

# FLORA: Unsupervised Knowledge Graph Alignment by Fuzzy Logic

Yiwen Peng<sup>[0009-0007-7902-4097]</sup>, Thomas Bonald<sup>[0000-0003-0468-0384]</sup>, and Fabian M. Suchanek<sup>[0000-0001-7189-2796]</sup>

Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France  
{yiwen.peng, thomas.bonald, fabian.suchanek}@telecom-paris.fr

**Abstract.** Knowledge graph alignment is the task of matching equivalent entities (that is, instances and classes) and relations across two knowledge graphs. Most existing methods focus on pure entity-level alignment, computing the similarity of entities in some embedding space. They lack interpretable reasoning and need training data to work. In this paper, we propose FLORA, a simple yet effective method that (1) is unsupervised, i.e., does not require training data, (2) provides a holistic alignment for entities and relations iteratively, (3) is based on fuzzy logic and thus delivers interpretable results, (4) provably converges, (5) allows dangling entities, i.e., entities without a counterpart in the other KG, and (6) achieves state-of-the-art results on major benchmarks.

**Keywords:** Knowledge Graphs · Entity Alignment · Holistic Matching · Symbolic Reasoning · Fuzzy logic

## 1 Introduction

The task of Knowledge Graph (KG) alignment consists in matching both entities and relations in one KG to their equivalents in another. Figure 1 shows a toy example of alignment between DBpedia and Wikidata. KG alignment is useful for knowledge fusion [53,49] and thus for a wide range of downstream applications like question answering [12,2], common-sense reasoning [32], and recommender systems [5,23]. It is a challenging task for several reasons. First, KGs are generally heterogeneous and incomplete: One KG may contain information that the other one does not contain. Second, entity names can differ vastly across KGs (in Figure 1 for instance, entities in Wikidata are represented by non-readable IDs, unlike in DBpedia). Third, training data are most often not available in practice, due to the high cost of manual annotation. Finally, the alignment of entities and relations is interdependent.

Many approaches exist for KG alignment, but most focus on the alignment of entities only. Moreover, they generally require training data [59] or assume a one-to-one entity mapping from one KG to the other (i.e, ignore dangling entities) [44]. A notable exception is PARIS [39], an approach based on the iterative computation of probabilistic equations that has been shown to outperform other methods in recent studies [29,43]. However, this approach (1) lacks convergence

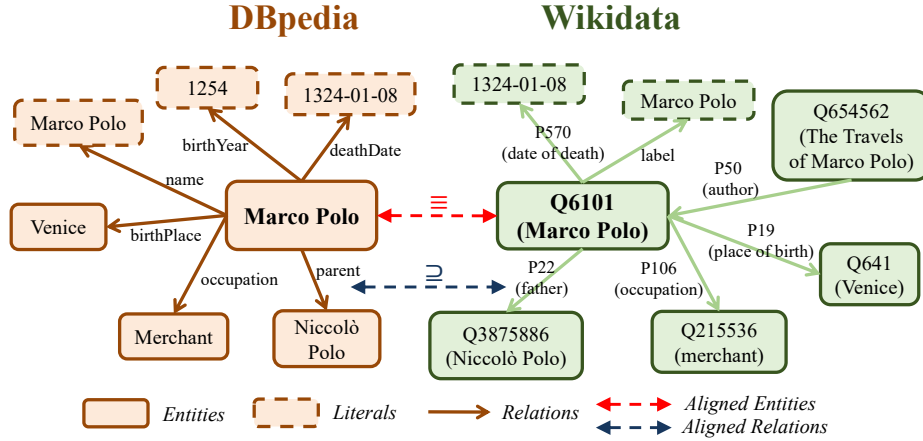


Fig. 1: An excerpt of DBpedia and Wikidata KGs. Note the asymmetric mapping of the relation *parent* to the sub-relation *father*.

guarantees, (2) does not achieve good performance when functional relations are absent, and (3) cannot see similarities between literals beyond a strict identity.

To address these shortcomings, we present FLORA<sup>1</sup>, a novel unsupervised approach for aligning both entities and relations. FLORA formalizes the alignment in Fuzzy Logic, in a principled framework that integrates various signals, both semantic and structural, and yields human-interpretable explanations for the results. FLORA works even in the presence of dangling entities and in the absence of functional relations. It relies on an iterative algorithm that provably converges. Our contributions are the following:

- FLORA, an approach based on Fuzzy Logic for aligning entities and relations that accounts for the asymmetry of relation alignments and the incompleteness of KGs. Although unsupervised by nature, FLORA can optionally exploit training data if available.
- An iterative algorithm that provably converges towards the optimal solution of the problem, as formalized in Fuzzy Logic.
- Extensive experiments on five entity alignment datasets across four languages, as well as two KG alignment datasets from the OAEI KG Track, showing that FLORA consistently outperforms both unsupervised and supervised baselines.

In what follows, Section 2 reviews related work, Section 3 provides preliminaries, Section 4 introduces our Fuzzy Logic formalism, Section 5 presents our method, Section 6 shows experimental results, and Section 7 concludes and outlines future work. All our code and data are publicly available at <https://github.com/dig-team/FLORA>.

<sup>1</sup> Fuzzy-Logic based Object and Relation Alignment

## 2 Related Work

**Entity Alignment.** Many existing approaches to entity alignment (EA) work in a supervised setting [59]. They use either translation-based embedding [40,41,48,24,42], GNN encoders [50,51,33,54,31] or fine-tuned language models [45,60,61,25]. Different from these approaches, FLORA does not require training data. In the unsupervised setting, non-neural methods such as LogMap [26], PARIS [39], PRASE [37], and NALA [55] exploit logical reasoning and lexical matching. Other methods transform the EA task into an optimal assignment problem [11,34,44], or leverage names, attributes and relations to improve the alignment of entities [58]. Neural methods, such as ICLEA [57], SelfKG [30], LLM4EA [8] and HLMEA [27], adopt a self-supervised setting or leverage large language models (LLMs) in a zero-shot manner. Unlike FLORA, these methods are *black-box* methods, i.e., they cannot explain why a given alignment was found. Besides, we show in our experiments that FLORA outperforms all of the above approaches on standard benchmark datasets. Moreover, FLORA aligns not just entities, but also relations at the same time.

**Knowledge Graph Alignment.** The task of KG alignment (also known as KG matching or holistic KG alignment) extends EA by aligning not only instances, but also classes and relations. Unlike Ontology Alignment [13], KG Alignment takes a simplistic view on classes, and typically does not consider asymmetric class subsumptions across KGs. The Ontology Alignment Evaluation Initiative (OAEI)<sup>2</sup> organizes a KG track tailored for KG alignment. Methods such as OLaLa [21] and SOBERT [18] align only classes. Our work, in contrast, aims at performing a holistic matching of instances, classes, and relations. Some methods like PARIS [39], LogMap [26], ATMatcher [20], and AgreementMakerLight [14] are based on lexical matching. Others, such as Wiktionary [36] or ALOD2Vec [35] incorporate external background knowledge. We can show that FLORA achieves better results than these methods without requiring additional external datasets.

**Fuzzy Logic.** Very few approaches use Fuzzy Logic for KG alignment. FuzzyEA [24] relies on a standard embedding approach for entity alignment and Fuzzy Logic is used only to filter candidate entity pairs. This is very different from our approach, where Fuzzy Logic is central in the alignment process, for both entities and relations. In other works [22,28,47,16], Fuzzy Logic is used for ontology alignment, but is limited to either class matching or instance matching.

## 3 Preliminaries

A knowledge graph  $\mathcal{G}$  is based on a set of entities  $\mathcal{E}$  and a set of relations  $\mathcal{R}$ . In this paper, in line with [9,49,6], we see a KG as a set of triples (or facts)  $\mathcal{T} \subset \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ . Each triple  $(h, r, t) \in \mathcal{T}$  represents a semantic relationship between the head entity  $h$  and the tail entity  $t$  with the relation  $r$ . We denote

<sup>2</sup> <https://oaei.ontologymatching.org/>

by  $r(h, t)$  the corresponding boolean proposition, which evaluates to *true* if and only if  $(h, r, t) \in \mathcal{T}$ . Without loss of generality, we assume that each relation  $r \in \mathcal{R}$  has an inverse relation  $r^{-1} \in \mathcal{R}$ , such that  $r(h, t)$  if and only if  $r^{-1}(t, h)$ . In this paper, the entity set  $\mathcal{E}$  includes literals, instances, and classes.

