

Results of OWL2Vec4OA in the OAEI 2025

Sevinj Teymurova, Ernesto Jimenez-Ruiz and Tillman Weyde

¹City St George's, University of London, EC1V 0HB, UK

Abstract

This paper presents an enhancement of OWL2Vec4OA for ontology alignment, focusing on regression-based local matching. The enhancement integrates Sentence-BERT (SBERT) and Word2Vec embeddings, each combined with lexical and URI-based features. Different fusion strategies—concatenation, averaging, and merging—are explored to create richer embedding representations. The regression model leverages these hybrid embeddings to improve similarity prediction for local ontology matching. Regression models using these hybrid embeddings are evaluated with Hits@K and Mean Reciprocal Rank (MRR) metrics.

Keywords

Ontology alignment, Machine Learning, Knowledge Graph Embeddings

1. Introduction

Ontology alignment is essential for enabling interoperability across heterogeneous knowledge systems. Traditional approaches such as LogMap [1] and AML [2] have proven effective but struggle when the ontologies lack enough lexical information.

Recent years have seen a shift from traditional lexical and structural methods [3] toward machine-learning-based approaches, particularly those leveraging neural networks and large language models. Systems such as LogMap-ML [4], BertMap [5], DeepAlignment [6], VeeAlign [7], SORBETMatcher [8], and OLaLa [9] exemplify this trend, prompting the OAEI to introduce the Bio-ML track [10] to evaluate such systems systematically.

This paper presents a regression-based ontology alignment framework built on OWL2Vec4OA embeddings [11] combining both lexical and URI-based features. OWL2Vec4OA has been extended to integrate both SBERT and Word2Vec. Hybrid embeddings are created through multiple fusion strategies and used to estimate fine-grained similarity scores between entity pairs, improving alignment accuracy under data sparsity and imbalance.

2. Methodology

Our ontology alignment framework enhances OWL2Vec4OA embeddings for local ontology matching. It consists of two main steps: generating hybrid embeddings from SBERT and Word2Vec across lexical, URI, and combined channels, and predicting similarity scores between entity pairs using a regression model. This approach captures both semantic and structural information from the ontologies.

The approach consists of two main components:

1. **Hybrid Embedding Generation** – combining SBERT and Word2Vec representations at lexical, URI, and mixed levels, both individually and in combination.
2. **Regression-Based Similarity Prediction** – computing similarity scores for candidate alignments and evaluating local and global matching.

OM 2025: The 20th International Workshop on Ontology Matching collocated with the 24th International Semantic Web Conference (ISWC 2025), November 2nd, 2025, Nara, Japan

✉ sevinj.teymurova@city.ac.uk (S. Teymurova); Ernesto.Jimenez-Ruiz@citystgeorges.ac.uk (E. Jimenez-Ruiz);

T.E.Weyde@citystgeorges.ac.uk (T. Weyde)

🌐 <https://github.com/Sevinjt/OWL2Vec4OA/> (S. Teymurova)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

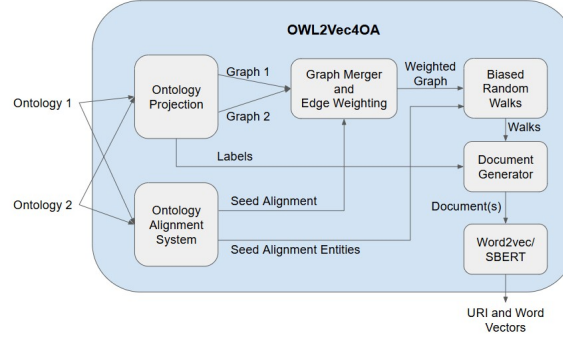


Figure 1: OWL2Vec4OA Architecture: The framework generates ontology embeddings by projecting input ontologies into RDF graphs and aligning them via seed mappings. A merged graph undergoes biased random walks, and the resulting sequences with labels form a corpus to train a Word2Vec model. The output vectors capture structural and textual semantics for ontology alignment.

2.1. OWL2Vec4OA: Input Ontologies and Preprocessing

OWL2Vec4OA (depicted in Figure 1) extends OWL2Vec* [12] by generating embeddings that capture both structural and textual semantics tailored to the ontology alignment task. Input ontologies are converted into RDF graphs, with simple axioms mapped directly and complex axioms transformed appropriately. Seed alignments from LogMap and AML identify high-confidence (“good”) and rejected (“bad”) mappings, which provide both positive and negative training data. The individual graphs are merged into a weighted graph, where edges reflect source confidence. Biased random walks on this graph produce sequences that train a Word2Vec model, resulting in embeddings that serve as features for machine learning models to predict semantic similarity between ontology entities.

