

An Eye on Representation Learning in Ontology Matching

Guilherme Sousa¹, Rinaldo Lima² and Cassia Trojahn¹

¹*Institut de Recherche en Informatique de Toulouse, Toulouse, France*

²*Universidade Rural de Pernambuco, Recife, Brazil*

Abstract

Representation learning has received increased attention in the last few years in several tasks, including knowledge graph completion, entity resolution, and ontology matching. This paper presents an overview of representation learning approaches applied to the ontology matching task. It proposes to classify such approaches into the following dimensions: lexical unit segmentation, training strategy, and information representation complexity. A discussion on them is presented together with their pros and cons. Perspectives for further developments are also discussed.

Keywords

Ontology Matching, Representation Learning, Embeddings

1. Introduction

The advances in machine learning have promoted the emergence of new architectures and methods capable of learning complex patterns on different types of data. Some applications show impressive results on NLP tasks, such as question answering and summarisation [1], image generation [2], or still in general tasks such as game playing and image captioning [3]. In knowledge representation, representation learning methods show significant results in tasks such as graph completion [4] and link prediction [5]. In ontology matching, a wave of representation learning systems has appeared in the last few years. Word embeddings were one of the first representation learning strategies adopted [6]. These models are based on the distributional hypotheses stating that similar words appear in similar contexts. The well-known Word2Vec [7] model has been used to compute the semantic similarity of ontology entities that can improve systems performance compared to classical lexical similarity methods [8].


Since word embeddings do not consider the ontology structure, better encoding strategies were designed. The RDF2Vec[9] is used to integrate background knowledge in ontology matching systems such as ALOD2Vec [10] and [11]. It generates sentences by randomly walking the paths in ontology and uses the Word2Vec algorithm to generate entity embeddings. Besides its representation capabilities, RDF2Vec does not fully represent OWL constraints, without considering as well the word composition in labels, harming the model's generality. To address these problems, the OWL2Vec model [12] was designed by using a set of rules to map OWL

Woodstock'22: Symposium on the irreproducible science, June 07–11, 2022, Woodstock, NY

✉ guilherme.santos-sousa@irit.fr (G. Sousa); rinaldo.jose@ufrpe.br (R. Lima); cassia.trojahn@irit.fr (C. Trojahn)

🆔 0000-0002-2896-2362 (G. Sousa); 0000-0002-1388-4824 (R. Lima); 0000-0003-2840-005X (C. Trojahn)

© 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)

predicates to RDF equivalents and considering the word composition of entity labels. Since these models generate independent embeddings for entities in two different ontologies, post-processing strategies like the alignment of embeddings using linear transformations [11] or Siamese networks [13] were used. Besides the performance of random walk-based strategies, as adopted in RDF2Vec and OWL2Vec, Graph Neural Networks [14] are more flexible and show impressive performance in generating graph embeddings [15, 16]. The Graph Attention Network [17], for instance, can filter irrelevant neighbors in entity encoding leading to more robust similarity metrics.

For better representation of terminological layers, Transformers [18] and BERT [19] language models were introduced in ontology matching systems [20, 21, 22]. They have achieved high performance and capacity of learning language representations in NLP tasks [23]. They have attracted interest in generating embeddings for entity labels. BERT is a language model based on the Transformer architecture [18] and has outperformed Recurrent Neural Networks as they rely on an attention mechanism to filter irrelevant data. BERT enables the architecture development of huge models with billions of parameters [1]. It is pre-trained by masking some input tokens and by predicting the next sentence. This task provides the resulting model with higher context notion capabilities in different NLP tasks by fine-tuning the model for each task. In ontology matching, transformer-based models are used to learn representations for the ontology terminological layer (labels, comments, etc).

The goal of this paper is to review ontology matching systems based on representation learning. These systems have some properties in common that enable them to be categorized based on the following dimensions: lexical unit segmentation, training, and information complexity. A discussion on the performance of variant systems to ontology matching is presented together with their pros and cons. The perspectives for further developments are also discussed. We start by presenting a categorization framework (§2), followed by the description of the reviewed proposals (§3). We then discuss their impact on ontology matching (§4) together with an analysis for further work (§5).

