0.810	0.892	0.961	0.910	0.670	0.770	0.710	0.763	0.863	0.801
DAT	-	-	-	-	-	-	-	-	-	0.758	0.899	0.810	0.876	0.955	0.900
DGMC	0.801	0.875	-	0.848	0.897	-	0.933	0.960	-	-	-	-	-	-	-
AttrGNN	0.796	0.929	0.845	0.783	0.920	0.834	0.919	0.979	0.910	-	-	-	-	-	-
CEA	0.787	-	-	0.863	-	-	0.972	-	-	0.962	-	-	0.971	-	-
EPEA	0.885	0.953	0.911	0.924	0.969	0.942	0.955	0.986	0.967	-	-	-	-	-	-
SEU(word)	0.816	0.923	0.854	0.865	0.952	0.896	0.953	0.989	0.967	0.812	0.902	0.843	0.902	0.951	0.920
SEU(char)	0.870	0.947	0.897	0.947	0.984	0.961	0.986	0.998	0.990	0.979	0.994	0.985	0.980	0.994	0.985
SEU(w+c)	<b>0.900</b>	<b>0.965</b>	<b>0.924</b>	<b>0.956</b>	<b>0.991</b>	<b>0.969</b>	<b>0.988</b>	<b>0.999</b>	<b>0.992</b>	<b>0.982</b>	<b>0.995</b>	<b>0.986</b>	<b>0.983</b>	<b>0.996</b>	<b>0.987</b>

Table 2: Main experimental results on DBP15K and SRPRS. Baselines are separated in accord with the three groups described in Section 5.2. Most results are from the original papers. Some recent papers are failed to run on missing datasets or do not release the source code yet. We will fill in these blanks after contacting their authors.

## 5.2 Baselines

We compare our method against the following three groups of advanced EA methods: (1) **Structure**: These methods only use the structure information (i.e., triples): GCN-Align (Wang et al., 2018), MuGNN (Cao et al., 2019a), BootEA (Sun et al., 2018), MRAEA (Mao et al., 2020), JEANS (Chen et al., 2021). (2) **Word-level**: These methods average the pre-trained entity name vectors to construct the initial features: GM-Align (Xu et al., 2019), RDGCN (Wu et al., 2019a), HGCN (Wu et al., 2019b), DAT (Zeng et al., 2020b), DGMC (Fey et al., 2020). (3) **Char-level**: These EA methods further adopt the char-level textual features: AttrGNN (Liu et al., 2020), CEA (Zeng et al., 2020a), EPEA (Wang et al., 2020).

For our proposed method, SEU(word) and SEU(char) represent the model only using the word and char features as the inputs, respectively. SEU(w+c) represents concatenating the word and char features together as the inputs.

## 5.3 Settings

**Metrics.** Following convention, we use  $Hits@k$  and *Mean Reciprocal Rank* (MRR) as our evaluation metrics. The  $Hits@k$  score is calculated by measuring the proportion of correct pairs in the top- $k$ . In particular,  $Hits@1$  equals accuracy.

**Hyper-parameter.** In the main experiments, we use the Sinkhorn operation to solve the assignment problem. For all dataset, we use a same default setting: the depth  $L = 2$ ; the iterations  $k = 10$ ; the temperature  $\tau = 0.02$ .

## 5.4 Main Experiments

Table 2 shows the main experimental results of all EA methods. Numbers in **bold** denote the best results among all methods.

**SEU vs. Baselines.** According to the results, our method consistently achieves the best performance across all datasets. Compared with the previous SOTA methods, SEU (w+c) improves the performance on  $Hits@1$  and  $MRR$  by 1.5% and 1.3% at least. More importantly, SEU outperforms the supervised competitors as an unsupervised method, which is critical in practical applications.