**Functionality.** Like PARIS [39], our approach exploits the notion of functionality of a relation, i.e., the degree to which a relation tends to map the head entity to a unique tail entity. Formally, the functionality of a relation  $r \in \mathcal{R}$  is:

$$fun(r) = \frac{|\{h \mid \exists t : r(h, t)\}|}{|\{(h, t) \mid r(h, t)\}|}$$

Note that  $fun(r) \in [0, 1]$  with  $fun(r) = 1$  if and only if the relation  $r$  is functional: It maps any head entity to a unique tail entity. For example, the relation *hasCapital* generally assigns exactly one capital city to each country, so that  $fun(hasCapital) \approx 1$ . If the entities referring to *France* in each KG are aligned, the corresponding entities referring to *Paris* can then be aligned as well.

Our approach also uses the notion of *local* functionality, which is specific to the head entity  $h$ :

$$fun(r, h) = \frac{1}{|\{t \mid r(h, t)\}|}$$

Again,  $fun(r, h) \in [0, 1]$  with  $fun(r, h) = 1$  if and only if there is a unique entity  $t$  such that  $r(h, t)$  is true. The local functionality is needed to deal with relations that are globally functional, but have local exceptions. For instance, *South Africa* has three capitals, *Pretoria*, *Bloemfontein*, and *Cape Town*, so that the relation *hasCapital* is not locally functional for the head entity *South Africa*.

**KG Alignment.** Given two knowledge graphs  $\mathcal{G} = \langle \mathcal{E}, \mathcal{R}, \mathcal{T} \rangle$ ,  $\mathcal{G}' = \langle \mathcal{E}', \mathcal{R}', \mathcal{T}' \rangle$ , entity alignment (EA) is the task of finding the set

$$\mathcal{M}_e = \{(e, e') \mid e \equiv e', e \in \mathcal{E}, e' \in \mathcal{E}'\}$$

where  $e \equiv e'$  means that entity  $e$  is semantically equivalent to  $e'$ . Relation alignment is the task of finding the set

$$\mathcal{M}_r = \{(r, O, r') \mid rOr', r \in \mathcal{R}, r' \in \mathcal{R}', O \in \{\subseteq, \supseteq, \equiv\}\},$$

where  $\subseteq, \supseteq, \equiv$  refer to subrelation, superrelation, or equivalence, respectively. KG alignment is the task of doing the alignment of both entities and relations at the same time.

## 4 Fuzzy Logic Framework

Our method, FLORA, is based on Fuzzy Logic. This allows us to integrate signals of different nature, semantic and structural, in a single principled framework. Our framework can use arbitrary aggregation functions, as long as they are continuous and non-decreasing. This distinguishes our framework from Probabilistic Soft Logic [4], which commonly requires linear aggregation functions (usually Łukasiewicz T-norms, or Gödel T-norms for negatively weighted rules [10]), and from standard propositional Gödel logic, which uses *min* aggregation [3].

**Fuzzy Inference System.** A *fuzzy set*  $A$  over a universe of discourse  $X$  is given by a membership function  $\mu_A$ , which maps every element of  $X$  to a value in  $[0, 1]$ . In an often-used example,  $X$  is the range of temperatures in degrees Celsius and  $A$  is the fuzzy set that describes “hot temperatures”. The membership function of  $A$  is then, for example,  $\mu_A(x) = 0$  for  $x \leq 10^\circ$ ,  $\mu_A(x) = (x - 10^\circ)/15^\circ$  for  $x$  between  $10^\circ$  and  $25^\circ$ , and  $\mu_A(x) = 1$  for  $x \geq 25^\circ$ . In its simplest form, a Mamdani-style Fuzzy Rule Based Reasoning System [38], or Fuzzy Inference System (FIS), consists of a set of rules of the form

$$p_1 \text{ is } P_1 \wedge \dots \wedge p_n \text{ is } P_n \Rightarrow c \text{ is } C$$

Here,  $p_1, \dots, p_n$  are input variables (we call them *premises*),  $P_1, \dots, P_n$  are fuzzy sets,  $c$  is the output variable (which we call the *conclusion*), and  $C$  is a fuzzy set. For example,  $p_1$  can be the temperature of the room in degrees Celcius,  $P_1$  is the fuzzy set that describes “hot temperature”,  $p_2$  is the number of people in the room,  $P_2$  is the fuzzy set of “a crowded room”,  $c$  is the desired speed of the fan (say, in rotations per minute), and  $C$  is the fuzzy set of “a fast rotation”. The expression “ $p_i \text{ is } P_i$ ” (for  $i = 1, \dots, n$ ) evaluates to  $\mu_{P_i}(p_i)$  (in a process called *fuzzification*). In our example, a room temperature of  $p_1 = 20^\circ$  belongs to the fuzzy set  $P_1$  of hot temperatures to a degree of  $\frac{2}{3}$ , and a number of people  $p_2 = 10$  could belong to “a crowded room”  $P_2$  to a degree of 0.5.

The *firing strength* of the rule is then computed by aggregating the individual fuzzified values of each premise using a custom aggregation function  $\phi$ . Usually,  $\phi$  is a T-norm [17], but we will not impose this restriction in this paper. A common choice for  $\phi$  is *min*, which yields 0.5 in our example. The output of the rule is a fuzzy set for  $c$ , which is given by the membership function of  $C$ , capped at the firing strength of the rule. In our example,  $\mu_C$  could assign a value of 1 to a speed of more than 60 rotations per minute, and would be capped to 0.5. If several rules have the same output variable ( $c$  in our example), their fuzzy sets ( $C$  in our example) are combined, usually using a point-wise max operation on the membership functions. This yields one final fuzzy set for the output variable, which is then *defuzzified* into a single scalar value  $c^*$ .

**Simple Positive FIS.** For our work, we need a simpler form of FIS, where all membership functions are identity functions, there is no negation, and defuzzification is done by the First of Maxima (FoM) method, i.e., the aggregated fuzzy set  $C$  is defuzzified into the scalar value  $c^* = \min\{c | \forall x, \mu_C(c) \geq \mu_C(x)\}$ . Our example rule can thus be written more simply as:

$$\text{temperature} \wedge \text{crowded} \Rightarrow \text{fanspeed}$$

Here, *temperature* is a normalized value for the temperature in  $[0, 1]$ , *crowded* is a normalized value for the degree of crowdedness in  $[0, 1]$ , and *fanspeed* is the normalized speed of the fan in  $[0, 1]$ . If we use *min* as aggregation, this rule means that *fanspeed* must be identical to the minimum of *temperature* and *crowdedness*. If several rules have *fanspeed* as conclusion, then it will be the maximum of the firing strengths of the rules. More formally:

**Definition 1 (Simple Positive FIS).** A Simple Positive FIS  $F$  is a set of rules of the form  $P_1 \wedge \dots \wedge P_n \xrightarrow{\phi} C$ , where  $P_1, \dots, P_n$  are input variables (premises) that have a given value in  $[0, 1]$ ,  $C$  (the conclusion) is an output variable whose value is to be computed, and  $\phi : \mathbb{R}^n \rightarrow [0, 1]$  is an aggregation function. A rule of  $F$  is satisfied if the value of the conclusion is greater than or equal to the rule strength  $\phi(P_1, \dots, P_n)$ . A solution of  $F$  is an assignment of each output variable to a value so that (1) all rules are satisfied, and (2) no output variable can be assigned a smaller value and still satisfy all rules.

**Recursive FIS.** We now extend the Simple Positive FIS by allowing the conclusions to appear as premises of rules:

**Definition 2 (Recursive FIS).** A recursive FIS is a Simple Positive FIS where the output variables can appear as premises.