2. Categorisation framework

The works in the literature are categorized into three groups of common features: lexical unit segmentation, training, and information complexity. In the next sections, these groups are discussed in detail. The reviewed works are listed in Table 1.

2.1. Lexical unit segmentation

The lexical unit segmentation categorizes the systems in the way they represent the terminological layer of the input ontologies: entity (no tokenization), words, and character. An example of the categories is illustrated in Figure 1. **Entity level** refers to the systems where no tokenization is applied in the terminological layer (id, labels, comments, etc), viewing these terminological features as a unique symbol (ex. "New_York_City"). It acts as an identifier and has low generalization since ontologies may combine the same words in a different order to represent the same entity. **Word level** refers to the splitting of the terminological features in their word components (ex. ["New", "York", "City"]). In order to generate embeddings with

Work	Lexical Unit			Learning			Information level			
	Entity	Word	Char	Ref.	Trained	Pre-trained	No Context	Context	Graph	BK
Zhang [6]		X				X	X			
Xiang [24]		X			X			X		
Kolyvakis [25]		X			X	X				X
Kolyvakis [26]		X			X	X				X
Nkisi [27]		X		X	X	X		X		
Gromann [28]		X				X	X			
Jimenez [29]		X			X		X			
OntoEmma [30]		X	X	X	X	X				X
Li [31]		X			X			X		X
HISDOM [32]		X				X		X		
Tounsi [33]		X				X		X		
DOME [34]		X			X			X		
Li [8]		X				X	X			
DeepFCA [35]		X	X		X	X		X		X
Bento [36]			X	X	X	X		X		
ALOD2Vec [10]					X					X
LogMap* [13]	X	X			X				X	
DAEOM [16]		X		X	X	X			X	
BERTMap [20]		X			X	X				X
OntoConnect [15]			X		X	X			X	
VeeAlign [37]		X		X	X	X			X	
AMD [38]	X				X				X	
Fine-TOM [21]		X		X	X	X	X			

Table 1

The reviewed works are categorized according to Lexical Unit Segmentation, Training, and Information level, sorted by publication year.

the same dimension, an aggregation step needs to be applied like summing or averaging the embeddings of each word. Most works in ontology matching use this approach since many pre-trained word embeddings are available and the input ontologies in the same domain share overlapping vocabulary. **Character level** is more generalist as it can be applied to any domain with the same alphabet. This type of approach requires a complex model to generate relevant representations as word boundaries need to be learned by the combination of characters.

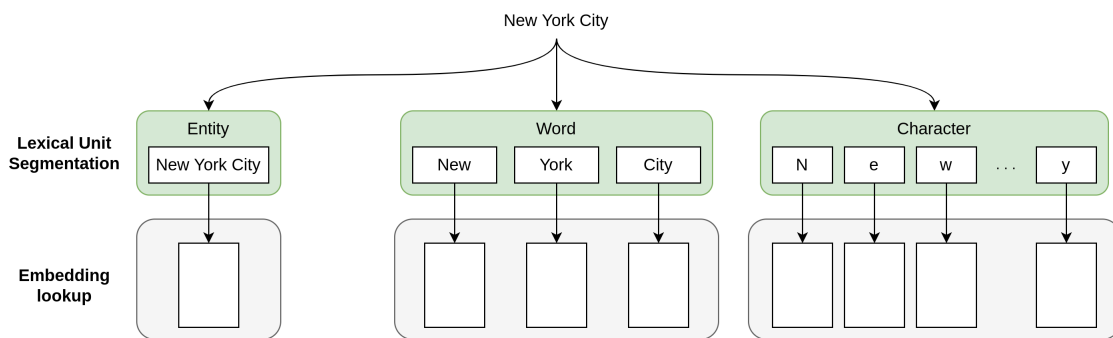


Figure 1: An example of the label "New York City" segmented in the Entity, Word, and Character levels. Optional lower casing can be applied.

2.2. Learning

Some works are supervised and need labeled data for training while others are based on pre-trained models or still on unsupervised learning. **Reference** refers to those systems that need reference alignments for training the model. The **Training** category classifies those systems that need to fine-tune or learn weights for each ontology before the production of final alignments. **Pre-trained** refers to the systems that contain background knowledge in the form of pre-trained word embeddings or models.