In addition to the better performances, SEU also has better interpretability and stability: (1) When solving with the Hungarian algorithm, we can trace the reasons for each decision by the augmenting path, which brings better interpretability. (2) As we all know, neural networks optimized by SGD usually have some performance fluctuations. Since both the Hungarian algorithm and Sinkhorn operation are deterministic, multiple runs of these algorithms remain unchanged under the same hyper-parameters, which means better stability.




Figure 4: F1-score of SEU(w+c) and PARIS.

algorithm	DBP <sub>ZH-EN</sub>	DBP <sub>JA-EN</sub>	DBP <sub>FR-EN</sub>
Hungarian	0.907	0.963	0.993
Sinkhorn	0.900	0.956	0.988

Table 3: *Hits*@1 of Hungarian and Sinkhorn. <sup>1</sup>

algorithm	DBP <sub>ZH-EN</sub>	DBP <sub>JA-EN</sub>	DBP <sub>FR-EN</sub>
Hungarian	43.4s	19.8s	7.6s
Sinkhorn(CPU)	6.1s	6.1s	6.2s
Sinkhorn(GPU)	1.8s	1.7s	1.8s

Table 4: Time costs of Hungarian and Sinkhorn.

**Word vs. Char.** From Table 2, we observe that the char-level SEU greatly outperforms the word-level SEU. Especially in SRPRS<sub>FR-EN</sub>, the performance gap on *Hits*@1 is more than 16%. As mentioned in Section 4.3.1, the main reason is that these datasets contain extensive OOV proper nouns. For example, in DBP15K, 4-6% of the words are OOV; while in SRPRS<sub>DE-EN</sub> and SRPRS<sub>FR-EN</sub>, more than 12% and 16% of the entity names are OOV, respectively.

Note that the performance difference between SEU(word) and SEU(char) is vast, but these two features still complement to each other, so the combination of them still improves the performances (especially on DBP<sub>ZH-EN</sub> dataset). We believe the hidden reason is synonyms. For example, *soccer* and *football* refer to the same Chinese phrase, but there is almost no overlap in the char-level between these two English words. However, the word-level features could bridge such semantic gap via pre-trained cross lingual word vectors.

**SEU vs. PARIS.** As mentioned in Section 1, a recent summary (Zhang et al., 2020) notes that several "advanced" EA methods are even beaten by the conventional methods. To make this study more comprehensive, we also compare

<sup>1</sup>Since the Hungarian algorithm only outputs the assigned entity pairs, instead of a probability matrix  $P$ , we can only report the *Hits*@1 performance.

Method	DBP15K	SRPRS
GCN-Align (Wang et al., 2018)	103	87
MuGNN (Cao et al., 2019a)	3,156	2,215
BootEA (Sun et al., 2018)	4,661	2,659
MRAEA (Mao et al., 2020)	3,894	1,248
GM-Align (Xu et al., 2019)	26,328	13,032
RDGCN (Wu et al., 2019a)	6,711	886
HGCN (Wu et al., 2019b)	11,275	2,504
<b>SEU(CPU)</b>	<b>22.1</b>	<b>13.8</b>
<b>SEU(GPU)</b>	<b>16.2</b>	<b>9.6</b>

Table 5: Time costs of EA methods (seconds).<sup>2</sup>

SEU against a representative conventional method PARIS (Suchanek et al., 2011) in Figure 4, which is a holistic unsupervised solution to align KGs based on probability estimates. Since PARIS may not always output a target entity for every source entity, we use the F1-score as the evaluation metric to deal with entities that do not have a match. In our method, the F1-score is equivalent to *Hits*@1. Consistent with Zhang’s summary, PARIS is better than most GNN-based EA methods. On the other hand, SEU outperforms PARIS significantly on these public datasets except for DBP<sub>ZH-EN</sub>.

