

# OAEI Machine Learning Dataset for Online Model Generation

Sven Hertling<sup>1</sup>[0000-0003-0333-5888], Ebrahim Norouzi<sup>1</sup>[0000-0002-2691-6995], and Harald Sack<sup>1,2</sup>[0000-0001-7069-9804]

<sup>1</sup> FIZ Karlsruhe – Leibniz Institute for Information Infrastructure,  
Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany  
`firstname.lastname@fiz-karlsruhe.de`

<sup>2</sup> Karlsruhe Institute of Technology (AIFB), Kaiserstr. 89, 76133 Karlsruhe, Germany  
`firstname.lastname@partner.kit.edu`

**Abstract.** Ontology and knowledge graph matching systems are evaluated annually by the Ontology Alignment Evaluation Initiative (OAEI). More and more systems use machine learning-based approaches, including large language models. The training and validation datasets are usually determined by the system developer and often a subset of the reference alignments are used. This sampling is against the OAEI rules and makes a fair comparison impossible. Furthermore, those models are trained offline (a trained and optimized model is packaged into the matcher) and therefore the systems are specifically trained for those tasks. In this paper, we introduce a dataset that contains training, validation, and test sets for most of the OAEI tracks. Thus, online model learning (the systems must adapt to the given input alignment without human intervention) is made possible to enable a fair comparison for ML-based systems. We showcase the usefulness of the dataset by fine-tuning the confidence thresholds of popular systems.

**Keywords:** Ontology Matching · Machine Learning · Online Model Generation.

## 1 Introduction

If applications want to use two or more knowledge graphs (KGs) simultaneously, the corresponding ontologies and instances must be aligned. This process is called ontology alignment (more generally, KG alignment). The inputs are two KGs ( $KG_1$  and  $KG_2$ ) as well as an input alignment  $A_{in}$ . The produced result is an improved alignment  $A_{out}$ . Each alignment consists of (possibly) multiple correspondences in the form  $\langle e_1, e_2, r, c \rangle$  where  $e_1 \in KG_1$  and  $e_2 \in KG_2$ .  $r$  represents the relation between the entities, such as equivalence or subsumption relation (in case of class correspondences), and  $c \in [0, 1]$  is a confidence value given by the matching process.

Starting from 2004, the Ontology Alignment Evaluation Initiative (OAEI) evaluates matching systems each year and allows for a fair comparison between

them. Over the years, more and more systems required correspondences for training and validating their machine learning-based approaches. Especially in 2023, systems for the conference track heavily relied on some form of training data<sup>3</sup> (due to the low number of correspondences in the reference alignment, this makes a huge difference in the final result metrics) but also in other tracks those systems would like to adapt their matching behavior to the task at hand.

In this paper, we introduce train, validation, and test splits for common OAEI datasets and argue that the models need to adapt to the given input alignment during the execution of the systems (online model generation) instead of downloading and generating a model by hand and package it into a matching system (offline model generation). With the latter, the systems will hardly be applicable to new datasets, which is one huge advantage of all matching systems participating at OAEI.

## 2 Related Work

Many systems submitted to OAEI require training data, and in the following, we list those systems together with their setup for generating training examples. In 2023, *GraphMatcher* [1] used the reference alignment in its 5-fold cross-validation and *TOMATO* [11] splitted the created dataset into 60% for training and 40% for testing. *SEBMatcher* [3] in 2022 generated a training dataset by reference and string alignments and sampled 20 % of the reference alignment to create positive cases. *Fine-TOM* [10] (participating in OAEI 2021) used a transformer architecture that needed fine-tuning. Their dataset included 20% of each reference alignment from the Anatomy, Conference, and Knowledge Graph track.

One first direction in creating a machine learning dataset for OAEI is done by He et al. [4]. They provide a train, validation, and test split for a new track called Bio-ML. In their evaluation, they used systems that are trained offline, which means the developers download the training alignments, tune the model, and upload a matcher including that model.

## 3 Generation of the dataset

We use the defined tracks of OAEI and split the reference alignment into 20% for training, 10% for validation, and 70% for testing to align with He et al. [4] and most other systems that generate training data. The reason why the training and validation fraction is so low is the extreme imbalance between correct and incorrect correspondences. Furthermore, in the real-world setting, only a few positive correspondences will be provided as a training signal due to the high effort of creating those correspondences.