2.3. Information complexity

This group describes how much information is considered in the matching process. The **No Context** category describes the systems that do not consider the entity neighborhood. **Context** characterizes the systems that consider neighborhood information without considering the graph structure (e.g., predicates or graph edge directions). The **Graph** category groups the architectures that consider the neighbors and edge features of the relations that connect them in arbitrary depth. The **Background Knowledge** category groups the systems that can aggregate information that is not present in the ontology (e.g., external dictionaries or thesauruses).

3. Works on the Information Complexity Category

We chose to describe the works that are grouped according to the category *information level*, which is the most discriminant one. This category classifies the works according to four main features (Table 1): i) considering element-level based strategies (no context); ii) systems exploiting context; iii) systems using deeper context (graphs); and iv) systems using external background knowledge.

3.1. Element-level based systems (no context)

This category groups the works that map an entity's terminological layer to embeddings without considering its neighborhood. The work of [6] is claimed to be the first work in ontology matching that used word embeddings, which were learned on Wikipedia data using two techniques: Word2Vec [39] and Latent Semantic Analysis [40]. Another relevant work using word embeddings is presented in [8]. The entity label is tokenized and each token is mapped to a word embedding from a biomedical domain corpus. Despite its simplicity, this method improved lexical similarities like edit distance. While these works have focused on monolingual ontologies, the system in [28] uses pre-trained word embeddings to align multilingual ontologies. A cross similarity between each word embedding, corresponding to the compared entity labels, is used to build an intermediate vector to produce the final similarity. To deal with the *Out Of Vocabulary* (OOV) problem, this work uses the average of the similarity of the present words.

While few works expose their implementation as modules or libraries, the MELT platform [41] has support for the use of pre-trained transformers in a pipeline for ontology matching. FineTOM [21, 22] is a system based on MELT that uses transformers and two pipelines. The first step is a training pipeline to fine-tune a pre-trained transformer and the second is a matching

pipeline that uses the transformer to measure the similarity between entities. Despite the capacity of transformer-based models to generate contextualized word embeddings, the TOM and Fine TOM were classified in the No Context category as these systems don't consider information from the entity neighborhood in the entity embedding.

Finally, with a focus on large ontologies, the work in [29] proposes a technique to cluster entities using an inverse word index to reduce the required matching space. Two clustering strategies are applied to generate the subtasks based on the inverse word index. The first is the random splitting of the entities in clusters of the same size and the second is the use of the StarSpace [42] embedding model to cluster related entities while learning embeddings for individual words.

3.2. Systems exploiting context

This category refers to the works that include information about the surroundings of each entity, without exploring the full graph structure. The work of [27] uses a combination of automatically learned and manually engineered features containing word embedding similarities. Four similarity metrics are proposed, combining context and entity labels similarity that are used to build a feature vector. The final alignment is produced using a random forest classifier to predict the alignment using the generated features. In [33], word embeddings are used to represent the entities as the average of the embeddings related to entity labels words and its context as the average of the embeddings in the entity lineage. The final alignment is generated by comparing the cosine similarity between the entity and context embeddings on the source and target ontologies. The system can infer the specific relation between the alignment using the radius measure of their contexts. Similar work, but using character level, is proposed in [36]. The entity embedding is generated using a CNN on the characters of the entity labels, its parents, and children. The character embedding is pre-trained and the last layer of the CNN is connected to a single neuron with sigmoid activation to predict the alignment probability.

The architecture ERSOM (Entity Representation and Structure-based Ontology Matching) [24] uses a stacked auto-encoder to learn features from entities based on different methods to encode classes, properties, and instances. Each entity representation aggregates context information based on its surroundings and hidden representations are learned using the auto-encoder. In the same sense, HISDOM [32] uses different similarity metrics for instances, names, attributes, structure, and comments. The name similarity is a weighted sum between edit distance and embedding cosine similarity in the entity labels. In particular, the structural similarity is calculated as the weighted sum of the Jaccard similarity of the children and parent concepts, with the comment similarity being calculated embedding the sentence of the entity comments using a CNN. The final similarity is the weighted sum of all similarities. Finally, DOME [34] applies a pipeline of filters. The first one is a string similarity matcher that compares the textual description of entities. The next step is a confidence adjustment using an embedding generated by Doc2vec [43], a generalization from Word2vec, to calculate cosine similarity. After instance and property alignment were performed, the class alignment is revisited using the matched instances as context to the class similarity since classes of matched individuals tend to be the same.