2.2. Data Preprocessing

2.2.1. Benchmark Datasets and Alignment Generation

In this study, we utilise multiple benchmark datasets from the Ontology Alignment Evaluation Initiative (OAEI)¹. The bio-ml benchmark consists of biomedical ontologies that are widely used in the biological and medical domains. These datasets provide diverse complexity levels and domain specificity, enabling comprehensive evaluation of our methodology.

2.2.2. Alignment Generation Methodology

We first obtain alignments between the ontologies using two established matching systems: LogMap [1] and AML [2].

From these systems, we generate only four distinct sets of alignments which we use in OWL2Vec4OA:

1. **LogMap overestimation (M_{LO}):** Comprises all correspondences identified by LogMap before applying its mapping repair techniques and logical consistency checks.
2. **Intersection of AML and LogMap mappings (M_I):** Contains only those correspondences that are identified by both AML and LogMap, representing high-confidence alignments that are likely to be correct:

$$M_I = M_{AML} \cap M_{LM} = \{(e_1, e_2, c) \mid (e_1, e_2, c_1) \in M_{AML} \wedge (e_1, e_2, c_2) \in M_{LM}\} \quad (1)$$

where M_{AML} and M_{LM} are the sets of mappings produced by AML and LogMap respectively. In the intersection, we retain the confidence score from LogMap.

¹The Ontology Alignment Evaluation Initiative (OAEI) is a coordinated international initiative to evaluate ontology matching technologies.

3. **LogMap output (M_{LM}):** Contains the final set of correspondences produced by LogMap after applying its repair techniques and logical consistency checks.
4. **Union of AML and LogMap mappings (M_U):** Combines all correspondences identified by either AML or LogMap, providing a more comprehensive but potentially less precise set of alignments:

$$M_U = M_{AML} \cup M_{LM} = \{(e_1, e_2, c) \mid (e_1, e_2, c) \in M_{AML} \vee (e_1, e_2, c) \in M_{LM}\} \quad (2)$$

In case of overlapping mappings, we retain the higher confidence score.

2.2.3. Negative Sample Generation

To train our supervised machine learning models effectively, we generate negative samples from LogMap’s processing pipeline. These include:

- **LogMap hard discards:** Correspondences explicitly rejected by LogMap’s lexical and structural filters.
- **LogMap discards:** Correspondences initially considered but later rejected during the mapping process.
- **Conflicted mappings:** Correspondences that cause logical inconsistencies in the aligned ontology.

The number of negative samples is **increased** to improve the model’s discriminative power. Class imbalance is addressed via **weighted loss functions**, introducing non-determinism and reducing reliance on exhaustive negative sets. The regression loss is **adjusted/lowered** to stabilise training while balancing contributions from positive and negative samples. The training dataset includes approximately 1,000 to 44,000 negative samples, depending on the specific ontology pair.

2.2.4. Embedding Generation using enhanced OWL2Vec4OA

In the second phase of our preprocessing pipeline, we utilise enhanced OWL2Vec4OA to generate vector representations of the ontology entities and their alignments. For each entity, embeddings are generated using SBERT and Word2Vec across three textual channels: lexical labels, URI tokens, and a combination of both. These embeddings are integrated using simple fusion strategies, such as concatenation, averaging or merging, to produce hybrid vectors that capture complementary semantic and structural information. Optionally, dimensionality reduction is applied to improve computational efficiency without sacrificing performance.

For each pair of ontologies, we process:

1. The input ontologies themselves, capturing their hierarchical structure and semantic relationships.
2. The four alignment sets described above (M_{LO} , M_I , M_{LM} , and M_U).

The OWL2Vec4OA framework generates embeddings by:

1. Converting the OWL ontologies into RDF graphs.
2. Applying random walks on the RDF graphs to generate sequences of entities and relationships.
3. Training a Word2Vec model on these sequences to produce vector representations, or fine-tuning a SBERT models on these sequences.
4. Incorporating alignment information to enhance the semantic alignment between embedding spaces of different ontologies.

2.3. OWL2vec4OA Embedding Enhancement Pipeline

2.3.1. Input Ontologies and Preprocessing

Input ontologies O_1 and O_2 are projected into RDF graphs using OWL2Vec4OA’s axiom transformation rules. Seed mappings from **LogMap** and **AML** provide high-confidence positive alignments. Each entity’s URI, label, and annotation are extracted to form textual corpora for embedding generation.