In our example, a fast-running fan may attract more people seeking chilling air:

$$\text{fanspeed} \Rightarrow \text{crowded}$$

The notions of rule satisfaction and solution carry over to recursive FISs.

**Fixed point iteration.** Algorithm 1 computes the solution of a Recursive FIS. It starts by assigning all output variables the value of zero, computes the firing strength of each rule, and updates the value of the output variable of this rule. This process is iterated until convergence.

---

**Algorithm 1** Solve Recursive FIS

---

**Require:** A recursive FIS  $F$

**Ensure:** The solution  $v$  of  $F$

```

 $v(x) \leftarrow 0$  for all output variables  $x$  of  $F$ 
while no convergence do
  for rule  $r$  of  $F$  do
     $\text{strength}(r) \leftarrow$  firing strength of  $r$  with values  $v$ 
     $x \leftarrow$  output variable of  $r$ 
     $v(x) \leftarrow \max(v(x), \text{strength}(r))$ 
  end for
end while
return  $v$ 

```

---

**Theorem 1.** If each aggregation function is continuous and non-decreasing, then Algorithm 1 converges to the solution of the input recursive FIS.

*Proof.* Let  $K$  be the number of output variables. Let us see the assignment  $v$  of Algorithm 1 as a vector  $\vec{v} \in [0, 1]^K$ , with one value per output variable. Each iteration of the algorithm corresponds to the update  $\vec{v} \leftarrow f(\vec{v})$  for some

mapping  $f : [0, 1]^K \rightarrow [0, 1]^K$ . A solution to the FIS must satisfy  $\vec{v} \geq f(\vec{v})$  component-wise so that all rules are satisfied. Since  $\vec{v} \leq f(\vec{v})$  component-wise, this means that  $\vec{v} = f(\vec{v})$ , i.e.,  $\vec{v}$  is a fixed point of  $f$ . Now since  $f$  is non-decreasing, because so are the aggregation functions involved in  $f$ , as well as the *max* operator, it follows from the Knaster-Tarski fixed point theorem [46] that the set of fixed points is a lattice and that there is a unique least fixed point. It remains to prove that Algorithm 1 converges to this least fixed point.

The algorithm computes a sequence of vectors  $\vec{v}_0 = 0, \vec{v}_1, \vec{v}_2, \dots$  by applying  $f$  at each iteration. Since  $\vec{v} \leq f(\vec{v})$  component-wise for any vector  $\vec{v}$ , we have  $\vec{v}_0 \leq \vec{v}_1$  component-wise. Since  $f$  is non-decreasing, we deduce that  $f(\vec{v}_0) \leq f(\vec{v}_1)$  component-wise, that is  $\vec{v}_1 \leq \vec{v}_2$ . By induction, we obtain that the sequence  $\vec{v}_0, \vec{v}_1, \vec{v}_2, \dots$  is non-decreasing. Since this sequence is upper bounded by the vector of ones, it has some limit  $\vec{v}^*$  in  $[0, 1]^K$ , showing that the algorithm converges.

To prove that  $\vec{v}^*$  is a fixed point of  $f$ , we use the continuity of  $f$ , which again follows from the fact that the aggregation functions involved in  $f$  and the *max* operator are continuous. We have:

$$f(\vec{v}^*) = f(\lim_{t \rightarrow \infty} \vec{v}_t) = \lim_{t \rightarrow \infty} f(\vec{v}_t) = \lim_{t \rightarrow \infty} \vec{v}_{t+1} = \vec{v}^*.$$

Finally,  $\vec{v}^*$  is the least fixed point of  $f$  because for any other fixed point of  $f$ , say  $\vec{v}$ , we have  $\vec{v}_0 = 0 \leq \vec{v}$  and thus  $\vec{v}_1 = f(\vec{v}_0) \leq f(\vec{v}) = \vec{v}$  component-wise. By induction, we get  $\vec{v}_{t+1} = f(\vec{v}_t) \leq f(\vec{v}) = \vec{v}$  component-wise for all  $t \geq 0$ . Taking the limit gives  $\vec{v}^* \leq \vec{v}$  component-wise.  $\square$

Our result extends to the scenario where some or all output variables have an initial value that shall not be undercut. It suffices to add a rule of the form  $c \Rightarrow x$  for each output variable  $x$  with initial value  $c$ . This allows for an alternative definition of a recursive FIS: A recursive FIS is a set of rules of the form  $P_1 \wedge \dots \wedge P_n \xrightarrow{\phi} C$ , where  $\phi$  is an aggregation function, and  $P_1, \dots, P_n$  and  $C$  are variables, each of which has an initial value in  $[0, 1]$ . A solution is then an assignment of each variable to a value so that (1) all variables have values that are larger than or equal to their initial value, (2) all rules are satisfied, and (3) no variable can be assigned a smaller value and still satisfy all rules. Such an alternative FIS can be translated to a FIS according to Definition 2 by considering all variables that appear in conclusions as output variables, considering the others as input variables, and adding a rule  $c \Rightarrow x$  for any output variable  $x$  with initial value  $c$ .

## 5 FLORA

We now apply our framework of Fuzzy Logic to the task of KG alignment. We first extend the notion of functionality to relation lists, which is exploited in FLORA to align entities that are objects of multiple facts.

### 5.1 Relation lists

For an alphabetically sorted list  $R = (r_1, \dots, r_n)$  of relations, a list of head-entities  $H = (h_1, \dots, h_n)$ , and a tail entity  $t$ , we define  $R(H, t)$  as the boolean proposition that all corresponding facts exist:

$$R(H, t) := r_1(h_1, t) \wedge \dots \wedge r_n(h_n, t).$$

By convention, the head entities corresponding to the same relation are sorted in alphabetic order. This ensures that there is a unique way to write  $R(H, t)$  even if several relations in  $R$  are identical. We can then define the functionality of the relation list  $R$  as:

$$fun(R) = \frac{|\{H \mid \exists t, R(H, t)\}|}{|\{(H, t) \mid R(H, t)\}|}$$

For instance, the relations *BirthDateOf* and *FamilyNameOf* are themselves not functional, but their combination is, as few people have both the same birth date and family name. Similarly, the local functionality of relation list  $R$  and entity list  $H$  is defined as:

$$fun(R, H) = \frac{1}{|\{t \mid R(H, t)\}|}$$

We have  $fun(R, H) = 1$  if and only if there is a unique tail entity  $t$  such that  $R(H, t)$  is *true*. Again, this notion is used to cope with relation lists that are globally functional, but have local exceptions.

### 5.2 Alignment Rules

In order to align two KGs  $\mathcal{G}$  and  $\mathcal{G}'$ , we construct a recursive Fuzzy Inference System. The key insight, similar to PARIS [39], is the following: If there exist facts  $r(h, t)$  and  $r'(h', t')$  such that entities  $h, h'$  have already been matched and  $r, r'$  are functional, then entities  $t, t'$  are matched. FLORA goes beyond PARIS by generalizing this insight to lists of relations, by embedding it in the framework of Fuzzy Logic, and by proving its convergence.

**Entity alignment.** For each pair of entities  $t \in \mathcal{E}$  and  $t' \in \mathcal{E}'$  that are not literals, and for each list of relations  $R$  and list of head entities  $H$  with  $R(H, t)$  in  $\mathcal{G}$ , and for each list of relations  $R'$  and list of head entities  $H'$  with  $R'(H', t')$  in  $\mathcal{G}'$ , we use:

$$R(H, t) \wedge R'(H', t') \wedge H \equiv H' \wedge R \cong R' \wedge fun(R) \wedge fun(R, H) \wedge fun(R') \wedge fun(R', H') \xrightarrow{\min} t \equiv t' \quad (1)$$

Here,  $R(H, t)$  and  $R'(H', t')$  are input variables whose value is, by construction, 1.  $H \equiv H'$  is an output variable that is implied by the equivalences of each pair of head entities of  $H$  and  $H'$  through the following rule:

$$h_1 \equiv h'_1 \wedge \dots \wedge h_n \equiv h'_n \xrightarrow{\text{hmean}} H \equiv H'$$

The premises of this rule are either themselves output variables of other rules, or literals, whose similarity is computed upfront (see Section 5.3). The aggregation function of this rule is the harmonic mean  $hmean$ , which allows taking into account high evidence (which would be ignored if we used the minimum) while penalizing strong conjunctions with low evidence (unlike the arithmetic mean).

