

A-LION - Alignment Learning through Inconsistency negatives of the aligned Ontologies

Sarah M. Alghamdi^{1,2}, Fernando Zhapa-Camacho¹ and Robert Hoehndorf¹

¹Computational Bioscience Research Center, Computer, Electrical & Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology, 4700 KAUST, 23955 Thuwal, Saudi Arabia

²King Abdul-Aziz University, Faculty of Computing and Information Technology, Rabigh, 25732, Kingdom of Saudi Arabia

Abstract

Ontologies play an important role in sharing and reusing knowledge. Several ontologies have been developed to describe a particular domain but from different perspectives from communities of developers and users. This has led to the existence of multiple ontologies covering the same or a different domain with varying degrees of variability. Ontology Alignment is typically used to identify correspondences between semantically related elements of two or more ontologies in order to address this problem.

We propose A-LION a system that learns alignments by combining lexical and semantic approaches as well as machine learning. The system utilizes OWL EL reasoning for negative sampling which is iteratively used to inform the correction of the learning of the alignments. We demonstrate that A-LION produces alignments that are coherent with respect to OWL EL.

Keywords

Ontology Alignments, Ontology matching, Inconsistency negatives

1. Presentation of the System

Alignment Learning through Inconsistency negatives of the aligned Ontologies (A-LION) is a system that discovers alignments between ontologies by combining various matching techniques, ranging from entity-level label matching to structure-level taxonomy learning and graph projection to logical reasoning and inconsistency detection and learning. This is the first participation of A-LION in the Ontology Alignment Evaluation Initiative (OAEI).

2. Proposed Methods

An ontology \mathcal{O} can be defined over a signature $\mathcal{O} := (C, R, I; ax)$, where C is a set of concept names, R is a set of relation names, I is a set of individual names, and ax is a set of axioms.

Given two ontologies $\mathcal{O}_s, \mathcal{O}_t$, the purpose of ontology alignment is to find the pairs of entities $(e_{\mathcal{O}_s}, e_{\mathcal{O}_t}) \in C_{\mathcal{O}_s} \times C_{\mathcal{O}_t}$ that are considered as being equivalent or standing in a subclass relation

OAEI 2022

✉ sarah.alghamdi.1@kaust.edu.sa (S. M. Alghamdi); fernando.zhapacamacho@kaust.edu.sa (F. Zhapa-Camacho); robert.hoehndorf@kaust.edu.sa (R. Hoehndorf)

🆔 0000-0001-5544-7166 (S. M. Alghamdi); 0000-0002-0710-2259 (F. Zhapa-Camacho); 0000-0001-8149-5890 (R. Hoehndorf)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

within certain contexts.

A graph is defined as a tuple $G = (E, R, T)$, where E is a set of entities names, R is a set of relations names and $T \subseteq E \times R \times E$ is a set of triples of the form $(head, relation, tail)$.

A projection of an ontology into a graph is a mapping $f : \mathcal{O} \rightarrow G$ that maps the ontology classes into graph nodes, ontology roles as graph relations, and ontology axioms as graph triples following a particular set of rules.

Our method A-LION combines different matching techniques and consists of four main components (see Figure 1):

- Learning lexical matching seeds.
- Graph construction from source and target ontology.
- Graph embedding and transformation learning.
- Consistency checking.

Those components cover element-wise, structure-wise, and formal semantics learning techniques. (a) Element-wise techniques consider the entities in the ontology in *isolation* in order to find alignments disregarding the fact that they are part of the ontology's structure. This means that we use the information belonging to an ontology class itself such as its textual annotations and labels only. (b) On the other hand, structure-wise techniques analyze the entities as part of their structure. In our case, we focus on adjacency structure within the ontology and extract the structure in the form of a graph. (c) Finally, the semantic component consists of employing formal semantics learning techniques and logical inference to identify correspondences and repair inconsistencies.

2.1. Learning Lexical matching seeds

To begin learning ontology alignment, we need some known-to-be-positive seed alignments. We chose to align the classes of both ontologies with the same IRI, or lexically matched labels and relative IRIs. For lexical matching, we utilize fuzzy lexical matching, a method for finding approximate string matching with a retrieved score representing the similarity between one string to another. We begin with an exact matching score and then we decrease the threshold iteratively until a sufficient number of seeds are obtained or a minimal accepted threshold is reached. The number of matching seeds required is a parameter of our method.