OWL2Vec4OA leverages structured random walks over the merged ontology graphs, generating sequences that capture both lexical and structural semantics. Alignment edges are weighted by source confidence (ontology axioms: 1.0; mappings: confidence value), allowing embeddings to encode semantic and relational context.

2.3.2. Multi-Channel Embedding Generation

For each entity, we generate **three textual representations**:

- **Lexical (lexc)**: tokenized rdfs:label and annotation strings.
- **URI**: tokenized local names extracted from entity IRIs.
- **Mixed (lexc + URI)**: combined lexical and URI tokens.

Each channel is processed using:

- **Word2Vec** - trained on the ontology corpus.
- **SBERT** - a transformer-based contextual embedding model.

We experiment with embeddings **individually** (SBERT-lex, SBERT-URI, Word2Vec-lex, Word2Vec-URI) and in **combinations** (SBERT-mix, Word2Vec-mix, hybrid SBERT+Word2Vec).

2.3.3. Fusion and Dimensionality Reduction

To integrate multiple channels, embeddings are fused using:

$$\mathbf{v}_{concat} = [\mathbf{v}_{lexc}; \mathbf{v}_{URI}; \mathbf{v}_{mix}] \quad (3)$$

$$\mathbf{v}_{avg} = \frac{1}{3}(\mathbf{v}_{lexc} + \mathbf{v}_{URI} + \mathbf{v}_{mix}) \quad (4)$$

This produces hybrid vectors encoding symbolic and contextual semantics. Principal Component Analysis (PCA) is applied to reduce dimensionality (2–96 dimensions), preserving variance while improving computational efficiency.

2.3.4. Regression-Based Similarity Score Prediction

Using the embeddings generated by enhanced OWL2Vec4OA, we train supervised machine learning models to improve LogMap’s output alignments. Our objective is to refine LogMap’s results by identifying false positives and recovering false negatives. Hybrid embeddings serve as input to a Siamese regression model that predicts continuous similarity scores between candidate entities. The architecture is designed to model semantic relatedness efficiently while addressing class imbalance via sample weighting. Training uses early stopping to prevent overfitting and ensures robust similarity estimation. The resulting similarity scores are then used for both local ranking of candidates and global alignment selection.

For training our models, we use:

- **Positive samples**: Alignments from LogMap’s output (M_{LM}), which are assumed to be correct but potentially incomplete.

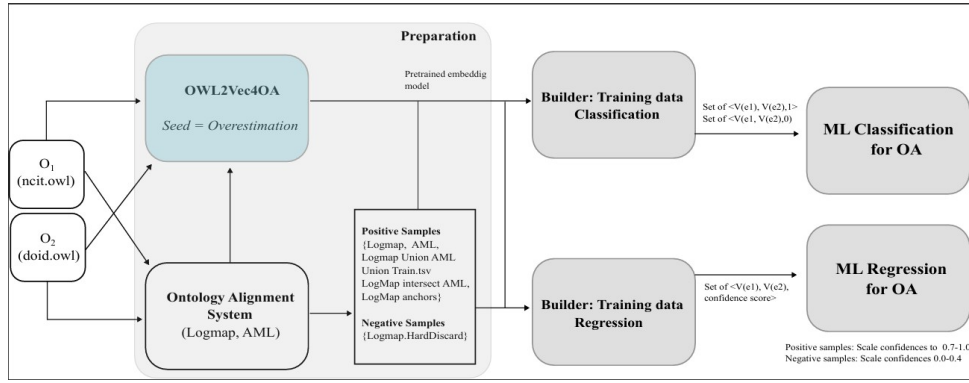


Figure 2: Training OA Model

- **Negative samples:** A combination of LogMap hard discards, LogMap discards, and conflicted mappings, which represent entity pairs that should not be aligned.

Class imbalance is addressed through weighting:

$$w_c = \frac{\sum_{i=1}^C n_i}{C \cdot n_c} \quad (5)$$

where w_c is the weight for class c , C is the number of classes, and n_c is the sample count for class c .

We employ established pre-trained models as the foundation for our machine learning approach, fine-tuning them on our alignment task. This transfer learning approach enables us to leverage general semantic understanding while adapting to the specific requirements of ontology alignment and predicts the scores between alignments (regression model)

In this research we use ML-based method to predict new alignments between two ontologies (classification model) and predict the scores between alignments (regression model). We employ two Siamese Neural Network architectures:

2.4. Regression Model Architecture

The regression model predicts continuous similarity scores between candidate ontology entities.