3.3. Graph as context

Graph-based category refers to the systems representing the ontology graph structure in depth. The OWL2Vec [12] model is similar to RDF2Vec[9], however, it differs in the lexical unit segmentation as OWL2Vec combines entity and word representation while RDF2vec only considers entity lexical units. LogMap* system [13] uses this embedding model, and since the generated embeddings are independent, a siamese network is used to learn a transformation that projects the embeddings into the same space.

OntoConnect [15] uses a recursive neural network to encode the graph structure of the entities in an unsupervised manner. The entity features are generated using FastText [44], an embedding model developed by Facebook based on character n-grams. A recursive neural network based on LSTM [45] is used to embed the graph structure of each entity. Also using graph neural networks DAEOM [16] (Deep Attentional Embedded Ontology Matching) divides its architecture into Ontology Attentional Encoder and Mapping Selection. The first step of the Encoder is the embedding of the terminological descriptions. A fine-tuned architecture based on BERT [19] is used to embed the words of textual data in the entities. Next, a Graph Attention Network (GAT) is used to aggregate the graph structure of the entity with its terminological layer embedding. Finally, VeeAlign [37] is a system that builds entity embeddings using the Universal Sentence Encoder [46] and aggregates four views to generate the context embedding: parents, child, properties, and datatype properties. A path and node attention is applied to select the most important representations. The entity embedding is then concatenated to the context embedding and passed through a feed-forward layer for down-sampling.

An example of the use of translational embedding to ontology matching [5] is AMD (AgreementMakerDeep) [38]. This type of embeddings was shown to be able to model some logical relations like inversions, symmetries, and some compositions that appear in ontologies entities (ex. Inverse Properties). It uses a modified version of RotatE [47] to generate embeddings of the ontologies and compare similarities.

3.4. Background knowledge

This category is dedicated to the works that use external knowledge in the matching process. Most systems use pre-trained embeddings and models as background knowledge as Alod2Vec [10] which uses embeddings encoded with RDF2Vec [9] on an external dataset. Other systems like DeepAlignment [25] refine word embeddings using synonyms to increase their semantic similarity. Pre-trained word embeddings are refined by contrasting synonym samples from background knowledge and negative samples from entities that are not explicitly stated as equivalent. BERTMap [20, 48] system also extract synonyms from background ontologies to fine-tune a BERT architecture. It also includes a mapping repair to increase the quality of final alignments. A different strategy is adopted in [26], where a Siamese CBOW [49] refines the word embeddings by extracting paraphrases from background knowledge and a Denoising Auto Encoder (DAE) [50] to learn entity embeddings.

In DeepFCA [35] word embeddings and formal concept analysis [51] are adopted. A pre-trained word embedding is used to map the words in the terminological layer using a character-level embedding as a fallback if the word is not present. The entity representation is created

by averaging the word vectors refined by synonyms extracted from semantic lexicons and contrasting them with negative samples from the ontology, guided by a lattice generated from formal concept analysis.

The system in [31] uses multi-view embedding, a strategy that combines embeddings from different views as parent and children ontology concepts, to calculate similarities between entities. Embeddings are learned using a negative sampling strategy to contrast entity synonyms extracted from external ontology with random samplings. Finally, OntoEmma [30], as in the previous work, combines representations from different views to generate entity embeddings. The first is the name view that concatenates character-level embeddings with pre-trained word embeddings and is aggregated using a bidirectional LSTM. The aliases, definition, and context views use the same strategy as the name view without the use of character embeddings and are enriched with background knowledge. The final alignment similarity is calculated by a neural network to predict the alignment probability.

4. Discussion

As introduced above, many representation learning strategies have been applied to ontology matching. Most of the works shown in Table 1 are based on word lexical segmentation and trained per ontology. They use pre-trained word embeddings and, some of them, rely on pre-trained models (e.g BERT), which take into account some types of contextual information. A qualitative analysis of the selected OM systems shows that not only a good capacity for generalization, but also a higher level of flexibility can be achieved when representation learning methods were applied. In the remainder of this section, we point out common difficulties found when applying representation learning methods in ontology matching. In addition, some improvements to address such difficulties are also discussed.