**Hungarian vs. Sinkhorn** Table 3 reports the performances of SEU(w+c) with the Hungarian algorithm and Sinkhorn operation, respectively. Theoretically, the Hungarian algorithm could generate the optimal solution precisely, while the Sinkhorn operation can only generate an approximate solution. Therefore, the Hungarian algorithm is always slightly better, but the performance gap is relatively small. Furthermore, we list the time costs of these two algorithms in Table 4. We observe that the time costs of the Hungarian algorithm are unstable, which depend on the dataset. Meanwhile, the time costs of the Sinkhorn operation are much more stable. Because the Sinkhorn operation is completely parallelizable, its time costs could be further reduced by the GPU. In general, the Sinkhorn operation is more suitable for large-scale EA because of its higher efficiency.

**Overall Time Efficiency** We specifically evaluate the overall time costs of some EA methods and report the results in Table 5. It is obvious that the efficiency of SEU far exceeds all advanced competitors. Typically, existing GNN-based methods require forward propagations on every batch, and the convergence of models usually requires hundreds of batches. Since SEU does not have any trainable parameters, it only requires forward propagation once, enabling SEU to achieve such acceleration.




Figure 5: *Hits@1* on  $DBP_{ZH-EN}$  with different  $\tau$ .




Figure 6: *Hits@1* with different depths  $L$ .

## 5.5 Auxiliary Experiments

To explore the behavior of SEU in different situations, we design the following experiments:

**Temperature  $\tau$ .** Similar to the temperature  $\tau$  in the softmax operation,  $\tau$  in the Sinkhorn operation is also used to make the distribution closer to one-hot. With the remaining config unchanged, we set  $\tau$  with different values and report the corresponding performances of SEU(w+c) on  $DBP_{ZH-EN}$  in Figure 5. If we choose an appropriate  $\tau$ , the Sinkhorn algorithm will converge quickly to the optimal solution. But if  $\tau$  is set too large, the algorithm will fail to converge.

**Depth  $L$ .** For depth  $L$ , we list the experimental results in Figure 6. In particular,  $L = 0$  is equivalent to aligning entities only according to their own features without the neighborhood information. SEU(w+c) with  $L = 2$  achieves the best performance on all subsets of DBP15K, which indicates the necessity of introducing neighborhood information. Similar to GNN-based EA methods, SEU is also affected by the over-smoothing problem. When stacking more layers, the performances begin to decrease slightly.

**Adjacency matrix  $A$ .** To distinguish different relation types in KGs, we adopt a simple strategy to generate the relational adjacency matrix  $A_{rel}$ . Table 6 reports the performances of SEU(w+c) with different types of adjacency matrices.  $A$  is the standard adjacency matrix,  $A_r = D^{-1}A$  is the

Method	$DBP_{ZH-EN}$		$DBP_{JA-EN}$		$DBP_{FR-EN}$	
	Hits@1	MRR	Hits@1	MRR	Hits@1	MRR
$A$	0.890	0.915	0.952	0.965	0.985	0.989
$A_r$	0.891	0.916	0.953	0.966	0.985	0.988
$A_L$	0.887	0.912	0.953	0.965	0.984	0.987
$A_{rel}$	<b>0.900</b>	<b>0.924</b>	<b>0.956</b>	<b>0.969</b>	<b>0.988</b>	<b>0.992</b>

Table 6: Performances with different types of adjacency matrices  $A$ .

equal probability random walk matrix and  $A_L = I - D^{-1/2}AD^{-1/2}$  is the symmetric normalized Laplacian matrix. The experimental results show that  $A_{rel}$  achieves the best performance across all these three subsets.

## 5.6 Discussion

From the experimental results, we observe that the supervised EA methods are even beaten by the unsupervised methods. In this section, we propose a hypothesis that the reason behind this counter-intuitive phenomenon is potential over-fitting.