Architecture Details:

To model semantic similarity between biomedical concepts, we leveraged a pre-trained model (OWL2Vec4OA*) to convert ontology terms into dense vector embeddings. Positive pairs were scaled into the similarity interval $[0.6, 1.0]$, while negative pairs were scaled into $[0.0, 0.6]$ using min-max normalization. Entities lacking pre-trained embeddings were discarded to maintain consistency in feature representation.

- Input: 256-dimensional embeddings (from fused SBERT/Word2Vec channels)
- Encoder: three Conv1D layers with tanh activation, batch normalisation, and max-pooling
- Output: a single regression head producing similarity in $[0, 1]$
- Loss: **mean absolute error (MAE)** between predicted and true similarity
- Optimisation: **Lion optimiser**, learning rate 0.001, batch size 1024, early stopping (patience = 5)
- Sample weights derived from class frequency to address imbalance
- PCA-compressed embeddings (2–96 dimensions) are used for robustness testing

2.5. Similarity Computation and Local Matching

Entity similarity is measured using standard vector-based metrics, which are combined to rank candidate target entities for each source entity. Alignments are determined by applying a similarity threshold




Figure 3: Selection of mappings for the Final alignments

to regression predictions, retaining only high-confidence pairs. This approach allows the system to prioritize precise matches while maintaining overall coverage.

Similarity between ontology entities is computed using multiple metrics implemented in our code:

$$\text{sim}_{\cos}(a, b) = 1 - \cos(\mathbf{v}_a, \mathbf{v}_b) \quad (6)$$

$$\text{sim}_{\text{euclid}}(a, b) = \frac{1}{1 + \|\mathbf{v}_a - \mathbf{v}_b\|_2} \quad (7)$$

$$\text{sim}_{\text{man}}(a, b) = \frac{1}{1 + \|\mathbf{v}_a - \mathbf{v}_b\|_1} \quad (8)$$

These similarity scores are used for:

- **Local Matching:** evaluating the ranking of candidate target entities for each source entity.
- **Global Matching:** selecting high-confidence alignments across all entity pairs.

Evaluation Metrics:

- **MRR:** measures the rank of the correct entity for each source.
- **Hits@K:** percentage of queries where the correct entity is within the top K candidates.

2.6. Final Alignment Selection

Final alignments are determined by applying a similarity threshold θ on regression predictions:

$$A = \{(e_1, e_2) \mid f_{\text{reg}}(e_1, e_2) \geq \theta\} \quad (9)$$

The study retained only entity pairs surpassing a defined similarity threshold in the final ontology alignment. Using two Siamese neural network architectures with OWL2Vec4OA embeddings—which integrate textual and structural ontology information—the model effectively captured rich semantic relationships between concepts. Training ran for up to 250 epochs with early stopping and balanced sampling, preserving the best-performing checkpoints. The framework used dual Conv1D towers and a similarity regression branch to estimate concept relatedness, keeping pairs with similarity scores above 0.60. All models, configurations, and results were fully documented and exported for reproducibility.

The regression model utilises an identical Siamese architecture with an additive attention mechanism for enhancing semantic similarity representation:

$$\text{Attention}(\mathbf{q}, \mathbf{k}, \mathbf{v}) = \sum_i \frac{\exp(\text{score}(\mathbf{q}, \mathbf{k}_i))}{\sum_j \exp(\text{score}(\mathbf{q}, \mathbf{k}_j))} \mathbf{v}_i \quad (10)$$

3. Results

3.1. Ontology Alignment Summary

The final alignment combines classification and regression confidence:

$$A = \{(e_1, e_2) \mid f_{class}(e_1, e_2) = 1 \wedge f_{reg}(e_1, e_2) \geq \theta\}$$

where θ is a confidence threshold.

3.2. OWL2vec4OA Enhancement

This section presents the consolidated evaluation results of the enhanced **OWL2Vec4OA** and retrieval experiments conducted across four datasets: **NCIT-DOID**, **OMIM-ORDO**, **SNOMED-Neoplas(NCIT)**, and **SNOMED-Neoplas (Pharm)**. Each experiment compared two embedding strategies—**Sentence-BERT (SBERT)** and **Word2Vec**—while employing multiple similarity metrics (Cosine, Euclidean, Manhattan) and varying embedding dimensions depending on the model configuration. Although similarity metrics and embedding dimensions were part of the evaluation pipeline, they are omitted from the table below for clarity and are discussed only within this descriptive text.

