



City Research Online

City St George's, University of London

Citation: Teymurova, S., Jiménez-Ruiz, E., Weyde, T. & Chen, J. (2025). OWL2Vec4OA: Tailoring Knowledge Graph Embeddings for Ontology Alignment. Paper presented at the Sixth International Knowledge Graph and Semantic Web Conference, 11-13 Dec 2024, Paris, France.

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/34052/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

OWL2Vec4OA: Tailoring Knowledge Graph Embeddings for Ontology Alignment

Sevinj Teymurova¹, Ernesto Jiménez-Ruiz^{1,2}, Tillman Weyde¹, and
Jiaoyan Chen³

¹ City St George’s, University of London, UK

² University of Oslo, Norway

³ University of Manchester, UK

Abstract. Ontology alignment is integral to achieving semantic interoperability as the number of available ontologies covering intersecting domains is increasing. This paper proposes OWL2Vec4OA, an extension of the ontology embedding system OWL2Vec*. While OWL2Vec* has emerged as a powerful technique for ontology embedding, it currently lacks a mechanism to tailor the embedding to the ontology alignment task. OWL2Vec4OA incorporates edge confidence values from seed mappings to guide the random walk strategy. We present the theoretical foundations, implementation details, and experimental evaluation of our proposed extension, demonstrating its potential effectiveness for ontology alignment tasks.

Keywords: ontology alignment · random walks · ontology embeddings
· knowledge graph embeddings

1 Introduction

Knowledge graphs (and ontologies) are increasingly recognized as essential for successful AI implementations across various data science applications [14]. Ontology alignment is a crucial task to enable semantic interoperability and enhance the application of knowledge graphs. The ontology alignment process involves finding and harmonizing semantic connections between different ontologies. While current methods have advanced ontology alignment considerably, there remain obstacles to their widespread implementation [20,31].

The ontology matching community has contributed to the evolution of ontology alignment systems for the last twenty years with the organization of the annual Ontology Alignment Evaluation Initiative (OAEI) [36,35]. In recent years there has been a shift from traditional systems [11], using lexical and structural techniques, to systems using machine learning and (large) language models. Prominent examples include LogMap-ML [5], BertMap [21], DeepAlignment [29], VeeAlign [26], SORBETMatcher [15] and OLaLa [24]. The OAEI, with the new Bio-ML track [23,22], has also evolved accordingly to attract and systematically evaluate such systems.

Knowledge Graph Embeddings (KGE) techniques [42,38] aim at capturing, in a low-dimensional continuous vector space, the structure and semantics of the graph. These low-dimensional representations enable the application of machine learning algorithms to graph-structured data in downstream tasks such as node classification, link prediction, or knowledge graph alignment [2]. Traditional KGE techniques commonly rely on one of the following: (i) geometric transformations, (ii) matrix factorization methods, and (iii) neural networks. Recent advancements in KGE have expanded to incorporate semantics beyond relational facts. These include encoding textual literals and integrating logical structures to capture richer semantic information within KG representations [6].

In this paper, we present OWL2Vec4OA, an extension of the ontology embedding system OWL2Vec* [4] tailored to the ontology alignment task. OWL2Vec* projects a given ontology into a graph, randomly walks over the graph to generate sequences of entities, and runs the language model Word2Vec [32] to generate embeddings of both entity URIs and words. OWL2Vec4OA, unlike OWL2Vec*, relies on potentially incomplete or inaccurate ontology alignments to bridge a given set of input ontologies. When projecting the ontologies and performing the random walks to create entity sequences, the confidence value of these seed mappings are used to bias the random walks (*i.e.*, edges with higher confidence values will have higher chances to be visited). Hence, OWL2Vec4OA allows for a tighter connection of the input ontologies given a set of seed mappings, while giving preference to edges with higher confidence.⁴ OWL2Vec4OA currently relies on the ontology matching systems LogMap [27] and AML [13] to produce seed mappings.

Our experiments show that the embeddings computed by OWL2Vec4OA are more suitable to the ontology alignment tasks than the original OWL2Vec* vectors. OWL2Vec4OA embeddings also lead to promising ranking results in the OAIE’s Bio-ML track by simply comparing the computed vectors between the relevant source and target entities.

The rest of the paper is organised as follows. Section 2 introduces the necessary notions behind OWL2Vec4OA. Section 3 presents relevant related work. OWL2Vec4OA is described in detail in Section 4. Section 5 provide experimental results of OWL2Vec4OA on the Bio-ML datasets. Finally, Section 6 concludes the paper and discusses potential lines for future work.

2 Preliminaries

Ontologies and knowledge graphs. *Ontologies* serve as structured, clearly defined representations of collectively agreed-upon concepts and relationships within a specific field or area of interest [17]. Widely applied in information retrieval, data integration, and knowledge-based systems, ontologies facilitate semantic interoperability and reasoning across diverse applications. Knowledge graphs [25] have recently gained attention to represent entities and relationships within a

⁴ Edges projected from ontology axioms are given the highest confidence.




Fig. 1. Fragment of an alignment between HeLiS and FoodOn (adapted from [5]). The green dash arrow denotes mappings with confidence values ranging from [0,1]. Blue arrows represent the inverse of the predicate `rdfs:subClassOf`.

graph-structured data model and have been very successful to improve search functionality, tailor user experiences, and inform strategic business choices [14]. Nonetheless, from the Semantic Web point of view, in essence, ontologies and knowledge graphs can be seen as equivalent notions (*i.e.*, OWL ontologies provide a formalization of the RDF graph data model [8], while knowledge graphs also imply the existence of such formalization). In this paper, we use knowledge graphs and ontologies interchangeably.