The term  $R \cong R'$  is an output variable that means “ $R$  is similar to  $R'$ ”. It is constrained by the following rule:

$$r_1 \cong r'_1 \wedge \dots \wedge r_n \cong r'_n \xrightarrow{hmean} R \cong R'$$

Here, each statement of the form  $r \cong r'$  is itself an output variable, which is implied by two rules as follows:

$$\begin{aligned} r \subseteq r' &\implies r \cong r' \\ r' \subseteq r &\implies r \cong r' \end{aligned}$$

This means that two relations are considered similar if one is a subrelation of the other (and aggregation is the identity function as there is only one premise).

The functionality terms are input variables whose values correspond to the functionality or local functionality of the respective lists of relations. Note that we cannot resort to local functionality only, due to the incompleteness of KGs. For example, people generally have two parents. If one person happens to have only a single parent in one KG, and a single parent in the other, this does not entitle us to match these two parents (as one could be, e.g. the mother and the other the father). The imposition of global functionality avoids such matches. We cannot use only global functionality either, as a globally functional relationship (such as *hasCapital*) can be non-functional in some cases (such as *South Africa*).

Finally,  $t \equiv t'$  is an output variable that determines to what degree  $t$  and  $t'$  are matched. We use the *min* aggregation function to make sure all premises of the rule are fulfilled (and a strong premise cannot make up for a weak one).

**Subrelation alignment.** To find that one relation  $r \in \mathcal{R}$  is a subrelation of another relation  $r' \in \mathcal{R}'$ , we first build the following rule for all entities  $h, t \in \mathcal{E}, h', t' \in \mathcal{E}'$  with  $r(h, t), r'(h', t')$ :

$$r(h, t) \wedge r'(h', t') \wedge h \equiv h' \wedge t \equiv t' \xrightarrow{\min} r \equiv_{h,t} r'$$

As before,  $r(h, t), r'(h', t')$  are input variables that evaluate to 1 by construction, while  $h \equiv h'$  and  $t \equiv t'$  are output variables of Equation 1. Then  $r \equiv_{h,t} r'$  is an output variable that says that  $r$  and  $r'$  coincide on  $h$  and  $t$ . The subrelation alignment can then be written as:

$$\bigwedge_{h,t:r(h,t)} r \equiv_{h,t} r' \xrightarrow{\alpha\text{-mean}} r \subseteq r' \quad (2)$$

This rule means that  $r$  is a subrelation of  $r'$  if all facts of  $r$  are also facts of  $r'$ . To make the degree of  $r \subseteq r'$  reflect the proportion of facts of  $r$  that are facts of  $r'$ ,

we use the arithmetic mean as the aggregation function. However, the arithmetic mean alone is not enough: Under the Open World Assumption, the KG  $\mathcal{G}'$  can be incomplete. Thus, some facts  $r'(h', t')$  might be missing, which would jeopardise our subrelation alignment. Therefore, we multiply the arithmetic mean by a constant  $\alpha$ , which we call the *benefit of the doubt* (we show different values of  $\alpha$  in our experiments). The aggregation function  $\alpha$ -mean is then the arithmetic mean multiplied by  $\alpha$  and capped to  $[0, 1]$ .

### 5.3 Algorithm

Our method for KG alignment is shown in Figure 2. It receives as input two knowledge graphs to be matched, optionally with training data.

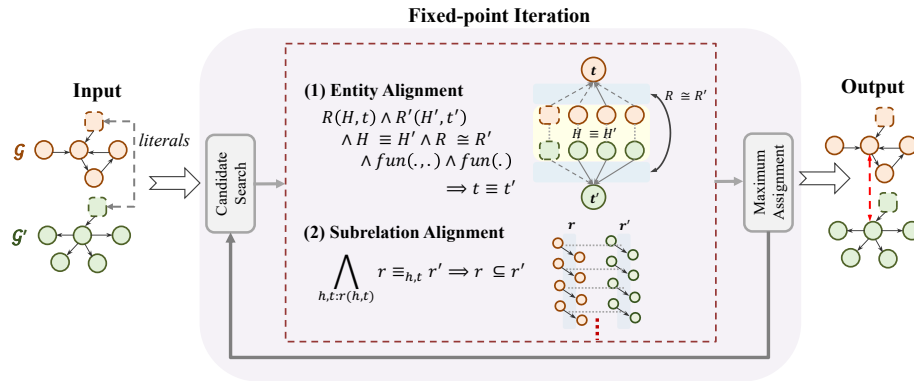


Fig. 2: Knowledge Graph Alignment with FLORA.

**Initialization.** We first compute the similarities between the literals of the two KGs (see details below). All literal similarity values are computed only once, belong to  $[0, 1]$ , and remain fixed throughout the algorithm. We integrate training data in the same way: If two entities are matched in the training data, we set their matching score to 1, and treat this match as an input variable that will never be changed. To bootstrap the rule of entity alignment of Equation 1, relation similarities are initially set to some small value  $\theta_r$ . This value is superseded by the value of the subrelation alignment of Equation 2 in the following iterations.

**Fixed point iteration.** After initialization, the steps of entity alignment (Equation 1) and subrelation alignment (Equation 2) are applied alternately. If multiple rules imply the same output variable, its score is the maximum of the rule strengths, in line with Algorithm 1. Since all our aggregation functions are continuous and non-decreasing, Theorem 1 applies after the first iteration (when

initial values have been set for both entity and relation alignments), proving the convergence of the iterative process to the solution of the corresponding FIS.

## 5.4 Implementation Details

**Similarity between literals.** The similarity between strings is computed by a (small) language model. To reduce the computational cost, we identify, for each string literal, its most similar string counterparts in the other KG by cosine similarity, and retain only those above some threshold  $\theta_s$ . Dates are matched exactly, and numbers are matched with a similarity score of 1 if they are approximately equal within a relative error of  $10^{-9}$ . While more complex matching strategies for dates are possible, experiments in Section 6 show that FLORA performs well even with simple exact matching.

**Candidate Search.** Performing a full search over all possible entity and relation combinations would be prohibitively expensive. Therefore, we first perform a candidate search to identify the most structurally similar counterpart for each entity. Specifically, we select for each entity  $t \in \mathcal{E}$  those entities  $t' \in \mathcal{E}'$  with the maximum number of counterpart triples. These triples are ranked in decreasing order of matching scores and then considered in this order.

**Maximum Assignment.** We assume that there are no duplicate entities within each KG, i.e., each entity  $e \in \mathcal{E}$  has at most one counterpart  $e' \in \mathcal{E}'$ . Therefore, we store only the highest-scored alignments in the matching set:

$$\mathcal{M}_e \leftarrow \left\{ (e, e') \mid e \equiv e' = \max_{e''} \{ e \equiv e'' \} \right\} \quad (3)$$

We do not impose this one-to-one assumption on relations, and thus store all relation pairs  $r \subseteq r'$  with scores greater than 0 in  $\mathcal{M}_r$ .

**Final alignment.** The iterations terminate when the total matching score increases by less than some threshold  $\varepsilon$  (early stopping). We keep only those entity matches that exceed a threshold  $\theta_e \in [0, 1]$ . The final alignment between two relations  $r$  and  $r'$  can be a subrelation  $r \subseteq r'$ , a superrelation  $r \supseteq r'$ , or an equivalence  $r \equiv r'$  if both  $r \subseteq r'$  and  $r' \subseteq r$  holds.