Across all datasets, the results consistently demonstrate that **SBERT** outperforms **Word2Vec** in terms of Mean Reciprocal Rank (MRR) and Hits@ k (for $k = 1, 5, 10$). The highest scores were achieved by SBERT in combination with cosine similarity, highlighting that contextualised sentence-level embeddings capture richer semantics than static word vectors.

Table 1

Consolidated comparison of SBERT and Word2Vec across all datasets. Best results per dataset are highlighted in bold.

| Dataset | Metric | SBERT | Word2Vec | Improvement (%) |
|------------------------|---------|--------------|----------|-----------------|
| NCIT-DOID | MRR | 0.879 | 0.731 | +20.2 |
| | Hits@1 | 84.15 | 57.84 | +45.5 |
| | Hits@5 | 92.53 | 94.12 | -1.7 |
| | Hits@10 | 95.27 | 98.14 | -2.9 |
| OMIM-ORDO | MRR | 0.707 | 0.594 | +19.0 |
| | Hits@1 | 63.45 | 52.48 | +20.9 |
| | Hits@5 | 78.62 | 67.45 | +16.5 |
| | Hits@10 | 83.80 | 70.06 | +19.6 |
| SNOMED-Neoplas(NCIT) | MRR | 0.884 | 0.414 | +113.6 |
| | Hits@1 | 82.18 | 27.97 | +193.7 |
| | Hits@5 | 95.45 | 55.12 | +73.2 |
| | Hits@10 | 96.68 | 64.77 | +49.3 |
| SNOMED-Neoplas (Pharm) | MRR | 0.831 | 0.431 | +92.8 |
| | Hits@1 | 76.92 | 35.37 | +117.4 |
| | Hits@5 | 89.86 | 49.97 | +79.8 |
| | Hits@10 | 91.95 | 54.69 | +68.1 |

Overall, The results show that embedding type and dimensionality greatly influence ontology alignment accuracy, with particularly strong gains in the SNOMED-Neoplas (NCIT). The largest improvements are observed in Hits@1 and MRR, indicating that SBERT is substantially more effective at ranking correct matches at the top of the retrieval list. While Word2Vec occasionally achieves slightly higher scores at broader cutoffs (Hits@5 or Hits@10), SBERT remains the dominant model, demonstrating superior semantic representation quality for ontology alignment tasks.

3.3. Regression-Based Similarity Prediction Results

The regression-based similarity prediction model was re-evaluated using the expanded BIO-ML datasets, incorporating full query coverage across all four ontology mapping tasks: **NCIT-DOID**, **OMIM-ORDO**, **SNOMED-Neoplas (NCIT)**, and **SNOMED-Neoplas (Pharm)**. The enlarged datasets improved overall robustness, yielding consistent gains in ranking quality.

Table 2 summarizes the best-performing configurations. The **NCIT-DOID** task produced the strongest results, achieving an **MRR of 0.8795** and **Hits@1 of 84.2%**, slightly improving over the previous maximum. **OMIM-ORDO** also benefited from the extended dataset, reaching **MRR 0.7079** and **Hits@1 63.6%**, with notable improvements at higher retrieval depths.

For **SNOMED-Neoplas (NCIT)**, the model attained an **MRR of 0.8286** and **Hits@1 of 77.2%**, reflecting stronger neoplasm alignment performance. The **SNOMED-Neoplas (Pharm)** configuration achieved the highest pharmacological alignment accuracy, with **MRR 0.8647** and **Hits@1 79.7%**, and strong precision across broader cutoffs.

Overall, these results demonstrate steady and consistent improvements across all tasks, confirming the model’s enhanced generalisation and ability to capture fine-grained biomedical semantic relationships.

Table 2

Updated regression-based performance across ontology mapping tasks.

| Dataset | MRR | Hits@1 | Hits@5 | Hits@10 | Hits@20 | Hits@30 |
|-------------------------------|--------|--------|--------|---------|---------|---------|
| NCIT-DOID | 0.8795 | 84.2 | 92.4 | 95.3 | 97.0 | 97.8 |
| OMIM-ORDO | 0.7079 | 63.6 | 78.6 | 83.8 | 88.5 | 91.1 |
| SNOMED-Neoplas (NCIT) | 0.8286 | 77.2 | 89.9 | 92.4 | 95.3 | 96.9 |
| SNOMED-Neoplas (Pharm) | 0.8647 | 79.7 | 94.2 | 95.9 | 97.1 | 97.7 |

4. Conclusion

Combining SBERT and Word2Vec embeddings with regression-based similarity scoring continues to yield strong results for ontology alignment. The updated experiments demonstrate consistent improvements across all BIO-ML tasks, with notable gains in MRR and Hits@1 for NCIT-DOID, OMIM-ORDO, and SNOMED-Neoplas (NCIT) mappings. These results confirm the effectiveness of embedding-based regression in capturing semantic correspondences across heterogeneous biomedical ontologies.