Ontology alignment is essential for data integration, semantic search, and cross-ontological reasoning. Ontology alignment can be defined as the process of identifying semantic relationships between elements (such as classes, attributes, and instances) of two or more ontologies. In this paper, we focus on atomic ontology matching where the goal is to establish equivalence or subsumption among atomic (*i.e.*, named) entities in the input ontologies [11,33]. Mappings are typically represented as a 4-tuple $\langle e, e', r, c \rangle$ where e and e' are entities from different ontologies; r is a semantic relation (*e.g.*, equivalence or subsumption); and c is a confidence value, usually, a real number within the interval $(0 \dots 1]$. For instance, Figure 1 shows a fragment of an alignment between the ontologies HeLiS [10] and FoodOn [9]. The alignment indicates that the concept HeLiS:Fructose is similar to the concept obo:FOODON_03301305 (fructose) with a confidence value of 0.9. Confidence values are typically provided by ontology alignment systems and represent the degree of certainty associated with a correspondence between entities. Alignment systems often employ sophisticated algorithms that consider lexical similarities, structural relationships, and external knowledge sources. Methods such as cross-referencing, semantic similarity measures, and machine learning techniques can be applied to establish the mappings.




Fig. 2. General architecture of LogMap.

LogMap [27,28] is an efficient ontology alignment tool for large-scale ontologies which employs lexical indexation, logic-based reasoning, and semantic similarity computation in a multi-stage process. LogMap has demonstrated effectiveness in various challenges and applications, especially in biomedicine. As shown in Figure 2, LogMap produces as output three different mappings sets: (i) LogMap overestimation are a large set of candidate mappings aiming for high recall while representing a manageable subset of all possible mappings; (ii) LogMap anchors are typically a highly precise set of mappings; and (iii) LogMap mappings are the final computed mappings aiming at a balanced Precision and Recall. These set of mappings will be used in our experiments.

Random walks are key for embedding systems like RDF2Vec [37], OWL2Vec* and node2vec [16] to capture the structural and contextual information of a graph or network, so that the system can learn about the relationships between different nodes and their local neighborhoods. The PageRank algorithm [30] revolutionized web page ranking by employing a random walk model on the web’s hyperlink structure. It simulates a “random surfer” traversing a graph of web pages, computing page importance based on the probability of the surfer landing on each page. This approach effectively captures the web’s complex link topology to determine page significance, becoming a cornerstone in information retrieval and influencing various fields beyond web search. Random walks allow to process massive graphs without needing to consider all possible paths. Seminal works like DeepWalk [34] and node2vec[16] utilize random walks to generate node sequences for training embeddings, effectively capturing network topology. Wei [43] proposed an extension of the Metropolis-Hastings algorithm for sampling from large-scale networks, introducing strategies such as early rejection and biased sampling. RDF2Vec [37] introduced random walk-based embeddings for RDF knowledge graphs, later extended by Steenwinckel [39] with new walk extraction strategies. Cochez et al. [7] also introduced biased walk strategies to RDF2Vec with twelve different edge weighting functions.

OWL2Vec. Inspired by RDF2Vec [37], OWL2Vec* [4] was designed following similar principles but adapted to create embeddings for OWL ontologies and to take into account the lexical information of the ontologies (*e.g.*, literals in the form of labels and synonyms). Figure 3 depicts the architecture of OWL2Vec* for




Fig. 3. General architecture of OWL2Vec*

generating ontology embeddings. OWL2Vec* projects the input ontology into a graph and generates entity sequences via random walks over the ontology graph, then it generates different types of documents by substituting none, some, or all entity URIs by its lexical representation. Finally, the word embedding model Word2vec [32] is applied over the generated documents to compute embeddings for both URIs and words. Note that URIs are unique and thus their embeddings are contextual. OWL2Vec* has shown to outperform other approaches in intra-ontology subsumption and class membership prediction when both structural and lexical information was critical. OWL2Vec*, however, focuses on creating embedding for a single ontology. Although a set of ontologies can also be given as input, their graph representation will not be connected and thus the random walks will not generate sequences involving elements from different ontologies.

3 Related Work

The ongoing research in the ontology matching community evidences the need for more sophisticated techniques, as shown in the annual OAEI campaign [35] new methodologies and systems are developed to address this challenging problem in the Semantic Web. Otero-Cerdeira et al. [33] provides a comprehensive survey of ontology matching techniques.

These diverse approaches demonstrate the complexity of ontology alignment and the variety of techniques employed [12]. Nonetheless, ontology matching tools can now process more efficiently even the most complex ontologies, including those with hundreds of thousands of classes, encompassing billions of possible connections. Recent advancements in ontology matching have incorporated machine learning, including embedding-based techniques and (large) language models, showing promising results in improving alignment accuracy (*e.g.*, [5], [21]). These approaches leverage vector representations of ontological elements to capture semantic relationships more effectively than traditional methods. By

utilizing language models or domain-specific embedding algorithms, these techniques can identify nuanced similarities between concepts, potentially leading to more accurate and comprehensive ontology alignments. This new generation of systems can broadly be categorized into three categories. (i) *Direct embedding comparison*: Methods like ERSOM [44] and DeepAlignment [29] calculate distances between concept embeddings directly. (ii) *Supervised mapping classifiers*: VeeAlign [26], MEDTO [18], LogMap-ML [5], and SORBET [15] train classifiers using concept embeddings as input. This approach adds a layer of learning specific to the OM task but still relies on independent embeddings for the input ontologies. (iii) *Based on language models*: Recent methods such as BERTMap [21], BERTSubs [3] and OLaLa [24] rely on language models to implement task-specific models. For instance, BERTMap fine-tunes pre-trained language models (PLMs) using synonyms from the ontologies, while BERTSubs focuses on subsumption mapping prediction using context-based information that is transformed into text by templates.