## 6 Experiments

### 6.1 Experimental Setup

**Datasets.** We evaluate FLORA on five widely-adopted entity alignment datasets, including two monolingual datasets (D-W-15K-V1/V2) from OpenEA [43] and three cross-lingual datasets from DBP15K [40]. D-W-15K is based on DBpedia and Wikidata and comes in two versions: Sparse version (D-W-15K-V1) and Dense version (D-W-15K-V2). DBP15K consists of three cross-lingual

KG pairs extracted from DBpedia: Chinese and English (DBP<sub>ZH-EN</sub>), Japanese and English (DBP<sub>JA-EN</sub>), French and English (DBP<sub>FR-EN</sub>). For holistic evaluation, we also include two datasets from the OAEI Knowledge Graph Track [19], memoryalpha-stexpanded (Mem-ST) and starwars-swtor (Star-SWT), which have non-trivial matches for relations, classes, and instances.

**Parameters.** As in PARIS [39], the relation similarity of FLORA is initialized to  $\theta_r = 0.1$ . We choose the benefit of doubt as  $\alpha = 3$  (we show other values in the ablation study). We choose  $\theta_s = 0.7$  and  $\theta_e = 0.1$ , which maximizes the F1 value in our experiments (a smaller value favors recall, a larger value favors precision). Early stopping is set to  $\varepsilon = 0.01$  (a smaller value improves precision but increases the execution time). Regarding the language models used for string similarity, we use LaBSE<sup>3</sup> [15] for the multilingual datasets, and PEARL<sup>4</sup> [7] for the monolingual datasets (because of its good performance on phrases).

**Evaluation Metrics.** On D-W-15K-V1/V2 and the OAEI KGs Track, we use standard classification-based metrics (Precision, Recall, F1), while on DBP15K, we follow existing EA baselines by reporting Hit@K and MRR, which measure top-K accuracy and mean reciprocal rank, respectively. To evaluate our model in ranking-based metrics, we directly output all possible entity pairs with scores instead of selecting the top-scoring ones.

**Competitors.** We select more than 30 competitors. For the EA task, we include the supervised methods JAPE [40], BootEA [41], RDGCN [51], OntoEA [52], PARIS+ [29], RHGN [31], GCNAlign [50], TransEdge [42], FuzzyEA [24], SelfKG [30], AttrGNN [33], TEA [60], SDEA [61], LLMEA [56], ChatEA [25] and the unsupervised methods AttrE [48], PARIS [39], PRASE [37], FGWEA [37], NALA [55], LLM4EA [8], CPL-OT [11], MultiKE [58], ICLEA [57], BERT-INT [45], HLMEA [27], and EmbMatch, a simple method using embedding similarity between entity names. For BERT-INT, which uses entity descriptions, we replace the descriptions with entity names for a fair comparison, following [61]. For holistic KG alignment, we choose classical methods LogMap [26], AML [14], and the most recent ATMatcher [20], ALOD2Vec [35], and Wiktionary [36]. We report the results available from the original papers.

## 6.2 Overall Results

Table 1 shows the results on monolingual EA in both unsupervised and supervised settings. FLORA achieves considerable improvements over all baselines, with and without training data. Table 2 shows the results for multi-lingual EA. FLORA consistently performs the best or the second-best across all datasets. FLORA does not beat FGWEA, which is designed for multilingual datasets and exploits the name bias present in such benchmarks [33]. Therefore, FGWEA

<sup>3</sup> <https://huggingface.co/sentence-transformers/LaBSE>

<sup>4</sup> [https://huggingface.co/Lihuchen/pearl\\_small](https://huggingface.co/Lihuchen/pearl_small)

Table 1: Performance comparison on the OpenEA datasets (monolingual)

Category	Method	D-W-15K-V1			D-W-15K-V2		
		Precision	Recall	F1	Precision	Recall	F1
Supervised	JAPE	0.250	0.250	0.250	0.262	0.262	0.262
	BootEA	0.572	0.572	0.572	0.821	0.821	0.821
	RHGN	0.560	0.753	0.644	-	-	-
	OntoEA	0.591	0.591	0.591	0.814	0.814	0.814
	PARIS+	<u>0.965</u>	<u>0.746</u>	<u>0.841</u>	<u>0.976</u>	<u>0.903</u>	<u>0.938</u>
	<b>Ours</b>	<b>0.967</b>	<b>0.856</b>	<b>0.908</b>	<b>0.985</b>	<b>0.966</b>	<b>0.975</b>
Unsupervised	EmbMatch	0.461	0.452	0.456	0.515	0.509	0.512
	AttrE	0.299	0.299	0.299	0.489	0.489	0.489
	PARIS	<u>0.953</u>	0.726	0.824	0.950	0.850	0.897
	PRASE	0.918	<u>0.809</u>	<u>0.860</u>	0.948	0.900	0.923
	FGWEA	0.933	0.725	0.816	<u>0.952</u>	0.903	<u>0.927</u>
	NALA	-	-	-	0.917	<u>0.908</u>	0.912
	LLM4EA	0.472	0.472	0.472	0.907	0.816	0.859
	<b>Ours</b>	<b>0.954</b>	<b>0.838</b>	<b>0.893</b>	<b>0.978</b>	<b>0.947</b>	<b>0.962</b>

performs worse than FLORA when entity names are incomplete, as observed in the monolingual EA datasets. FGWEA further assumes that each entity of one KG has one match in the other KG, which is not representative of real-world scenarios. Regarding ChatEA, it is supervised and additionally uses the background knowledge of LLMs to generate entity descriptions. In contrast, FLORA is unsupervised, relies solely on the given KGs data, and needs only a small language model for string similarity. Overall, the difference of FLORA to the best performer is less than 1 percentage point – a price we pay for a fully transparent unsupervised algorithm that works across different scenarios.

Table 3 further shows the performance of FLORA on the KG alignment task. FLORA achieves the highest overall F1 score with an average of 96%, outperforming competitors for the alignment of classes, relations, and instances. It falls within 2 percentage points of the best model for relation matching precision on the Star-SWT dataset. This slight drop is because FLORA finds structural matches (*apprentice* matches *padawan*, which also means apprentice) rather than the intended semantic matches (*apprentice* matches *apprentices*).

### 6.3 Analysis

**Ablation study.** We first experimented with different values of the benefit-of-doubt parameter  $\alpha$  (see Figure 3). We find that setting  $\alpha = 1$  is ineffective and leads to a significant reduction in performance, thus proving the necessity of  $\alpha > 1$ . Beyond  $\alpha = 2$ , the performance does not change much while computation time first decreases and then remains stable. However, a larger  $\alpha$  does not necessarily yield better results, as the performance degrades once  $\alpha$  goes beyond 100 (see Table 4). In practice, we set  $\alpha = 3$  to balance effectiveness and efficiency. We also

Table 2: Performance comparison on the DBP15K dataset across three language pairs (ZH-EN, JA-EN, FR-EN). Group (I) are structure-only methods, Group (II) uses both structure and entity names, Group (III) jointly uses relational triples, attribute triples, and entity names, and Group (IV) is LLM-based.