Lexical Unit Segmentation The use of representation learning in ontology matching has some common problems that need to be addressed. First, the non-determinism of neural learning-based methods leads to encoding entities in two different ontologies, that do not share vocabulary, producing meaningless similarity scores when compared [13]. Some works try to mitigate this issue by projecting the embeddings of the mapped ontologies into the same space using linear transformations [11] or learning a projection matrix using siamese neural networks [13]. The choice of entity lexical unit segmentation, as described in Section 2.1, emphasizes this problem since similar entity labels with the same words, but in a different order, may lead to distinct representations with a low similarity between them. Another attempt to solve this problem is the use of word and character lexical unit segmentation [52, 53]. Moreover, word-level lexical units are also prone to the known out-of-vocabulary (OOV) issue, when words not present in the training vocabulary may decrease embedding representation robustness. Thus the use of character embeddings to generate entity representations may be a promising research direction to mitigate the OOV issue when combined with other types of learning representations. Second, the performance of word embeddings in ontology matching is limited by the vocabulary coverage of the employed embeddings restricting matching systems developers to use domain-specific word embeddings. Since the alphabet size of some languages is orders of magnitude

lower than the vocabulary size and even different languages may share some part of their alphabet, the use of characters embedding seems to reduce the occurrence of OOV at the cost of bigger models [54].

Learning Concerning the choice of learning strategy, one difficulty of applying supervised representation learning techniques is the low amount of reference alignments present in OM benchmarking datasets [25], making unsupervised and self-supervised learning valuable techniques to model ontology concepts. However, the major challenge when employing such learning strategies consists in modeling a direct loss metric that reflects the similarity between two given entities. In fact, the general notion of similarity can change among ontology domains and is in the focus of ongoing research [55]. Concerning the training model dimension, recent works have demonstrated that relying only on word embeddings is not enough to fully address the matching problem of rejecting false positives so the contextual information derived from graph structures in ontologies needs to be taken carefully into consideration [16]. Actually, some graph embedding techniques achieve promising results using models supervised learning techniques [37, 16]. However, due to the problem of a low amount of labeled data, such embedding techniques may not achieve their highest performance in ontology matching without some adaptation [56]. One possible direction to deal with the low amount of labeled data is based on unsupervised graph embedding methods, including translational models [5], such as AMD [38]. On the other hand, unsupervised graph embedding methods have shown effective performance in encoding graph structures, especially for the Link Prediction task in Knowledge Graphs [4]. These models have the advantage of being able to learn logical relations between entities, including inversion, symmetrical, and some types of relation compositions obeying the locality principle [57] in ontology matching. However, only relying on this type of embedding method does not ensure acceptable performance, since it does not enforce high similarity scores between similar entities belonging to two distinct ontologies.

Information Level The use of embeddings brings the possibility to compare symbolic data similarities, by analogy, using metrics such as euclidean distance or cosine similarity. For designing metrics fully adapted to the ontology matching problem, the embeddings should be organized in a latent space enforcing that similar concepts have representations positioned near to each other in the same vector space. However, some models, e.g vanilla Auto Encoders do not enforce this type of organization in latent space. As a result, the learned embeddings of similar entities are not guaranteed. Some improvements in generative models such as Infogan [58], Variational Auto Encoder (VAE) [59] and Beta-VAE [60] can learn expressive interpretable representations and can encode background knowledge information in their models. They encode specific features in each embedding dimension making similar embeddings be placed near each other in latent space. Therefore, in ontology matching, these different dimensions could offer better generalization and interpretability of concepts. An interesting research direction is the use of these representation methods for encoding ontology entities, allowing the embedding representations to have similarity metrics suitable to the matching task. Another direction is to exploit such existing representation learning techniques aiming to generate expressive alignments. While many systems rely on context and the ontology graph structure

itself, in complex alignment, is difficult to define the boundaries of the entities as they can be composed of multiple distinct elements. This may cause the need either change existing architectures or propose new ones to deal with the complex alignment task.