Although systems based on language models have shown impressive results, they rely on pre-trained or large language models adding an important complexity layer for large matching tasks. Our approach OWL2Vec4OA leverages a simpler language model like Word2Vec to create tailored embeddings for the ontology matching task. Although Word2Vec does not create contextual embeddings (*i.e.*, same string with different meanings will get the same embedding), the sequences that OWL2Vec4OA creates include *things* (*i.e.*, entity URIs) in addition to strings, leading to contextual embeddings for the ontology entities as URIs are unique. The embeddings are tailored to the matching task as OWL2Vec4OA bridges the input ontologies with seed mappings computed by a (traditional) alignment system like LogMap [27] and AML [13] before performing the random walks to generate the URI and word sequences. Hence, unlike other approaches like LogMap-ML, the computed embeddings for both ontologies are in the same vector space and tightly related via the seed mappings.

Our evaluation uses the datasets of the OAEI’s Bio-ML track [23], specialized benchmarks for evaluating machine learning-based OM systems. Currently, our experiments focus on direct embedding comparison, as the main purpose of this exercise was to evaluate the quality of the embeddings without any additional layer. The reported results are promising. In the near future we also plan to train a model with the computed embeddings similarly to LogMap-ML and the approach presented by Hao et al. [19].

4 Ontology embeddings with OWL2Vec4OA

OWL2Vec4OA extends the OWL2Vec $\star$ system with a mechanism to tailor the embeddings to the ontology alignment task involving two or more input ontologies. The main steps of our OWL2Vec4OA are depicted in Figure 4,⁵ and summarised as follows.

⁵ Source codes of OWL2Vec4OA are available here: <https://github.com/Sevinjt/OWL2Vec4OA>




Fig. 4. General architecture of OWL2Vec4OA

Ontology projection. We use the same ontology projection rules as in OWL2Vec* [4]. The projection rules transform one or more ontology axioms into RDF triples (*i.e.*, labeled edges in the projected graph). Some axioms such as class subsumption (*e.g.*, `obo:CHEBI_28757 (fructose) rdfs:subClassOf obo:FOODON_03420108 (sugar)`) and annotations (*e.g.*, `obo:FOODON_03420108 rdfs:label "sugar"`) have a one-to-one triple transformation; while more complex axioms require the application of projection rules. For example, the axiom `obo:FOODON_03301391 (mushroom (canned)) rdfs:subClassOf RO_0001000 (derives from) some FOODON_03411261 (fungus)` is transformed into the triple $\langle \text{obo:FOODON_03301391}, \text{RO_0001000}, \text{FOODON_03411261} \rangle$. A directed labeled graph for each of the input ontologies is returned as the output of the projection.

Ontology alignment. OWL2Vec4OA currently relies on the traditional ontology matching systems LogMap [27] and AML [13] to produce seed mappings. For example, Figure 1 shows a subset of plausible mappings computed by an alignment system. These seed mappings are used to bridge the ontology graphs and thus enabling the execution of random walks over entities from different ontologies. Note that, seed mappings do not need to be accurate nor complete and their confidence values will be used to bias the random walks. As shown in Figure 2, LogMap produces three sets of mappings of different quality that will be used as seed: (*i*) an overestimation of potential mappings ($\text{LogMap}_{\text{over}}$), (*ii*) highly precise mappings or anchors ($\text{LogMap}_{\text{anch}}$), and (*iii*) the regular output mappings ($\text{LogMap}_{\text{out}}$). In addition, $\text{LogMap}_{\text{out}}$ mappings are combined with AML mappings in our experiments (*i.e.*, $\text{LogMap}_{\text{out}} \cup \text{AML}$, and $\text{LogMap}_{\text{out}} \cap \text{AML}$).

Input: Weighted graph $G = (V, E, U, W)$; seed entities S ; walk depth wd ; iterations $iter$

Output: Walks or entity sequences W

$W = \{\}$;

```

for  $k$  in  $\text{range}(iter)$  do
    // iterates over the seed entities
    for  $e$  in  $S$  do
         $current\_walk = []$ ;
        Append  $uri(G, e)$  to  $current\_walk$ ;
         $current\_size = 1$ ;
         $focus = e$ ;
        while  $current\_size < wd$  do
            Extract set of outer edges  $E_{focus}$  from  $focus$  vertex;
            if  $|E_{focus}| = 0$  then
                break;
            end
            // According to the probabilities of the edges in
             $E_{focus}$  as in Equation 1
            Randomly select an outer edge  $l$  such that  $l = (focus, v) \in E_{focus}$ ;
            Append  $uri(G, l)$  to  $current\_walk$ ; // URI of the link
            Append  $uri(G, v)$  to  $current\_walk$ ; // URI of the vertex
             $current\_size++$ ;
             $focus = v$ ;
        end
        Add  $current\_walk$  to  $W$ 
    end
end
return  $W$ ; // Set of walks/sequences.