2.2. Ontology Projection

We project each ontology as a graph in order to learn structure-level information from the source and the target ontologies. We evaluate two graph construction techniques:

- Subsumption hierarchy: in this method, we only utilized the subclass axioms asserted between the ontology classes to generate a directed graph for the source ontology and the target ontology. We evaluated this technique for Anatomy, Conference, Biodiversity and Ecology, and Material Sciences and Engineering tracks.




Figure 1: A-LION system component workflow

- **OWL Projection:** This method was proposed in [1] where OWL axioms are transformed directly into edges in the graph, and complex axioms are approximated in the graph to avoid the use of blank nodes. Despite the fact that this transformation method does not preserve exact logical relations, it enables correlation and learning alignments between classes of the source and target ontologies as well as within the same ontology. We evaluated this technique using the Phenotype ontology alignment.

2.3. Transformation Learning

After projecting an ontology, the result is a graph. Depending on the chosen projection method (Section 2.2), these graphs would encode the taxonomical structure or relational information found in the ontologies.

In our method, we start with two ontologies \mathcal{O}_s (source) and \mathcal{O}_t (target), which, after applying the graph projection, will become two graphs G_s, G_t , respectively. When we deal with two graphs, there are several graph alignment methods that can align two graphs from a small number of seed alignments; we follow the method in [2].

To learn representations of the two graphs G_s, G_t , we define two vector spaces V_s, V_t , where the entities (nodes and edges) of each graph will be processed separately. To learn the graph embeddings we rely on knowledge graph embeddings methods such as TransR [3], optimizing the following loss function:

$$O_{s_{\text{loss}}} = \|M_{r_s} \cdot c_{si} + r_s - M_{r_s} \cdot c_{sj}\| \quad (1)$$

for each relation r_s in the source graph where the triple (c_{si}, r_s, c_{sj}) exists.

$$O_{t_{\text{loss}}} = \|M_{r_t} \cdot c_{ti} + r_t - M_{r_t} \cdot c_{tj}\| \quad (2)$$

for each relation r_t in the target graph where the triple (c_{ti}, r_t, c_{tj}) exists.

Simultaneously, we use a transformation $M : E_{G_s} \rightarrow E_{G_t}$ that takes the entities from the seeds we found earlier (Section 2.1) from the source embedding space to the target space, using the following loss:

$$A_{\text{loss}} = \|M \cdot c_s - c_t\| \quad (3)$$

2.4. Inconsistency negatives learning

OWL ontologies are based on Description Logic and facilitate the use of automated reasoners, which in turn facilitate computing entailments of statements from the asserted ontology axioms. In addition, these inferences can be investigated to determine if a class in an ontology is satisfiable or unsatisfiable. A class is unsatisfiable if it cannot have any instances (i.e., the axioms constrain the class in a contradictory way); an ontology is inconsistent if it has at least one instance of a logical contradiction [4, 5]. We utilize the ELK reasoner [6] to find alignments that lead to unsatisfiable classes. In order to find unsatisfiable classes in aligning \mathcal{O}_s and \mathcal{O}_t , we first merge both ontologies (i.e., we combine their axioms into a new ontology) and add all alignments predicted by our model as equivalence class axioms to the merged ontology $\mathcal{O}_{\text{merged}}$. We define this ontology as $\mathcal{O}_{\text{merged}} := (C_s \cup C_t, R_s \cup R_t, I_s \cup I_t, ax_s \cup ax_t, A)$, where C_i is a set of concepts from ontology i , R_i is a set of relations from ontology i , I_i is a set of individuals from ontology i , and ax_i is a set of axioms from ontology i , A is the predicted alignment.

Then we use the ELK reasoner [6] to identify unsatisfiable classes in the merged ontology. If we identify an unsatisfiable class, we generate explanations for the entailment generated by ELK; an explanation consists of a small set of axioms from which the unsatisfiability follows directly; we specifically identify any of the equivalence class axioms we have added within the generated explanations, as these are likely causing the class to become unsatisfiable. We remove the equivalence class axioms causing unsatisfiable classes from the merged ontology and iterate.

Finally, we return to the transformation learning step with an updated loss to optimize for alignment learning as follows:

$$A_{loss} = \|(M \cdot c_s - c_t) - (M \cdot c_{ns} - c_{nt})\| \quad (4)$$

where c_s, c_t are positive class pairs from source ontology and target ontology, respectively, c_{ns}, c_{nt} are pairs of classes which gave rise to unsatisfiable classes and which we removed in the repair step. The new iteration of our method now uses these pairs as negatives during training in the alignment of both ontologies. We repeat this step until no more unsatisfiable classes remain.