Group	Method		DBP <sub>ZH-EN</sub>			DBP <sub>JA-EN</sub>			DBP <sub>FR-EN</sub>		
			Hit@1	Hit@10	MRR	Hit@1	Hit@10	MRR	Hit@1	Hit@10	MRR
I	GCNAlign	S	0.413	0.744	0.549	0.399	0.745	0.546	0.373	0.745	0.532
	BootEA	SS	0.629	0.848	0.703	0.622	0.854	0.701	0.653	0.874	0.731
	TransEdge	S	0.735	0.919	0.801	0.719	0.932	0.795	0.710	0.941	0.796
II	RDGCN	S	0.708	0.846	0.749	0.767	0.895	0.812	0.886	0.957	0.908
	FuzzyEA	S	0.863	0.984	0.909	0.898	0.985	0.933	0.977	0.998	0.986
	SelfKG	U	0.745	0.866	0.782	0.816	0.913	0.844	0.957	0.992	0.971
	CPL-OT	U	0.911	0.950	0.930	0.945	0.976	0.960	0.986	0.992	0.990
	SEU	U	0.900	0.965	0.924	0.956	<u>0.991</u>	0.969	0.988	<u>0.999</u>	0.992
III	AttrGNN	S	0.796	0.929	0.845	0.783	0.921	0.834	0.919	0.978	0.910
	TEA	S	0.941	0.983	0.960	0.941	0.979	0.959	0.979	0.997	0.991
	SDEA	S	0.871	0.966	0.912	0.848	0.952	0.890	0.969	0.995	0.981
	MultiKE	U	0.509	0.576	0.532	0.393	0.489	0.432	0.639	0.712	0.665
	ICLEA	U	0.884	0.972	-	0.924	0.978	-	0.991	0.999	-
	BERT-INT <sub>name</sub>	U	0.814	0.835	0.820	0.806	0.835	0.821	0.987	0.992	0.990
	FGWEA	U	<b>0.976</b>	<b>0.994</b>	<b>0.983</b>	<b>0.978</b>	<b>0.992</b>	<b>0.988</b>	<b>0.997</b>	<b>0.999</b>	<b>0.998</b>
	<b>Ours</b>	U	<b>0.976</b>	<u>0.988</u>	<u>0.981</u>	<u>0.967</u>	0.979	0.972	<u>0.992</u>	<u>0.997</u>	<u>0.993</u>
IV	HLMEA	U	0.930	-	0.934	0.938	-	0.950	<b>0.986</b>	-	0.989
	LLMEA	S	0.890	0.923	-	0.911	0.946	-	0.957	0.977	-
	ChatEA	S	<b>0.980</b>	-	<b>0.984</b>	<b>0.985</b>	-	<b>0.993</b>	-	-	-




Fig. 3: Sensitivity to the benefit of doubt  $\alpha$

replaced the language model used for string similarity by simple string identity (see Table 4). This leads to an overall performance drop, especially on DBP<sub>ZH-EN</sub>, where entities have very different literal values across languages. But even with identity literal initialization, the score drops by 3 percentage points at most, and FLORA still outperforms strong baselines such as RDGCN and SelfKG, underscoring the strength of our structural matching. Next, we restrict the list of relations in entity alignment (Equation 1) to length 1, thereby removing the structural signal from our rules. In this setting, our system can propagate only alignment scores through flat triples, similar to PARIS [39]. This significantly degrades performance, demonstrating the importance of structural information

Table 3: Performance comparison on datasets from OAEI

		Class			Relation			Instance			Overall		
		Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
Mem-ST	BaselineLabel	1.00	0.46	0.63	0.97	0.68	0.80	0.98	0.84	0.91	0.98	0.83	0.90
	LogMap	0.78	0.54	0.64	-	-	-	0.88	0.77	0.82	0.88	0.75	0.81
	AML	1.00	0.69	0.82	0.89	0.80	0.85	0.93	0.93	0.93	0.93	0.92	0.93
	ATMatcher	1.00	0.77	0.87	0.95	0.95	0.95	0.96	0.92	0.94	0.96	0.92	0.94
	ALOD2Vec	1.00	0.62	0.76	0.87	0.95	0.91	0.92	0.93	0.93	0.92	0.93	0.92
	Wiktionary	1.00	0.62	0.76	0.87	0.95	0.91	0.92	0.93	0.93	0.92	0.93	0.92
	<b>Ours</b>	<b>1.00</b>	<b>0.85</b>	<b>0.92</b>	<b>0.98</b>	<b>0.98</b>	<b>0.98</b>	<b>0.96</b>	<b>0.95</b>	<b>0.95</b>	<b>0.96</b>	<b>0.95</b>	<b>0.95</b>
Star-SWT	BaselineLabel	1.00	0.80	0.89	<b>1.00</b>	0.89	0.94	0.95	0.84	0.89	0.95	0.84	0.89
	LogMap	1.00	0.73	0.85	-	-	-	0.94	0.78	0.86	0.94	0.75	0.84
	AML	1.00	0.87	0.93	0.98	0.71	0.82	0.93	0.90	0.92	0.93	0.90	0.91
	ATMatcher	1.00	0.87	0.93	<b>1.00</b>	0.98	<b>0.99</b>	0.94	0.91	0.92	0.95	0.91	0.93
	ALOD2Vec	1.00	0.87	0.93	0.98	0.98	0.98	0.92	0.91	0.92	0.93	0.92	0.92
	Wiktionary	1.00	0.87	0.93	0.98	0.98	0.98	0.92	0.91	0.92	0.93	0.92	0.92
	<b>Ours</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	0.98	<b>0.98</b>	0.98	<b>0.98</b>	<b>0.96</b>	<b>0.97</b>	<b>0.98</b>	<b>0.96</b>	<b>0.97</b>

Table 4: Ablation studies of FLORA on various datasets

Method	D-W-15K-V1			D-Y-15K-V2			DBP <sub>FR-EN</sub>			DBP <sub>ZH-EN</sub>		
	Prec.	Rec.	F1	Prec.	Rec.	F1	Hit@1	Hit@10	MRR	Hit@1	Hit@10	MRR
<b>FLORA</b>	<b>0.954</b>	<b>0.838</b>	<b>0.893</b>	<b>0.978</b>	<b>0.947</b>	<b>0.962</b>	<b>0.992</b>	<b>0.997</b>	<b>0.993</b>	<b>0.976</b>	<b>0.988</b>	<b>0.981</b>
- w/o literal similarity	0.939	0.797	0.862	0.962	0.925	0.943	0.975	0.983	0.977	0.947	0.965	0.951
- w/o list of relations	0.894	0.604	0.721	0.918	0.628	0.756	0.933	0.964	0.944	0.700	0.852	0.757
- with <i>min</i> aggregation	0.564	0.524	0.543	0.659	0.626	0.641	0.894	0.949	0.912	0.547	0.778	0.632
- with $\alpha = 100$	0.940	0.807	0.868	0.956	0.928	0.937	0.974	0.985	0.976	0.871	0.935	0.898

in the absence of functional relations. Finally, using the *min* aggregation function for all inference rules significantly degrades the performance across all datasets.

**Explainability.** A salient feature of FLORA is that its results are interpretable by humans. For any pair of aligned entities, the corresponding rule with the highest firing strength, as given by Equation 1, explains the alignment. In the dataset D-W-15K-V1 for instance, the entities  $t=Lady\ Gaga$  of DBPedia and  $t'=Q19848$  of Wikidata are aligned with a value of 0.984. The rule leading to this alignment states that Lady Gaga was born on *1986-03-28* and featured in the song *3-Way*; it is the alignment of the relations  $musicalArtist \cong P175$  and  $birthDate^{-1} \cong P569^{-1}$ , the functionality of the corresponding relation set, and the alignment of the birth dates and of the entities *3-Way* and Q659417, that lead to the final alignment of *Lady Gaga* and Q19848.

**Scalability and Efficiency.** Finally, we evaluate the scalability and efficiency of our approach. For scalability, we use a larger-scale entity alignment dataset from OpenEA, D-W-100K (with 100k gold alignments). As shown in Table 5, FLORA outperforms both advanced unsupervised and supervised baselines.