```

Algorithm 1: Biased Random Walks Algorithm

Graph merger and Edge Weighting. Unlike our predecessor OWL2Vec*, OWL2Vec4OA builds a single graph taking as input the graph projections of the ontologies to be aligned and a set of seed mappings. Mappings (*i.e.*, $\langle e, e', r, c \rangle$) have a direct graph representation as triples, for example, in Figure 1 the following mapping ($r = \text{equivalence}$) was identified $\langle \text{HeLiS:Fructose}, \text{owl:equivalentClass}, \text{obo:FOODON_03301305 (fructose)} \rangle$ with confidence $c = 0.9$. OWL2Vec4OA assigns a weight to each edge or link as follows: (i) 1.0 if the edge was derived from ontology axioms; and (ii) c if the edge was derived from a mapping. The output is a labeled weighted graph $G = (V, E, U, W)$, where E is the set of edges built from the projected RDF triples, V is the set of vertices composed by the subjects and objects in these triples, U is the set of URIs associated to the vertices and edges, and W is the set of weights associated to the edges. The function $weight(G, l)$ returns the weight for a given edge l , while the function $uri(G, e)$ returns the URI of a given entity e .

Biased random walks. OWL2Vec4OA, inspired by Cochez et al. [7], implements the biased random walker summarised in Algorithm 1. The algorithm takes as input a weighted labeled graph G and a set of seed entities S and performs (biased) random walks of depth wd starting from each of the seed entities. It optionally iterates over the seed entities more than once to allow for different walks for the same seed entity. In our setting, the seed entities represent the entities involved in the seed mappings computed in the alignment step. This way the walks are tailored to the alignment task, without the need of exploring the whole input ontologies. The bias in the random walk takes into account the weight assigned to each of the edges or links (l) in the graph G to assign a probability to each of the potential paths. Given $E_u = \{l_i = (u, v_i) \text{ with } i = 1..n\}$ and l_i an outer edge for u , the probability for each edge is computed as in Equation 1.

$$Pr(l_j = (u, v_j)) = \frac{weight(G, l_j)}{\sum_{i=1}^n weight(G, l_i)} \quad (1)$$

For example, a walk of depth 3 starting from the seed entity `HeLiS:Fructose` (*i.e.*, an entity appearing in a mapping) could include the following sequence of URIs: `HeLiS:Fructose`, `owl:equivalentClass`, `obo:FOODON_03301305` (fructose), `rdfs:subClassOf`, `obo:FOODON_03420108` (sugar).

Document generator and Word2Vec embeddings. OWL2Vec4OA, as in OWL2Vec*, creates three types of documents from the generated walks W in the previous step: (i) structure document, (ii) lexical document, and (iii) combined document. The structure document is a direct representation of the walks as the sentences are composed by entity URIs. The lexical document replaces every URI occurrence in the walks by the respective lexical representation of the entity (*i.e.*, the occurrence of `obo:FOODON_03420108` is replaced by its lexical representation “sugar”, typically provided in the ontology via an `rdfs:label` annotation). Finally, the combined document randomly replaces in each walk some of the entity URI occurrences by its associated label. The three documents are merged and used to train a Word2Vec model with the skip-gram architecture. The trained Word2Vec model produces embeddings for both URI and word occurrences in the merged document. As introduced in Section 2, the URI embeddings can be seen as contextual embeddings as URIs are unique.

5 Evaluation

We have performed a preliminary evaluation of the suitability of the embeddings computed by OWL2Vec4OA in ontology alignment tasks. Particularly we have used the datasets provided by the OAEI’s 2023 Bio-ML track⁶. The Bio-ML track included several tasks (*e.g.*, OMIM-ORDO, NCIT-DOID, SNOMED-NCIT-Pharm and SNOMED-NCI-Neoplas), involving biomedical ontologies with

⁶ Bio-ML Challenge [23]: <https://krr-oxford.github.io/OAEI-Bio-ML/>. Bio-ML 2023 Datasets: <https://doi.org/10.5281/zenodo.8193375>

Table 1. OMIM-ORDO task with Walk depth 3, Walker iteration: iter=1 / iter=5.

Seed Mappings	Hits@1	Hits@5	Hits@10	MRR
Train-Validation	0.01 / 0.01	0.02 / 0.02	0.05 / 0.05	0.04 / 0.04
LogMap _{over}	0.27 / 0.28	0.51 / 0.53	0.60 / 0.62	0.38 / 0.40
LogMap _{anch}	0.11 / 0.26	0.28 / 0.41	0.36 / 0.46	0.20 / 0.33
LogMap _{out}	0.26 / 0.27	0.41 / 0.45	0.47 / 0.53	0.33 / 0.36
LogMap _{out} $\cup$ AML	0.30 / 0.31	0.54 / 0.54	0.61 / 0.61	0.41 / 0.41
LogMap _{out} $\cap$ AML	0.31 / 0.34	0.49 / 0.50	0.54 / 0.54	0.40 / 0.41

tens of thousands of classes, and reference alignments based on Mondo [41] and UMLS [1].

Bio-ML presents two evaluation settings: global matching and local ranking. Global matching is evaluated with the traditional measures Precision and Recall, comparing a set of system-computed mappings with the reference set of mappings; while local matching evaluates the capacity of a system to rank a correct mapping given a pool of potential candidates. Bio-ML uses Mean Reciprocal Rank (MRR) and Hits@K (*i.e.*, cases where the correct mapping was ranked within the top-k) in the local matching setting.

Scoring function and settings. We have applied OWL2Vec4OA embeddings into the local matching tasks of Bio-ML. Mappings are scored and ranked according to the cosine similarity of the computed URI embeddings for the entities in the mapping. We have computed OWL2Vec4OA embeddings for different walk depths and iterations. We fixed the Word2Vec hyperparameters — the number of epochs and embedding dimension to 70 and 100, respectively. The experiments were conducted on a High-Performance Computing cluster with access to up to 48 CPUs, using the Slurm workload manager to ensure efficient resource allocation and job scheduling. Generated resources are available in Zenodo [40].