To have a good segmentation of the reference alignment, we stratify it by the following criteria: (1) entity type (class, property, instance) of the source and target entity, (2) relation type, and (3) difficulty of the mapping. For the

<sup>3</sup> <https://oaei.ontologymatching.org/2023/results/conference/index.html>

latter, we further differentiate between easy, medium, and hard matches. Easy matches can be found by simple label matching with a bit of string processing like lowercasing and camel case splitting. Medium matches share at least one token among the labels, and the rest are considered hard matches.

The stratification for all criteria is achieved by creating subgroups for each combination, e.g., the group class-class-equivalence-easy and instance-instance-equivalence-medium are two subgroups that we use for stratification, such that in training, validation, and test, we ensure the same distribution as in the original reference alignment. We include the easy matches in the training and validation alignment because, in the real-world use case, the provided matches will also contain those simple correspondences (especially if the mapping is created by finding matching entities for a random sample of entities). The input to a matching system is only one alignment. To differentiate between the training and validation set, we added an additional correspondence extension to indicate to which set it belongs. If a system ignores this distinction, it can use the whole input as training data.

Only positive training examples are provided as input. For a machine learning model, negative examples are often required. Unlike [4], we do not provide negatives (they need it for easily evaluating a ranking-based method) but let the system create them on the fly. This can be achieved by assuming that, for an entity in the source, at most one entity exists in the target graph (or vice versa), and both graphs are duplicate-free. Then, given a correspondence  $\langle A, B \rangle$  in the training alignment, the system can use its own distance function to search for hard negatives of A and B. The advantage is that those correspondences are hard negatives for the matching system at hand and reflects the amount and distribution during the matching of the whole input KGs.

We used only OAEI test cases with at least 350 correspondences in the reference alignment, such that the combined training and validation sets have at least 100 correspondences. This results in the following tracks: Anatomy, BioDiv, Knowledge Graph, and Bio-ML.

All datasets are well integrated with MELT [7] and can be downloaded via the track repository<sup>4</sup> (there is a separated section for it). Code for the generation and evaluation can be found on GitHub<sup>5</sup>.

## 4 Use Case

We analyze the usefulness of the presented dataset by using three matching systems from OAEI 2023, which return correspondences with a wide range of confidence values (not only 1.0). This results in the following systems: Matcha [2], LogMap [8,9], and OLaLa [6,5]. We use the training and validation data to adapt the confidence threshold to the task at hand and remove correspondences below the threshold, which could improve precision and eventually lower recall. Two automated approaches for finding the right threshold are implemented under

<sup>4</sup> <https://dwslab.github.io/melt/track-repository#ml-based-tracks>

<sup>5</sup> <https://github.com/dwslab/melt/tree/master/examples/mlDataset>

**Table 1.** Results of unsupervised (U) and two supervised approaches (SPart, SComp) for tuning the confidence threshold.

	Precision			Recall			F <sub>1</sub> -Measure		
	U	SPart	SComp	U	SPart	SComp	U	SPart	SComp
Anatomy									
Matcha	93.09	93.42	<b>94.72</b>	<b>92.65</b>	92.37	91.23	92.87	92.89	<b>92.94</b>
LogMap	88.34	89.86	<b>96.85</b>	<b>85.01</b>	84.35	69.46	86.65	<b>87.02</b>	80.90
OLaLa	89.04	89.27	<b>95.93</b>	<b>89.63</b>	89.35	82.09	<b>89.34</b>	89.31	88.47
BioDiv									
Matcha	62.30	63.32	<b>63.69</b>	<b>99.12</b>	99.07	99.05	76.51	77.26	<b>77.53</b>
LogMap	61.02	61.86	<b>62.72</b>	<b>99.25</b>	98.30	93.84	75.58	<b>75.93</b>	75.19
OLaLa	63.37	63.69	<b>90.37</b>	<b>99.27</b>	99.18	90.63	77.36	77.57	<b>90.50</b>
KnowledgeGraph									
Matcha	7.01	13.35	<b>33.09</b>	<b>80.64</b>	74.88	56.63	12.90	22.66	<b>41.78</b>
LogMap	45.86	45.90	<b>46.31</b>	<b>72.65</b>	<b>72.65</b>	72.56	56.23	56.26	<b>56.54</b>
Bio-ML									
Matcha	61.37	61.45	<b>63.78</b>	<b>60.29</b>	60.18	57.31	60.80	<b>60.81</b>	60.37
LogMap	58.22	58.46	<b>64.51</b>	<b>58.43</b>	58.27	53.46	58.33	58.36	<b>58.47</b>
OLaLa	36.42	37.31	<b>51.11</b>	<b>42.30</b>	42.04	34.33	39.14	39.53	<b>41.07</b>