Table 5: Performance comparison on the OpenEA 100K datasets (monolingual)

Method	Supervised	D-W-100K-V1			D-W-100K-V2		
		Precision	Recall	F1	Precision	Recall	F1
RDGCN	Yes	0.362	0.362	0.362	0.421	0.421	0.421
BootEA	Yes	0.516	0.516	0.516	0.766	0.766	0.766
PARIS	No	0.934	0.644	0.762	0.933	0.796	0.859
PRASE	No	-	-	-	0.927	0.855	0.890
NALA	No	0.732	0.732	0.732	0.919	0.919	0.919
<b>FLORA</b>	No	<b>0.937</b>	<b>0.773</b>	<b>0.847</b>	<b>0.970</b>	<b>0.929</b>	<b>0.949</b>

For efficiency, we measured the run time of FLORA on datasets of different sizes from our previous experiments, using a Linux server with an AMD EPYC 9374F CPU and 519 GB of RAM. FLORA takes approximately 7 min on D-W-15K(V1/V2) and 48 min on the larger dataset D-W-100K-V1. This is comparable to most state-of-the-art methods [43]. However, the other methods are executed in GPUs, while our approach needs only CPUs. On the DBP15K datasets, FLORA is much slower than on the D-W-15K datasets, with around 2 hours on average. This is likely due to the amount of auxiliary information (i.e., attribute values), which is about five times larger in the DBP15K datasets than in the D-W-15K datasets. On the OAEI KG Track datasets, FLORA completes the alignment process in about 1h45min. This is much slower than the runtime of lexical-based methods (e.g, 4 min on average for ALOD2Vec) that has been reported by the OAEI platform on a virtual machine with 32GB of RAM and 16 vCPUs (2.4 GHz) [1]. This is because FLORA relies more on structural reasoning, which, in turn, leads to better performance, allows for explainability, and can find not only relation equivalence but also relation subsumption.

## 7 Conclusion and Future Work

In this paper, we have introduced FLORA, an iterative method based on Fuzzy Logic that aligns entities and relations across knowledge graphs. FLORA provably converges to the solution of the corresponding Fuzzy Inference System. It can find subrelation matches, it can deal with dangling entities, it does not need training data, and it provides interpretable results. Comparative experiments on major EA and KG alignment benchmarks show that FLORA outperforms all competitors on nearly all datasets. Future work could extend FLORA to more complex taxonomy structures, for example, by exploring cross-KGs subclasses, computing deductive closures, and taking into account the graph structure.

**Supplementary Material Statement.** The dataset, source code, supplementary material, and hyperparameters are available at our GitHub repository <https://github.com/dig-team/FLORA>.

## References

1. Abd, M., Algergawy, A., Amardeilh, F., Amini, R., Fallatah, O., Faria, D., Fundulaki, I., Harrow, I., Hertling, S., Hitzler, P., Huschka, M., Ibanescu, L., Jiménez-Ruiz, E., Karam, N., Laadhar, A., Lambrix, P., Li, H., Li, Y., Michel, F., Nasr, E., Paulheim, H., Pesquita, C., Portisch, J., Roussey, C., Saveta, T., Shvaiko, P., Splendiani, A., Trojahn, C., Vataščinová, J., Yaman, B., Zamazal, O., Zhou, L.: Results of the Ontology Alignment Evaluation Initiative 2021  $\star$ . In: Shvaiko, P., Digitale, T., Euzenat, I.J., Alpes, I.U.G., Jiménez-Ruiz, F.E., City, of London, U., SIRIUS, U., of Oslo, U., Hassanzadeh, N.O., Research, I., Trojahn, U.C., IRIT, France (eds.) Proceedings of the 16th International Workshop on Ontology Matching co-located with the 20th International Semantic Web Conference (ISWC 2021). vol. 3063, pp. 62–108. Ceur Workshop Proceedings, Virtual conference, United States (Oct 2021), <https://hal.inrae.fr/hal-05007937>
2. Abi Akl, H.: PSYCHIC: A neuro-symbolic framework for knowledge graph question-answering grounding. In: ISWC 2023-International semantic web conference (2023)
3. Baaz, M., Preining, N., Zach, R.: First-order gödel logics. *Annals of Pure and Applied Logic* **147**(1-2), 23–47 (2007)
4. Bach, S.H., Broecheler, M., Huang, B., Getoor, L.: Hinge-loss markov random fields and probabilistic soft logic. *Journal of Machine Learning Research* **18**(109), 1–67 (2017)
5. Chen, H., Bei, Y., Shen, Q., Xu, Y., Zhou, S., Huang, W., Huang, F., Wang, S., Huang, X.: Macro graph neural networks for online billion-scale recommender systems. In: Proceedings of the ACM web conference 2024. pp. 3598–3608 (2024)
6. Chen, J., Li, Z., Zhao, P., Liu, A., Zhao, L., Chen, Z., Zhang, X.: Learning short-term differences and long-term dependencies for entity alignment. In: The Semantic Web–ISWC 2020: 19th International Semantic Web Conference, Athens, Greece, November 2–6, 2020, Proceedings, Part I 19. pp. 92–109. Springer (2020)
7. Chen, L., Varoquaux, G., Suchanek, F.M.: Learning High-Quality and General-Purpose Phrase Representations (PEARL). In: EACL (2024)
8. Chen, S., Zhang, Q., Dong, J., Hua, W., Li, Q., Huang, X.: Entity alignment with noisy annotations from large language models. In: The Thirty-eighth Annual Conference on Neural Information Processing Systems (2024)
9. Chen, Z., Guo, L., Fang, Y., Zhang, Y., Chen, J., Pan, J.Z., Li, Y., Chen, H., Zhang, W.: Rethinking uncertainly missing and ambiguous visual modality in multi-modal entity alignment. In: International Semantic Web Conference (ISWC). pp. 121–139. Springer (2023)
10. Dickens, C., Augustine, E.: Negative weights in hinge-loss markov random fields. In: Workshop on Tractable Probabilistic Modeling (TPM) (2021)
11. Ding, Q., Zhang, D., Yin, J.: Conflict-aware pseudo labeling via optimal transport for entity alignment. In: 2022 IEEE International Conference on Data Mining (ICDM). pp. 915–920. IEEE (2022)
12. Dong, J., Zhang, Q., Huang, X., Duan, K., Tan, Q., Jiang, Z.: Hierarchy-aware multi-hop question answering over knowledge graphs. In: Proceedings of the ACM web conference 2023. pp. 2519–2527 (2023)
13. Euzenat, J., Shvaiko, P.: *Ontology matching*. Springer (2007)
14. Faria, D., Santos, E., Balasubramani, B.S., Silva, M.C., Couto, F.M., Pesquita, C.: Agreementmakerlight. *Semantic Web* **16**(2), SW–233304 (2025)

15. Feng, F., Yang, Y., Cer, D., Arivazhagan, N., Wang, W.: Language-agnostic bert sentence embedding. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 878–891 (2022)
16. Fernández, S., Velasco, J.R., Marsa-Maestre, I., Lopez-Carmona, M.A.: Fuzzyalign—a fuzzy method for ontology alignment. In: International Conference on Knowledge Engineering and Ontology Development. vol. 2, pp. 98–107. SciTePress (2012)
17. Fodor, J.: Left-continuous t-norms in fuzzy logic: an overview. *Acta Polytechnica Hungarica* **1**(2), 35–47 (2004)
18. Gosselin, F., Zouaq, A.: SORBET: a siamese network for ontology embeddings using a distance-based regression loss and bert. In: International Semantic Web Conference (ISWC). pp. 561–578. Springer (2023)
19. Hertling, S., Paulheim, H.: The knowledge graph track at OAEI-gold standards, baselines, and the golden hammer bias. In: The Semantic Web - 17th International Conference, ESWC (2020)
20. Hertling, S., Paulheim, H.: ATBox results for OAEI 2021. In: CEUR Workshop Proceedings. vol. 3063, pp. 137–143. RWTH Aachen (2021)
21. Hertling, S., Paulheim, H.: OLaLa: Ontology matching with large language models. In: Proceedings of the 12th Knowledge Capture Conference 2023. pp. 131–139 (2023)
22. Hnatkowska, B., Kozierekiewicz, A., Pietranik, M.: Fuzzy logic framework for ontology instance alignment. In: International Conference on Computational Science. pp. 653–666. Springer (2022)
23. Huang, S., Hu, C., Kong, W., Liu, Y.: Disentangled contrastive learning for knowledge-aware recommender system. In: International Semantic Web Conference (ISWC). pp. 140–158. Springer (2023)
24. Jiang, W., Liu, Y., Deng, X.: Fuzzy entity alignment via knowledge embedding with awareness of uncertainty measure. *Neurocomputing* **468**, 97–110 (2022)
25. Jiang, X., Shen, Y., Shi, Z., Xu, C., Li, W., Li, Z., Guo, J., Shen, H., Wang, Y.: Unlocking the power of large language models for entity alignment. In: Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 7566–7583 (2024)
26. Jiménez-Ruiz, E., Cuenca Grau, B.: LogMap: Logic-based and scalable ontology matching. In: International Semantic Web Conference (ISWC). pp. 273–288. Springer (2011)
27. Jin, X., Wang, Z., Chen, J., Yang, L., Oh, B., Hwang, S.w., Li, J.: HLMEA: Un-supervised entity alignment based on hybrid language models. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 11888–11896 (2025)
28. Kozierekiewicz, A., Pietranik, M., Jankowiak, W.: Fuzzy logic framework for ontology concepts alignment. In: Asian Conference on Intelligent Information and Database Systems. pp. 170–182. Springer (2023)
29. Leone, M., Huber, S., Arora, A., García-Durán, A., West, R.: A critical re-evaluation of neural methods for entity alignment. *Proceedings of the VLDB Endowment* **15**(8), 1712–1725 (2022)
30. Liu, X., Hong, H., Wang, X., Chen, Z., Kharlamov, E., Dong, Y., Tang, J.: SelfKG: Self-supervised entity alignment in knowledge graphs. In: Proceedings of the ACM web conference 2022. pp. 860–870 (2022)
31. Liu, X., Zhang, K., Liu, Y., Chen, E., Huang, Z., Yue, L., Yan, J.: RHGN: Relation-gated heterogeneous graph network for entity alignment in knowledge graphs. In: Findings of the Association for Computational Linguistics: ACL 2023. pp. 8683–8696 (2023)