5. Conclusion and Future Work

This paper presented an overview of ontology matching systems from the perspective of representation learning. The aforementioned systems were classified according to three categories revealing the most relevant aspects of them. Our analysis pointed out that not only a good capacity for generalization, but also a higher level of flexibility can be achieved when representation learning methods were applied. However, there is still room for improvement as the selected works still do not fully explore all the possibilities that representation learning can provide. In this direction, new approaches should explore deep learning models with higher parameter dimensions as well as the integration of background knowledge and model composition.

References

- [1] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *Advances in neural information processing systems* 33 (2020) 1877–1901.
- [2] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, M. Chen, Hierarchical text-conditional image generation with CLIP latents, 2022. URL: <https://arxiv.org/abs/2204.06125>. doi:10.48550/ARXIV.2204.06125. arXiv:2204.06125.
- [3] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, et al., A generalist agent, *arXiv preprint arXiv:2205.06175* (2022).
- [4] Z. Chen, Y. Wang, B. Zhao, J. Cheng, X. Zhao, Z. Duan, Knowledge graph completion: A review, *IEEE Access* 8 (2020) 192435–192456.
- [5] S. Ji, S. Pan, E. Cambria, P. Marttinen, S. Y. Philip, A survey on knowledge graphs: Representation, acquisition, and applications, *IEEE Transactions on Neural Networks and Learning Systems* 33 (2021) 494–514.
- [6] Y. Zhang, X. Wang, S. Lai, S. He, K. Liu, J. Zhao, X. Lv, Ontology matching with word embeddings, in: *Chinese computational linguistics and natural language processing based on naturally annotated big data*, Springer, 2014, pp. 34–45.
- [7] L. Gutiérrez, B. Keith, A systematic literature review on word embeddings, in: *International Conference on Software Process Improvement*, Springer, 2018, pp. 132–141.
- [8] G. Li, Improving biomedical ontology matching using domain-specific word embeddings, in: *Proceedings of the 4th International Conference on Computer Science and Application Engineering*, 2020, pp. 1–5.
- [9] P. Ristoski, J. Rosati, T. Di Noia, R. De Leone, H. Paulheim, RDF2Vec: RDF graph embeddings and their applications, *Semantic Web* 10 (2019) 721–752.
- [10] J. Portisch, M. Hladik, H. Paulheim, ALOD2Vec matcher results for OAEI 2020, in: *Proceedings of the 15th Workshop on Ontology Matching*, 2020, pp. 147–153.

- [11] J. Portisch, G. Costa, K. Stefani, K. Kreplin, M. Hladik, H. Paulheim, Ontology matching through absolute orientation of embedding spaces, arXiv preprint arXiv:2204.04040 (2022).
- [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.
- [13] J. Chen, E. Jiménez-Ruiz, I. Horrocks, D. Antonyrajah, A. Hadian, J. Lee, Augmenting ontology alignment by semantic embedding and distant supervision, in: R. Verborgh, K. Hose, H. Paulheim, P. Champin, M. Maleshkova, Ó. Corcho, P. Ristoski, M. Alam (Eds.), *The Semantic Web - 18th International Conference, ESWC 2021, Virtual Event, June 6-10, 2021, Proceedings*, volume 12731 of *Lecture Notes in Computer Science*, Springer, 2021, pp. 392–408. URL: https://doi.org/10.1007/978-3-030-77385-4_23. doi:10.1007/978-3-030-77385-4_23.
- [14] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, S. Y. Philip, A comprehensive survey on graph neural networks, *IEEE transactions on neural networks and learning systems* 32 (2020) 4–24.
- [15] J. Chakraborty, S. K. Bansal, L. Virgili, K. Konar, B. Yaman, Ontoconnect: Unsupervised ontology alignment with recursive neural network, in: *Proceedings of the 36th Annual ACM Symposium on Applied Computing*, 2021, pp. 1874–1882.
- [16] J. Wu, J. Lv, H. Guo, S. Ma, Daeom: A deep attentional embedding approach for biomedical ontology matching, *Applied Sciences* 10 (2020) 7909.
- [17] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Lio, Y. Bengio, Graph attention networks, *stat* 1050 (2017) 20.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in Neural Information Processing Systems* 30 (2017).
- [19] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, arXiv preprint arXiv:1810.04805 (2018).
- [20] Y. He, J. Chen, D. Antonyrajah, I. Horrocks, Bertmap: A bert-based ontology alignment system, arXiv preprint arXiv:2112.02682 (2021).
- [21] L. Knorr, J. Portisch, Fine-TOM matcher results for OAEI 2021, in: *Proceedings of the 16th Workshop on Ontology Matching*, volume 3063, 2022, pp. 144–151.
- [22] D. Kossack, N. Borg, L. Knorr, J. Portisch, TOM matcher results for OAEI 2021, in: *Proceedings of the 16th Workshop on Ontology Matching*, 2021, pp. 144–151.
- [23] X. Qiu, T. Sun, Y. Xu, Y. Shao, N. Dai, X. Huang, Pre-trained models for natural language processing: A survey, *CoRR* abs/2003.08271 (2020). URL: <https://arxiv.org/abs/2003.08271>. arXiv:2003.08271.
- [24] C. Xiang, T. Jiang, B. Chang, Z. Sui, Ersom: A structural ontology matching approach using automatically learned entity representation, in: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2015, pp. 2419–2429.
- [25] P. Kolyvakis, A. Kalousis, D. Kiritsis, Deepalignment: Unsupervised ontology matching with refined word vectors, in: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2018, pp. 787–798.
- [26] P. Kolyvakis, A. Kalousis, B. Smith, D. Kiritsis, Biomedical ontology alignment: an approach based on representation learning, *Journal of Bio. Semantics* 9 (2018) 1–20.
- [27] I. Nkisi-Orji, N. Wiratunga, S. Massie, K.-Y. Hui, R. Heaven, Ontology alignment based