3. Results

For this year’s evaluation, we tested A-LION in three tracks: Anatomy, Conference and Material Sciences and Engineering (MSE). We have also tested our system on the phenotypes track using last year’s evaluation tests.

3.1. Participation in OAEI

We selected tracks that align ontologies that contain disjoint class assertion axioms. Disjoint class assertion axioms are a common cause of inconsistencies, and, therefore, we will be able to observe the performance of our method in correcting and training to avoid inconsistencies. In the anatomy track, the ontology file `human.owl` contains 17 disjoint class assertion axioms. In the conference track, the number of disjoint class assertion axioms are as follows: 81 in `cmt.owl`, 42 in `Conference.owl`, 15 in `confious.owl`, 129 in `confOf.owl`, 36 in `crs_dr.owl`, 1,221 in `edas.owl`, 222 in `ekaw.owl`, 3 in `iasted.owl`, 12 in `MICRO.owl`, 384 in `MyReview.owl`, 237 in `OpenConf.owl`, 396 in `paperdyne.owl`. Finally, in the Material Sciences and Engineering track 158 disjoint class axioms was found in `Matonto`. All the results can be found in OAEI 2022 campaign page <http://oaei.ontologymatching.org/2022/results/>.

3.1.1. Anatomy

In terms of precision, recall, and F-measure, the matching performance of A-LION in the anatomy track were below the string equivalence baseline. The main issue that affected our performance in this track is the small number of the predicted alignments and the small number of inconsistencies discovered using the OWL EL reasoner. The main issue affecting the performance of A-LION in the anatomy track is the limited number of initial seeds discovered based on the parameters settings we used, which substantially affected recall. To overcome this limitation, an adaptive method that uses a specific pairs of ontologies to determine parameters (such as for seed matching) could be developed and used to overcome this limitation.

3.1.2. Conference

The Conference track contains information about conference organization. This track comes in two versions: standard and uncertain. The standard version of the Conference track contains a reference alignment which was the result of a “Consensus Workshop” in 2008. However,

some of these alignments may not be possible to detect either by a computational algorithm or manually by humans [7, 8]. For that reason, the uncertain version of the conference track was generated by consulting a group of experts and computing the ratio of agreement on each match. As a consequence, the uncertain track is more realistic because it removes the *controversial* alignments (i.e., the ones for which the experts could not reach a consensus). For that reason, when the evaluation is done on the uncertain version of the track, it is expected that systems increase their performance. A-LION has the highest increase with respect to the standard version among all the systems. This suggests that A-LION is capable to detect *non-controversial* alignments more easily than the *controversial* ones.

The current version of A-LION uses OWL EL reasoning to detect and exclude alignments that cause inconsistencies. However, the results show that A-LION does not detect all inconsistent alignments. The main reason for this lack of removing all incoherent alignment is the use of more expressive description logics than OWL EL. A-LION only uses OWL EL reasoning because computing entailments in expressive description logics has a high computational complexity and may not always be successful for larger ontologies, such as those used in the biomedical domain. However, the ontologies used in the Conference track are small compared to ontologies in other tracks such as Anatomy. In a future version of A-LION, we may include additional reasoners, including reasoners for more expressive logics.

3.1.3. Material Sciences and Engineering

There are three test cases for the Material Sciences and Engineering track. The first and second test cases align MatOnto ontology to the Material Information ontology, and the third case EMMO ontology to Material Information ontology. MatOnto contains 158 disjoint class axioms and could thus introduce useful inconsistencies that can be exploited by our method. The results indicate that A-LION had the highest recall, and an F-measure comparable to the other tested methods. However, there is one test case where A-LION failed to parse the labels in the ontology (the EMMO ontology); consequently, A-LION failed to produce any alignments.

3.2. Phenotype matching use case

We tested the OWL projection method in the problem of aligning phenotype ontologies. To test this approach, we utilized the datasets provided last year [9] for aligning Human phenotype ontology (HP) [10] and Mammalian Phenotype Ontology (MP) [11]. The seed alignments we used are exactly matching IRIs of classes, as well as lexical alignments for HP and MP classes only. We tested two different approaches for generating the graphs from source and target ontologies (Section 2.2). Results are shown in Table 3.2 where we included the results for some of the participating systems from last year for comparison [12, 13, 14, 15]. Comparing the results of the various graph generation techniques, we found that using the OWL projection in the problem of phenotype mappings allows for the discovery of more mappings, whereas the subsumption hierarchy produces alignments with high precision but finds fewer alignments, thereby decreasing the recall.