Impact of seed mappings and number of iterations. Table 1 shows the results over the OMIM-ORDO for different sets of seed mappings as introduced in Section 4. We also used as seed the mappings provided as training and validation in Bio-ML, as shown in the first row, using training-validation mappings in isolation did not lead to promising results given their reduced size. The set of seed mappings leading to the best results was the union of LogMap_{out} and AML mappings, with LogMap_{out} $\cap$ AML and LogMap_{over} leading to similar results. The results also show that OWL2Vec4OA is also able to handle noisy set of seed mappings like LogMap_{over}. The impact of additional iterations over the seed entities did not lead to a significantly increased performance.

Comparison with OWL2Vec.* Following the results in Table 1, we set LogMap_{out} $\cup$ AML as the seed mappings and the number of iterations to 1 in the subsequent experiments. We experimented with walk depths ranging from 2 to 4. We compared the performance of the embeddings computed with OWL2Vec4OA with

Table 2. Results of OWL2Vec4OA and OWL2Vec* over four Bio-ML tasks, with different walk depths (wd).

Task	System	wd	MRR	Hits@1	Hits@5	Hits@10	Hits@20	Hits@30
OMIM-ORDO	OWL2Vec*	2	0.074	0.018	0.091	0.178	0.332	0.393
		3	0.073	0.018	0.090	0.170	0.318	0.381
		4	0.071	0.019	0.078	0.320	0.321	0.387
	OWL2Vec4OA	2	0.586	0.533	0.637	0.657	0.672	0.693
		3	0.402	0.306	0.512	0.587	0.650	0.685
		4	0.215	0.132	0.281	0.359	0.446	0.532
NCIT-DOID	OWL2Vec*	2	0.218	0.110	0.306	0.448	0.631	0.746
		3	0.175	0.074	0.251	0.377	0.561	0.690
		4	0.105	0.035	0.121	0.225	0.409	0.541
	OWL2Vec4OA	2	0.195	0.064	0.310	0.508	0.709	0.812
		3	0.358	0.181	0.573	0.741	0.872	0.924
		4	0.609	0.442	0.840	0.928	0.970	0.984
SNOMED-NCIT-N	OWL2Vec*	2	0.063	0.014	0.075	0.134	0.231	0.309
		3	0.068	0.017	0.079	0.142	0.238	0.308
		4	0.055	0.011	0.052	0.114	0.218	0.305
	OWL2Vec4OA	2	0.648	0.543	0.767	0.831	0.888	0.904
		3	0.605	0.484	0.746	0.813	0.872	0.899
		4	0.805	0.747	0.872	0.888	0.902	0.910
SNOMED-NCIT-P	OWL2Vec*	2	0.079	0.018	0.094	0.184	0.302	0.675
		3	0.078	0.018	0.092	0.181	0.292	0.667
		4	0.055	0.011	0.052	0.114	0.218	0.305
	OWL2Vec4OA	2	0.436	0.342	0.534	0.583	0.609	0.967
		3	0.311	0.190	0.435	0.502	0.558	0.933
		4	0.291	0.204	0.355	0.434	0.521	0.944

those computed with the original OWL2Vec* version (using its multi-ontology setting). As expected, Table 2 shows that the ranking with OWL2Vec4OA embeddings considerably outperforms the ranking with OWL2Vec* embeddings in all tasks and for all evaluated walk depths, indicating that the OWL2Vec4OA embeddings are more suitable for ontology alignment tasks. The best results are obtained for the task SNOMED-NCIT-Neoplas with walk depth 4 where Hits@1 reach more than 80% of the cases. In other tasks, the results are also promising indicating that the embeddings computed by OWL2Vec4OA capture relevant features of the original entities that could be exploited by a subsequent machine learning model.

Impact of the walk depth. Increasing the walk depths has a positive impact in the tasks NCIT-DOID and SNOMED-NCIT-Neoplas; while for OMIM-ORDO and SNOMED-NCIT-Pharm longer paths seem to add noise to the embeddings. This is inline with the results obtained in the original OWL2Vec* paper [4] where longer paths did not seem to lead to better results. It is worth mentioning that longer paths also increase the computation times. In the near future, we plan to perform an extended evaluation to better understand the impact of longer walks on different ontologies and matching tasks.

6 Conclusions and Future Work

We have presented OWL2Vec4OA, an extension of the ontology embedding system OWL2Vec* [4]. OWL2Vec4OA has been tailored to the ontology alignment task by using a preliminary set of ontology alignments, possibly incomplete or inaccurate, to bridge a given set of input ontologies. These seed mappings and their confidence are key when performing biased random walks to create sequences of entities from both ontologies. The results section shows promising results where the OWL2Vec4OA embeddings lead to much better-ranking results than those computed by OWL2Vec*.

Currently, our experiments rely on direct embedding comparison which leads to good similarity results. However, predicting equivalent or subsumption mappings is a more complex task. In the near future, we aim at training machine learning models to better benefit from the features of the OWL2Vec4OA embeddings for an ontology alignment task. Prominent examples in the literature are LogMap-ML [5], which successfully applied a Siamese Neural Network; and Hao et al. [19], which explored the use of Graph Neural Networks (GNN). These approaches, however, created embeddings that were independent for each input ontology, unlike those computed by OWL2Vec4OA.