As mentioned in Section 5.2, existing EA methods could be divided into structure-based and text-based according to the input features. The only difference between them is that the structure-based methods use randomly initialized vectors as the entity features, while the text-based methods use pre-mapped textual features as the inputs. Let us consider the vanilla GCN as a sample:

$$H^{l+1} = \sigma(A_L H^l W^l) \quad (13)$$

where  $\sigma$  represents the activation function. For the structure-based methods, since the input features  $H$  and the transformation matrix  $W$  are both randomly initialized, they could be simplified into one matrix, i.e.,  $H^{l+1} = \sigma(A_L H^l)$ . This idea has been proved by many structure-based EA methods (Cao et al., 2019a; Mao et al., 2020), which propose to diagonalize or remove the transformation matrix  $W$ . In this situation, GCN is reduced to a simple fully connected neural network with adjacency matrices as its input features. The essence of structure-based EA methods is to map the features of adjacency matrices into a unified vector space. Therefore, these structure-based EA methods require supervised data to learn the parameters.

As for the text-based EA methods, the textual features of entities have already been pre-mapped into a unified semantic space by machine translation or cross-lingual word vectors. Therefore, these



text-based EA methods are equivalent to further fitting these pre-mapped features on a few aligned entity pair seeds, which could cause potential overfitting. Considering that we could directly align entities as an assignment problem, it is unnecessary to further fit entity features via neural networks.

As a simple unsupervised method, our proposed SEU achieves excellent performances on several EA datasets, which confirms the above analysis from the empirical side. It is noted that this section only proposes a possible explanation, not rigorous proof. We will continue to explore in this direction.

## 6 Conclusion

In this paper, we successfully transform the cross-lingual EA problem into the assignment problem. Based on this finding, we propose a frustratingly Simple but Effective Unsupervised EA method (SEU) without neural networks. Experiments on widely used public datasets indicate that SEU outperforms all advanced competitors and has high efficiency, interpretability, and stability.

## References

- Ryan Prescott Adams and Richard S. Zemel. 2011. [Ranking via sinkhorn propagation](#). *CoRR*, abs/1106.1925.
- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 2787–2795.
- Yixin Cao, Zhiyuan Liu, Chengjiang Li, Zhiyuan Liu, Juanzi Li, and Tat-Seng Chua. 2019a. [Multi-channel graph neural network for entity alignment](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1452–1461.
- Yixin Cao, Xiang Wang, Xiangnan He, Zikun Hu, and Tat-Seng Chua. 2019b. [Unifying knowledge graph learning and recommendation: Towards a better understanding of user preferences](#). In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 151–161.
- Muhao Chen, Weijia Shi, Ben Zhou, and Dan Roth. 2021. [Cross-lingual entity alignment with incidental supervision](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 645–658.
- Marco Cuturi. 2013. [Sinkhorn distances: Lightspeed computation of optimal transport](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 2292–2300.
- Matthias Fey, Jan Eric Lenssen, Christopher Morris, Jonathan Masci, and Nils M. Kriege. 2020. [Deep graph matching consensus](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Lingbing Guo, Zequn Sun, and Wei Hu. 2019. [Learning to exploit long-term relational dependencies in knowledge graphs](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 2505–2514.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. [Dimensionality reduction by learning an invariant mapping](#). In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006), 17-22 June 2006, New York, NY, USA*, pages 1735–1742.
- Ernesto Jiménez-Ruiz and Bernardo Cuenca Grau. 2011. [Logmap: Logic-based and scalable ontology matching](#). In *The Semantic Web - ISWC 2011 - 10th International Semantic Web Conference, Bonn, Germany, October 23-27, 2011, Proceedings, Part I*, pages 273–288.
- Roy Jonker and A. Volgenant. 1987. [A shortest augmenting path algorithm for dense and sparse linear assignment problems](#). *Computing*, 38(4):325–340.
- Thomas N. Kipf and Max Welling. 2016. [Semi-supervised classification with graph convolutional networks](#). *CoRR*, abs/1609.02907.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Harold W Kuhn. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97.
- Zhiyuan Liu, Yixin Cao, Liangming Pan, Juanzi Li, and Tat-Seng Chua. 2020. [Exploring and evaluating attributes, values, and structures for entity alignment](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6355–6364. Association for Computational Linguistics.