Table 1

Phenotype use case test results on last year 5-consensus. We show the results for different variations of A-LIO_n, starting with the use of the subsumption hierarchy graph (SH), projected graph (P)

	number of alignments	Precision	Recall	F-score
LogMap	2,136	0.767	0.908	0.831
AML	2,029	0.810	0.910	0.857
ATMatcher	769	0.968	0.412	0.578
TOM	306	0.101	0.140	0.117
A-LIO _n - (SH)	700	0.986	0.382	0.551
A-LIO _n - (P)	1078	0.822	0.732	0.774

4. Conclusion

A-LIO_n is a system that incorporates both entity-level and structure-level information in learning alignments between two ontologies; A-LIO_n also uses logical reasoning to correct alignments that are likely faulty because they lead to unsatisfiable classes, and incorporates the results of this symbolic step in the learning process to generate new negatives. In the future, we plan to make our system able to learn better parameters based on the input ontologies features and self-evaluate the predicted alignment. For example, using a different set of parameters for anatomy and the first task on Material Sciences and Engineering tracks allowed us to increase the F-score by 10% and 3.3% respectively. A further improvement will be the use of language models in seed selection.

References

- [1] J. Chen, P. Hu, E. Jimenez-Ruiz, O. M. Holter, D. Antonyrajah, I. Horrocks, Owl2vec*: Embedding of owl ontologies, *Machine Learning* 110 (2021) 1813–1845.
- [2] M. Chen, Y. Tian, M. Yang, C. Zaniolo, Multilingual knowledge graph embeddings for cross-lingual knowledge alignment, *arXiv preprint arXiv:1611.03954* (2016).
- [3] Y. Lin, Z. Liu, M. Sun, Y. Liu, X. Zhu, Learning entity and relation embeddings for knowledge graph completion, in: *Twenty-ninth AAAI conference on artificial intelligence*, 2015.
- [4] L. T. Slater, G. V. Gkoutos, R. Hoehndorf, Towards semantic interoperability: finding and repairing hidden contradictions in biomedical ontologies, *BMC Medical Informatics and Decision Making* 20 (2020) 1–13.
- [5] J. Martinez-Gil, S. Yin, J. Küng, F. Morvan, Matching large biomedical ontologies using symbolic regression, in: *The 23rd International Conference on Information Integration and Web Intelligence*, 2021, pp. 162–167.
- [6] Y. Kazakov, M. Krötzsch, F. Simančík, Elk: a reasoner for owl el ontologies, *System Description* (2012).
- [7] M. Cheatham, P. Hitzler, Conference v2. 0: An uncertain version of the oaei conference benchmark, in: *International Semantic Web Conference*, Springer, 2014, pp. 33–48.
- [8] J. Bock, C. Dänschel, M. Stumpp, Mappso and mapevo results for oaei 2011, in: *Proceedings*

of the 6th International Conference on Ontology Matching - Volume 814, OM'11, CEUR-WS.org, Aachen, DEU, 2011, p. 179–183.

- [9] M. Pour, A. Algergawy, F. Amardeilh, R. Amini, O. Fallatah, D. Faria, I. Fundulaki, I. Harrow, S. Hertling, P. Hitzler, et al., Results of the ontology alignment evaluation initiative 2021, in: CEUR Workshop Proceedings 2021, volume 3063, CEUR, 2021, pp. 62–108.
- [10] S. Köhler, M. Gargano, N. Matentzoglou, L. C. Carmody, D. Lewis-Smith, N. A. Vasilevsky, D. Danis, G. Balagura, G. Baynam, A. M. Brower, et al., The human phenotype ontology in 2021, *Nucleic acids research* 49 (2021) D1207–D1217.
- [11] C. L. Smith, J. T. Eppig, The mammalian phenotype ontology as a unifying standard for experimental and high-throughput phenotyping data, *Mammalian genome* 23 (2012) 653–668.
- [12] E. Jiménez-Ruiz, Logmap family participation in the oaei 2021, in: CEUR Workshop Proceedings, volume 3063, 2021, pp. 175–177.
- [13] D. Faria, B. Lima1, F. Couto, M. Silva, C. Pesquita, Aml and amlc results for oaei 2021, in: The 23rd International Conference on Information Integration and Web Intelligence, 2021, pp. 131–136.
- [14] S. Hertling, H. Paulheim, Atbox results for oaei 2021, in: CEUR Workshop Proceedings, volume 3063, RWTH Aachen, 2021, pp. 137–143.
- [15] D. Kossack, N. Borg, L. Knorr, J. Portisch, Tom matcher results for oaei 2021, in: CEUR Workshop Proceedings, volume 3063, RWTH, 2022, pp. 193–198.