In addition, we also plan to conduct additional experiments to better understand the impact of the walk depth with different strategies to create entity sequences (*i.e.*, focusing on concepts and/or avoiding OWL constructs). Entity embedding can also be constructed using the word embedding associated to their labels, which may bring additional features with respect to the URI embeddings. Finally, once we have an end-to-end ontology alignment system in place, we aim to participate in the OAEI campaign and perform an extensive comparison with the state-of-the-art.

Acknowledgements

This research is funded by the Ministry of Education and Science of Azerbaijan Republic with support from City St George's, University of London. This work has also been partially supported by the Academy of Medical Sciences Network Grant (Neurosymbolic AI for Medicine, NGR1\1857), the project "XAI4SOC: Explainable Artificial Intelligence for Healthy Aging and Social Wellbeing" funded by the Agencia Estatal de Investigación (AEI), the Spanish Ministry of Science, Innovation and Universities and the European Social Funds (PID2021-123152OB-C22), the EPSRC project OntoEm (EP/Y017706/1), and the EU Projects: RE4DY (101058384, HORIZON-CL4-2021), Plooto (101092008, HORIZON-CL4-2022), SM4RTENANCE (101123490, DIGITAL-2022), and Tec4MaasEs (101138517, HORIZON-CL4-2023).

References

1. Bodenreider, O.: The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Res.* **32**(Database-Issue), 267–270 (2004). <https://doi.org/10.1093/NAR/GKH061>, <https://doi.org/10.1093/nar/gkh061>
2. Cai, H., Zheng, V.W., Chang, K.C.C.: A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE transactions on knowledge and data engineering* **30**(9), 1616–1637 (2018)
3. Chen, J., He, Y., Geng, Y., Jiménez-Ruiz, E., Dong, H., Horrocks, I.: Contextual semantic embeddings for ontology subsumption prediction. *World Wide Web* **26**(5), 2569–2591 (2023)
4. Chen, J., Hu, P., Jimenez-Ruiz, E., Holter, O.M., Antonyrajah, D., Horrocks, I.: OWL2vec*: Embedding of OWL ontologies. *Machine Learning* **110**(7), 1813–1845 (2021)
5. Chen, J., Jiménez-Ruiz, E., Horrocks, I., Antonyrajah, D., Hadian, A., Lee, J.: Augmenting ontology alignment by semantic embedding and distant supervision. In: *The Semantic Web: ESWC*. pp. 392–408. Springer (2021)
6. Chen, J., Mashkova, O., Zhapa-Camacho, F., Hoehndorf, R., He, Y., Horrocks, I.: Ontology Embedding: A Survey of Methods, Applications and Resources. *arXiv preprint arXiv:2406.10964* (2024)
7. Cochez, M., Ristoski, P., Ponzetto, S.P., Paulheim, H.: Biased graph walks for RDF graph embeddings. In: Akerkar, R., Cuzzocrea, A., Cao, J., Hacid, M. (eds.) *Proceedings of the 7th International Conference on Web Intelligence, Mining and Semantics*. pp. 21:1–21:12. ACM (2017). <https://doi.org/10.1145/3102254.3102279>, <https://doi.org/10.1145/3102254.3102279>
8. Cuenca Grau, B., Horrocks, I., Motik, B., Parsia, B., Patel-Schneider, P.F., Sattler, U.: OWL 2: The next step for OWL. *J. Web Semant.* **6**(4), 309–322 (2008). <https://doi.org/10.1016/J.WEBSEM.2008.05.001>, <https://doi.org/10.1016/j.websem.2008.05.001>
9. Dooley, D.M., Griffiths, E.J., Gosal, G.S., Buttigieg, P.L., Hoehndorf, R., Lange, M.C., Schriml, L.M., Brinkman, F.S., Hsiao, W.W.: FoodOn: A harmonized food ontology to increase global food traceability, quality control and data integration. *npj Science of Food* **2**(1), 1–10 (2018)
10. Dragoni, M., Bailoni, T., Maimone, R., Eccher, C.: HeLis: An ontology for supporting healthy lifestyles. In: *ISWC*. pp. 53–69. Springer (2018)
11. Euzenat, J., Meilicke, C., Stuckenschmidt, H., Shvaiko, P., dos Santos, C.T.: Ontology alignment evaluation initiative: Six years of experience. *J. Data Semant.* **15**, 158–192 (2011). https://doi.org/10.1007/978-3-642-22630-4_6, https://doi.org/10.1007/978-3-642-22630-4_6
12. Faria, D., Pesquita, C., Mott, I., Martins, C., Couto, F.M., Cruz, I.F.: Tackling the challenges of matching biomedical ontologies. *J. Biomed. Semant.* **9**(1), 4:1–4:19 (2018). <https://doi.org/10.1186/S13326-017-0170-9>, <https://doi.org/10.1186/s13326-017-0170-9>
13. Faria, D., Pesquita, C., Santos, E., Palmonari, M., Cruz, I.F., Couto, F.M.: The agreementmakerlight ontology matching system. In: *On the Move to Meaningful Internet Systems. Lecture Notes in Computer Science*, vol. 8185, pp. 527–541. Springer (2013). https://doi.org/10.1007/978-3-642-41030-7_38, https://doi.org/10.1007/978-3-642-41030-7_38
14. Fensel, D.A., Simsek, U., Angele, K., Huaman, E., Kärle, E., Panasiuk, O., Toma, I., Umbrich, J., Wahler, A.: *Knowledge Graphs: Methodology, Tools and Selected Use Cases*. Springer (2020), <https://api.semanticscholar.org/CorpusID:210975360>