SBERT remains particularly strong in producing high top-ranked precision, while Word2Vec contributes additional contextual diversity, enhancing overall coverage. The approach scales efficiently to larger datasets, maintaining high retrieval accuracy even when the number of candidate mappings increases substantially.

The improved recall across all mappings indicates that the Graph-AI-based alignment system effectively reduces the candidate space to a small set of high-probability matches—an essential property for semi-automatic ontology integration workflows.

Overall, these enhanced results validate the robustness and scalability of the proposed regression-based approach. Future work will focus on integrating these regression-based representations into the broader Graph-AI alignment pipeline to further improve contextual reasoning and cross-domain generalisation.

Declaration on Generative AI

During the preparation of this work, the authors used Grammarly in order to grammar and spell check, and improve the text readability. After using the tool, the authors reviewed and edited the content as needed to take full responsibility for the publication’s content.

References

- [1] E. Jimenez-Ruiz, B. Cuenca Grau, LogMap: Logic- Based and Scalable Ontology Matching, The Semantic Web – ISWC (2011) vol 7031.
- [2] D. Faria, C. Pesquita, E. Santos, M. Palmonari, I. F. Cruz, F. M. Couto, The agreementmakerlight ontology matching system, in: On the Move to Meaningful Internet Systems, volume 8185 of *Lecture Notes in Computer Science*, Springer, 2013, pp. 527–541. URL: https://doi.org/10.1007/978-3-642-41030-7_38. doi:10.1007/978-3-642-41030-7_38.
- [3] J. Euzenat, P. Shvaiko, Ontology Matching, 2nd ed., Springer, 2013. URL: <https://doi.org/10.1007/978-3-642-38721-0>.
- [4] J. Chen, E. Jiménez-Ruiz, I. Horrocks, D. Antonyrajah, A. Hadian, J. Lee, Augmenting ontology alignment by semantic embedding and distant supervision, in: The Semantic Web: ESWC, Springer, 2021, pp. 392–408.
- [5] Y. He, J. Chen, D. Antonyrajah, I. Horrocks, BERTMap: A BERT-Based Ontology Alignment System, in: Thirty-Sixth AAAI Conference on Artificial Intelligence, 2022.
- [6] P. Kolyvakis, A. Kalousis, D. Kiritsis, 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.
- [7] V. Iyer, A. Agarwal, H. Kumar, VeeAlign: Multifaceted Context Representation Using Dual Attention for Ontology Alignment, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP), Association for Computational Linguistics, 2021, pp. 10780–10792. URL: <https://doi.org/10.18653/v1/2021.emnlp-main.842>. doi:10.18653/v1/2021.EMNLP-MAIN.842.
- [8] F. Gosselin, A. Zouaq, SORBET: A Siamese Network for Ontology Embeddings Using a Distance-Based Regression Loss and BERT, in: International Semantic Web Conference, Springer, 2023, pp. 561–578.
- [9] S. Hertling, H. Paulheim, OLaLa: Ontology Matching with Large Language Models, in: K. B. Venable, D. Garijo, B. Jalaian (Eds.), Proceedings of the 12th Knowledge Capture Conference (K-CAP), ACM, 2023, pp. 131–139. URL: <https://doi.org/10.1145/3587259.3627571>. doi:10.1145/3587259.3627571.
- [10] Y. He, J. Chen, H. Dong, E. Jiménez-Ruiz, A. Hadian, I. Horrocks, Machine Learning-Friendly Biomedical Datasets for Equivalence and Subsumption Ontology Matching, in: 21st International Semantic Web Conference, volume 13489 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 575–591. URL: https://doi.org/10.1007/978-3-031-19433-7_33. doi:10.1007/978-3-031-19433-7_33.
- [11] S. Teymurova, E. Jiménez-Ruiz, T. Weyde, J. Chen, OWL2Vec4OA: Tailoring Knowledge Graph Embeddings for Ontology Alignment 15459 (2024) 168–182. URL: https://doi.org/10.1007/978-3-031-81221-7_12. doi:10.1007/978-3-031-81221-7_12.
- [12] 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.