the assumption that the input alignment is partial (supervised partial - SPart) or complete (supervised complete - SComp). It computes the optimal threshold by using the training and validation data as a ground truth and finds the value that gives the highest F<sub>1</sub>-Measure.

Table 1 shows the micro-averaged results for all tracks and matchers. The numbers represent the class matches except for the Knowledge Graph track, where the instance matches are shown (OLaLa is not capable of matching instances). The unsupervised case is evaluated on the same test set so that all numbers are comparable.

In most cases, there is only a slight improvement because the systems are highly tuned, and the alignment filtering relies solely on the confidence values. Nevertheless, in some cases, huge improvements can be observed. For example, when setting the right confidence threshold for OLaLa in the BioDiv dataset, there is a huge improvement of over 13% in terms of F<sub>1</sub>-Measure. We hope to see larger improvements if systems use it to train their whole approach using more features to differentiate between correct and incorrect mappings.

## 5 Conclusion

In this paper, we introduced a new dataset for machine learning systems based on existing OAEI tracks. With the presented dataset, we hope to encourage matching system developers to use a given input alignment to tune and optimize their parameters online and only fall back to default/pre-trained parameters if no input alignment is given.

In the future, we plan to create datasets for the conference track that do not have many correspondences. Thus, one complete test case can be used as training data and another one as a test (in-domain transfer learning).

## References

1. Efeoglu, S.: Graphmatcher system presentation. In: *Int. Workshop on Ontology Matching (OM)*. vol. 3591, pp. 154–156. CEUR-WS.org (2023)
2. Faria, D., Silva, M.C., Cotovio, P., Ferraz, L., Balbi, L., Pesquita, C.: Results for matcha and matcha-dl in OAEI 2023. In: *Int. Workshop on Ontology Matching (OM)*. vol. 3591, pp. 164–169. CEUR-WS.org (2023)
3. Gosselin, F., Zouaq, A.: Sebmatcher results for OAEI 2022. In: *Int. Workshop on Ontology Matching (OM)*. vol. 3324, pp. 202–209. CEUR-WS.org (2022)
4. 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: *Int. Semantic Web Conference*. pp. 575–591. Springer (2022)
5. Hertling, S., Paulheim, H.: Olala: Ontology matching with large language models. In: *Knowledge Capture Conference (K-CAP)*. pp. 131–139. ACM (2023). <https://doi.org/10.1145/3587259.3627571>
6. Hertling, S., Paulheim, H.: Olala results for OAEI 2023. In: *Int. Workshop on Ontology Matching (OM)*. vol. 3591, pp. 170–177. CEUR-WS.org (2023)
7. Hertling, S., Portisch, J., Paulheim, H.: Melt-matching evaluation toolkit. In: *Int. Conference on Semantic Systems*. pp. 231–245 (2019)
8. Jiménez-Ruiz, E.: Logmap family participation in the OAEI 2023. In: *Int. Workshop on Ontology Matching (OM)*. vol. 3591, pp. 157–158. CEUR-WS.org (2023)
9. Jiménez-Ruiz, E., Cuenca Grau, B.: Logmap: Logic-based and scalable ontology matching. In: *International Semantic Web Conference (ISWC)*. pp. 273–288 (2011)
10. Knorr, L., Portisch, J.: Fine-tom matcher results for OAEI 2021. In: *Int. Workshop on Ontology Matching (OM)*. vol. 3063, pp. 144–151. CEUR-WS.org (2021)
11. Roussille, P., Teste, O.: TOMATO: results of the 2023 OAEI evaluation campaign. In: *Int. Workshop on Ontology Matching (OM)*. vol. 3591, pp. 191–199. CEUR-WS.org (2023)