15. Gosselin, F., Zouaq, A.: SORBET: A Siamese Network for Ontology Embeddings Using a Distance-Based Regression Loss and BERT. In: International Semantic Web Conference. pp. 561–578. Springer (2023)
16. Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 855–864 (2016)
17. Gruber, T.R.: Toward principles for the design of ontologies used for knowledge sharing? *International journal of human-computer studies* **43**(5-6), 907–928 (1995)
18. Hao, J., Lei, C., Efthymiou, V., Quamar, A., Özcan, F., Sun, Y., Wang, W.: Medto: Medical data to ontology matching using hybrid graph neural networks. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining. pp. 2946–2954 (2021)
19. Hao, Z., Mayer, W., Xia, J., Li, G., Qin, L., Feng, Z.: Ontology alignment with semantic and structural embeddings. *J. Web Semant.* **78**, 100798 (2023). <https://doi.org/10.1016/J.WEBSEM.2023.100798>, <https://doi.org/10.1016/j.websem.2023.100798>
20. Harrow, I., Balakrishnan, R., Jiménez-Ruiz, E., Jupp, S., Lomax, J., Reed, J.L., Romacker, M., Senger, C., Splendiani, A., Wilson, J., Woollard, P.M.: Ontology mapping for semantically enabled applications. *Drug discovery today* (2019)
21. He, Y., Chen, J., Antonyrajah, D., Horrocks, I.: BERTMap: A BERT-Based Ontology Alignment System. In: Thirty-Sixth AAAI Conference on Artificial Intelligence (2022)
22. He, Y., Chen, J., Dong, H., Horrocks, I.: Exploring large language models for ontology alignment. In: Fundulaki, I., Kozaki, K., Garijo, D., Gómez-Pérez, J.M. (eds.) *Proceedings of the ISWC 2023 Posters, Demos and Industry Tracks. CEUR Workshop Proceedings*, vol. 3632. CEUR-WS.org (2023), https://ceur-ws.org/Vol-3632/ISWC2023_paper_427.pdf
23. He, Y., Chen, J., Dong, H., Jiménez-Ruiz, E., Hadian, A., Horrocks, I.: Machine Learning-Friendly Biomedical Datasets for Equivalence and Subsumption Ontology Matching. In: 21st International Semantic Web Conference. *Lecture Notes in Computer Science*, vol. 13489, pp. 575–591. Springer (2022). https://doi.org/10.1007/978-3-031-19433-7_33, https://doi.org/10.1007/978-3-031-19433-7_33
24. Hertling, S., Paulheim, H.: OLaLa: Ontology Matching with Large Language Models. In: Venable, K.B., Garijo, D., Jalaian, B. (eds.) *Proceedings of the 12th Knowledge Capture Conference (K-CAP)*. pp. 131–139. ACM (2023). <https://doi.org/10.1145/3587259.3627571>, <https://doi.org/10.1145/3587259.3627571>
25. Hogan, A., Blomqvist, E., Cochez, M., d’Amato, C., Melo, G.D., Gutierrez, C., Kirrane, S., Gayo, J.E.L., Navigli, R., Neumaier, S., et al.: Knowledge graphs. *ACM Computing Surveys (Csur)* **54**(4), 1–37 (2021)
26. Iyer, V., Agarwal, A., Kumar, H.: VeeAlign: Multifaceted Context Representation Using Dual Attention for Ontology Alignment. In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pp. 10780–10792. Association for Computational Linguistics (2021). <https://doi.org/10.18653/V1/2021.EMNLP-MAIN.842>, <https://doi.org/10.18653/v1/2021.emnlp-main.842>
27. Jimenez-Ruiz, E., Cuenca Grau, B.: LogMap: Logic- Based and Scalable Ontology Matching. *The Semantic Web – ISWC* p. vol 7031 (2011)
28. Jiménez-Ruiz, E., Cuenca Grau, B., Zhou, Y., Horrocks, I.: Large-scale interactive ontology matching: Algorithms and implementation. In: Raedt, L.D., Bessiere, C., Dubois, D., Doherty, P., Frasconi, P., Heintz, F., Lucas, P.J.F. (eds.) *ECAI*