32. Liu, Y., Wan, Y., He, L., Peng, H., Yu, P.S.: KG-BART: Knowledge graph-augmented bart for generative commonsense reasoning. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 6418–6425 (2021)
33. Liu, Z., Cao, Y., Pan, L., Li, J., Chua, T.S.: Exploring and evaluating attributes, values, and structures for entity alignment. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 6355–6364 (2020)
34. Mao, X., Wang, W., Wu, Y., Lan, M.: From alignment to assignment: Frustratingly simple unsupervised entity alignment. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. pp. 2843–2853 (2021)
35. Portisch, J., Paulheim, H.: ALOD2vec matcher results for OAEI 2021. In: CEUR Workshop Proceedings. vol. 3063, pp. 117–123. RWTH Aachen (2022)
36. Portisch, J., Paulheim, H.: Wiktionary matcher results for OAEI 2021. In: CEUR Workshop Proceedings. vol. 3063, pp. 199–206. RWTH Aachen (2022)
37. Qi, Z., Zhang, Z., Chen, J., Chen, X., Xiang, Y., Zhang, N., Zheng, Y.: Unsupervised knowledge graph alignment by probabilistic reasoning and semantic embedding. In: International Joint Conference on Artificial Intelligence (2021)
38. Sabri, N.: Fuzzy inference system: Short review and design. *International Review of Automatic Control* (01 2013)
39. Suchanek, F.M., Abiteboul, S., Senellart, P.: PARIS: Probabilistic alignment of relations, instances, and schema. *Proceedings of the VLDB Endowment* **5**(3) (2011)
40. Sun, Z., Hu, W., Li, C.: Cross-lingual entity alignment via joint attribute-preserving embedding. In: International semantic web conference (ISWC). pp. 628–644. Springer (2017)
41. Sun, Z., Hu, W., Zhang, Q., Qu, Y.: Bootstrapping entity alignment with knowledge graph embedding. In: IJCAI. vol. 18 (2018)
42. Sun, Z., Huang, J., Hu, W., Chen, M., Guo, L., Qu, Y.: TransEdge: Translating relation-contextualized embeddings for knowledge graphs. In: The Semantic Web–ISWC 2019: 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part I 18. pp. 612–629. Springer (2019)
43. Sun, Z., Zhang, Q., Hu, W., Wang, C., Chen, M., Akrami, F., Li, C.: A benchmarking study of embedding-based entity alignment for knowledge graphs. *Proceedings of the VLDB Endowment* **13**(11) (2020)
44. Tang, J., Zhao, K., Li, J.: A fused gromov-wasserstein framework for unsupervised knowledge graph entity alignment. In: 61st Annual Meeting of the Association for Computational Linguistics, ACL 2023. pp. 3320–3334. Association for Computational Linguistics (ACL) (2023)
45. Tang, X., Zhang, J., Chen, B., Yang, Y., Chen, H., Li, C.: BERT-INT: a bert-based interaction model for knowledge graph alignment. In: Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence. pp. 3174–3180 (2021)
46. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* (1955)
47. Todorov, K., Hudelot, C., Popescu, A., Geibel, P.: Fuzzy ontology alignment using background knowledge. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **22**(01), 75–112 (2014)
48. Trisedya, B.D., Qi, J., Zhang, R.: Entity alignment between knowledge graphs using attribute embeddings. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 297–304 (2019)

49. Wang, Y., Cui, Y., Liu, W., Sun, Z., Jiang, Y., Han, K., Hu, W.: Facing changes: continual entity alignment for growing knowledge graphs. In: International Semantic Web Conference (ISWC). pp. 196–213. Springer (2022)
50. Wang, Z., Lv, Q., Lan, X., Zhang, Y.: Cross-lingual knowledge graph alignment via graph convolutional networks. In: Proceedings of the 2018 conference on empirical methods in natural language processing. pp. 349–357 (2018)
51. Wu, Y., Liu, X., Feng, Y., Wang, Z., Yan, R., Zhao, D.: Relation-aware entity alignment for heterogeneous knowledge graphs. arXiv preprint arXiv:1908.08210 (2019)
52. Xiang, Y., Zhang, Z., Chen, J., Chen, X., Lin, Z., Zheng, Y.: OntoEA: Ontology-guided entity alignment via joint knowledge graph embedding. In: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021. pp. 1117–1128 (2021)
53. Xin, K., Sun, Z., Hua, W., Hu, W., Qu, J., Zhou, X.: Large-scale entity alignment via knowledge graph merging, partitioning and embedding. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. pp. 2240–2249 (2022)
54. Xin, K., Sun, Z., Hua, W., Hu, W., Zhou, X.: Informed multi-context entity alignment. In: Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining. pp. 1197–1205 (2022)
55. Xu, C., Cheng, J., Zhang, F.: NALA: an effective and interpretable entity alignment method. In: Findings of the Association for Computational Linguistics: EMNLP 2024. pp. 13752–13772 (2024)
56. Yang, L., Chen, H., Wang, X., Yang, J., Wang, F.Y., Liu, H.: Two heads are better than one: Integrating knowledge from knowledge graphs and large language models for entity alignment. arXiv preprint arXiv:2401.16960 (2024)
57. Zeng, K., Dong, Z., Hou, L., Cao, Y., Hu, M., Yu, J., Lv, X., Cao, L., Wang, X., Liu, H., et al.: Interactive contrastive learning for self-supervised entity alignment. In: Proceedings of the 31st ACM International conference on information & knowledge management. pp. 2465–2475 (2022)
58. Zhang, Q., Sun, Z., Hu, W., Chen, M., Guo, L., Qu, Y.: Multi-view knowledge graph embedding for entity alignment. In: International Joint Conference on Artificial Intelligence (2019)
59. Zhao, X., Zeng, W., Tang, J., Wang, W., Suchanek, F.M.: An experimental study of state-of-the-art entity alignment approaches. *IEEE Transactions on Knowledge and Data Engineering* **34**(6), 2610–2625 (2020)
60. Zhao, Y., Wu, Y., Cai, X., Zhang, Y., Zhang, H., Yuan, X.: From alignment to entailment: A unified textual entailment framework for entity alignment. In: Findings of the Association for Computational Linguistics: ACL 2023. pp. 8795–8806 (2023)
61. Zhong, Z., Zhang, M., Fan, J., Dou, C.: Semantics driven embedding learning for effective entity alignment. In: 2022 IEEE 38th International Conference on Data Engineering (ICDE). pp. 2127–2140. IEEE (2022)