- Xin Mao, Wenting Wang, Huimin Xu, Man Lan, and Yuanbin Wu. 2020. [MRAEA: an efficient and robust entity alignment approach for cross-lingual knowledge graph](#). In *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020*, pages 420–428.
- Gonzalo E. Mena, David Belanger, Scott W. Linderman, and Jasper Snoek. 2018. [Learning latent permutations with gumbel-sinkhorn networks](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Yunqi Qiu, Yuanzhuo Wang, Xiaolong Jin, and Kun Zhang. 2020. [Stepwise reasoning for multi-relation question answering over knowledge graph with weak supervision](#). In *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020*, pages 474–482.
- Lyle Ramshaw and Robert E Tarjan. 2012. On minimum-cost assignments in unbalanced bipartite graphs. *HP Labs, Palo Alto, CA, USA, Tech. Rep. HPL-2012-40R1*.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. [Facenet: A unified embedding for face recognition and clustering](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 815–823. IEEE Computer Society.
- Richard Sinkhorn. 1964. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The annals of mathematical statistics*, 35(2):876–879.
- Fabian M. Suchanek, Serge Abiteboul, and Pierre Senellart. 2011. [PARIS: probabilistic alignment of relations, instances, and schema](#). *Proc. VLDB Endow.*, 5(3):157–168.
- Zequan Sun, Wei Hu, and Chengkai Li. 2017. [Cross-lingual entity alignment via joint attribute-preserving embedding](#). In *The Semantic Web - ISWC 2017 - 16th International Semantic Web Conference, Vienna, Austria, October 21-25, 2017, Proceedings, Part I*, volume 10587 of *Lecture Notes in Computer Science*, pages 628–644. Springer.
- Zequan Sun, Wei Hu, Qingheng Zhang, and Yuzhong Qu. 2018. [Bootstrapping entity alignment with knowledge graph embedding](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4396–4402.
- Hongwei Wang, Fuzheng Zhang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. 2019. [Multi-task feature learning for knowledge graph enhanced recommendation](#). In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 2000–2010.
- Zhichun Wang, Qingsong Lv, Xiaohan Lan, and Yu Zhang. 2018. [Cross-lingual knowledge graph alignment via graph convolutional networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 349–357.
- Zhichun Wang, Jinjian Yang, and Xiaoju Ye. 2020. [Knowledge graph alignment with entity-pair embedding](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 1672–1680.
- 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 Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5278–5284.
- 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 2019, Hong Kong, China, November 3-7, 2019*, pages 240–249.
- Kun Xu, Liwei Wang, Mo Yu, Yansong Feng, Yan Song, Zhiguo Wang, and Dong Yu. 2019. [Cross-lingual knowledge graph alignment via graph matching neural network](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3156–3161.
- 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 2019, Hong Kong, China, November 3-7, 2019*, pages 4430–4440.
- Weixin Zeng, Xiang Zhao, Jiuyang Tang, and Xuemin Lin. 2020a. [Collective entity alignment via adaptive features](#). In *36th IEEE International Conference on Data Engineering, ICDE 2020, Dallas, TX, USA, April 20-24, 2020*, pages 1870–1873.
- Weixin Zeng, Xiang Zhao, Wei Wang, Jiuyang Tang, and Zhen Tan. 2020b. [Degree-aware alignment for entities in tail](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 811–820. ACM.

- Ziheng Zhang, Hualuo Liu, Jiaoyan Chen, Xi Chen, Bo Liu, Yuejia Xiang, and Yefeng Zheng. 2020. [An industry evaluation of embedding-based entity alignment](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020 - Industry Track, Online, December 12, 2020*, pages 179–189.
- Chen Zhao, Chenyan Xiong, Xin Qian, and Jordan L. Boyd-Graber. 2020. [Complex factoid question answering with a free-text knowledge graph](#). In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, pages 1205–1216.
- X. Zhao, W. Zeng, J. Tang, W. Wang, and F. Suchanek. 2020. An experimental study of state-of-the-art entity alignment approaches. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–1.