- 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31, 2012. *Frontiers in Artificial Intelligence and Applications*, vol. 242, pp. 444–449. IOS Press (2012). <https://doi.org/10.3233/978-1-61499-098-7-444>, <https://doi.org/10.3233/978-1-61499-098-7-444>
29. Kolyvakis, P., Kalousis, A., Kiritsis, D.: DeepAlignment: Unsupervised ontology matching with refined word vectors. In: *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1-6 June 2018 (2018)
 30. Lawrence Page, R.: The PageRank citation ranking: Bringing order to the web. Technical report, Stanford Digital Library Technologies Project (1998)
 31. Li, H., Dragisic, Z., Faria, D., Ivanova, V., Jiménez-Ruiz, E., Lambrix, P., Pesquita, C.: User validation in ontology alignment: functional assessment and impact. *Knowl. Eng. Rev.* **34** (2019)
 32. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems* **26** (2013)
 33. Otero-Cerdeira, L., Rodríguez-Martínez, F.J., Gómez-Rodríguez, A.: Ontology matching: A literature review. *Expert Systems with Applications* **42**(2), 949–971 (2015)
 34. Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: Online learning of social representations. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 701–710 (2014)
 35. Pour, M.A.N., Algergawy, A., Buche, P., Castro, L.J., Chen, J., Coulet, A., Cufi, J., Dong, H., Fallatah, O., Faria, D., Fundulaki, I., Hertling, S., He, Y., Horrocks, I., Huschka, M., Ibanescu, L., Jain, S., Jiménez-Ruiz, E., Karam, N., Lambrix, P., Li, H., Li, Y., Monnin, P., Nasr, E., Paulheim, H., Pesquita, C., Saveta, T., Shvaiko, P., Sousa, G., Trojahn, C., Vatasacinová, J., Wu, M., Yaman, B., Zamazal, O., Zhou, L.: Results of the Ontology Alignment Evaluation Initiative 2023. In: Shvaiko, P., Euzenat, J., Jiménez-Ruiz, E., Hassanzadeh, O., Trojahn, C. (eds.) *Proceedings of the 18th International Workshop on Ontology Matching (OM)*. CEUR Workshop Proceedings, vol. 3591, pp. 97–139. CEUR-WS.org (2023), https://ceur-ws.org/Vol-3591/oeai23_paper0.pdf
 36. Pour, M.A.N., Algergawy, A., Buche, P., Castro, L.J., Chen, J., Dong, H., Fallatah, O., Faria, D., Fundulaki, I., Hertling, S., He, Y., Horrocks, I., 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., Saveta, T., Shvaiko, P., Trojahn, C., Verhey, C., Wu, M., Yaman, B., Zamazal, O., Zhou, L.: Results of the ontology alignment evaluation initiative 2022. In: Shvaiko, P., Euzenat, J., Jiménez-Ruiz, E., Hassanzadeh, O., Trojahn, C. (eds.) *Proceedings of the 17th International Workshop on Ontology Matching (OM)*. CEUR Workshop Proceedings, vol. 3324, pp. 84–128. CEUR-WS.org (2022), https://ceur-ws.org/Vol-3324/oeai22_paper0.pdf
 37. Ristoski, P., Rosati, J., Noia, T.D., Leone, R.D., Paulheim, H.: RDF2Vec: RDF graph embeddings and their applications. vol. 10, pp. 721–752 (2019). <https://doi.org/10.3233/SW-180317>, <https://doi.org/10.3233/SW-180317>
 38. Rossi, A., Barbosa, D., Firmani, D., Matinata, A., Merialdo, P.: Knowledge Graph Embedding for Link Prediction: A Comparative Analysis. *ACM Trans. Knowl. Discov. Data* **15**(2), 14:1–14:49 (2021). <https://doi.org/10.1145/3424672>, <https://doi.org/10.1145/3424672>

39. Steenwinckel, B., Vandewiele, G., Agozzino, T., Ongenae, F.: pyRDF2Vec: A Python Implementation and Extension of RDF2Vec. In: European Semantic Web Conference. pp. 471–483. Springer (2023)
40. Teymurova, S.: OWL2VecOA Resources (Aug 2024). <https://doi.org/10.5281/zenodo.13217801>, <https://doi.org/10.5281/zenodo.13217801>
41. Vasilevsky, N.A., Matentzoglou, N.A., Toro, S., Flack, J.E., Hegde, H., Unni, D.R., Alyea, G.F., Amberger, J.S., Babb, L., Balhoff, J.P., Bingaman, T.L., Burns, G.A., Buske, O.J., Callahan, T.J., Carmody, L.C., Cordo, P.C., Chan, L.E., Chang, G.S., Christiaens, S.L., Daugherty, L.C., Dumontier, M., Failla, L.E., Flowers, M.J., Garrett, H.A., Goldstein, J.L., Gration, D., Groza, T., Hanauer, M., Harris, N.L., Hilton, J.A., Himmelstein, D.S., Hoyt, C.T., Kane, M.S., Köhler, S., Lagorce, D., Lai, A., Larralde, M., Lock, A., Santiago, I.L., Maglott, D.R., Malheiro, A.J., Meldal, B.H.M., Munoz-Torres, M.C., Nelson, T.H., Nicholas, F.W., Ochoa, D., Olson, D.P., Oprea, T.I., Osumi-Sutherland, D., Parkinson, H., Pendlington, Z.M., Rath, A., Rehm, H.L., Remennik, L., Riggs, E.R., Roncaglia, P., Ross, J.E., Shadbolt, M.F., Shefchek, K.A., Similuk, M.N., Sioutos, N., Smedley, D., Sparks, R., Stefancsik, R., Stephan, R., Storm, A.L., Stupp, D., Stupp, G.S., Sundaramurthi, J.C., Tammen, I., Tay, D., Thaxton, C.L., Valasek, E., Valls-Margarit, J., Wagner, A.H., Welter, D., Whetzel, P.L., Whiteman, L.L., Wood, V., Xu, C.H., Zankl, A., Zhang, X.A., Chute, C.G., Robinson, P.N., Mungall, C.J., Hamosh, A., Haendel, M.A.: Mondo: Unifying diseases for the world, by the world. medRxiv (2022). <https://doi.org/10.1101/2022.04.13.22273750>, <https://www.medrxiv.org/content/early/2022/05/03/2022.04.13.22273750>
42. Wang, Q., Mao, Z., Wang, B., Guo, L.: Knowledge Graph Embedding: A Survey of Approaches and Applications. *IEEE Trans. Knowl. Data Eng.* **29**(12), 2724–2743 (2017). <https://doi.org/10.1109/TKDE.2017.2754499>, <https://doi.org/10.1109/TKDE.2017.2754499>
43. Wei, W., Erenrich, J., Selman, B.: Towards efficient sampling: Exploiting random walk strategies. In: AAAI. vol. 4, pp. 670–676. Citeseer (2004)
44. Xiang, C., Jiang, T., Chang, B., Sui, Z.: ERSOM: A structural ontology matching approach using automatically learned entity representation. In: Proceedings of the 2015 conference on empirical methods in natural language processing. pp. 2419–2429 (2015)