- on word embedding and random forest classification, in: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer, 2018, pp. 557–572.
- [28] D. Gromann, T. Declerck, Comparing pretrained multilingual word embeddings on an ontology alignment task, in: Proceedings of the eleventh international conference on language resources and evaluation (LREC), 2018.
- [29] E. Jiménez-Ruiz, A. Agibetov, M. Samwald, V. Cross, Breaking-down the ontology alignment task with a lexical index and neural embeddings, arXiv preprint arXiv:1805.12402 (2018).
- [30] L. L. Wang, C. Bhagavatula, M. Neumann, K. Lo, C. Wilhelm, W. Ammar, Ontology alignment in the biomedical domain using entity definitions and context, arXiv preprint arXiv:1806.07976 (2018).
- [31] W. Li, X. Duan, M. Wang, X. Zhang, G. Qi, Multi-view embedding for biomedical ontology matching., Proceedings of the 14th Workshop on Ontology Matching 2536 (2019) 13–24.
- [32] J. Liu, Y. Tang, X. Xu, HISDOM: A Hybrid Ontology Mapping System based on Convolutional Neural Network and Dynamic Weight, in: Proceedings of the 6th IEEE/ACM International Conf. on Big Data Computing, Applications and Tech., 2019, pp. 67–70.
- [33] M. Tounsi Dhouib, C. Faron Zucker, A. G. Tettamanzi, An ontology alignment approach combining word embedding and the radius measure, in: International Conference on Semantic Systems, Springer, Cham, 2019, pp. 191–197.
- [34] S. Hertling, H. Paulheim, DOME results for OAEI 2019, in: Proceedings of the 14th Workshop on Ontology Matching, 2019, pp. 123–130.
- [35] G. Li, Deepfca: Matching biomedical ontologies using formal concept analysis embedding techniques, in: Proceedings of the 4th International Conference on Medical and Health Informatics, 2020, pp. 259–265.
- [36] A. Bento, A. Zouaq, M. Gagnon, Ontology matching using convolutional neural networks, in: Proceedings of the 12th Language Resources and Evaluation Conference, 2020, pp. 5648–5653.
- [37] 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, 2021, pp. 10780–10792.
- [38] Z. Wang, I. F. Cruz, AgreementMakerDeep results for OAEI 2021, in: Proceedings of the 16th Workshop on Ontology Matching, 2021, pp. 124–130.
- [39] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, Advances in neural information processing systems 26 (2013).
- [40] S. T. Dumais, et al., Latent semantic analysis, Annu. Rev. Inf. Sci. Technol. 38 (2004).
- [41] S. Hertling, J. Portisch, H. Paulheim, Matching with Transformers in MELT, arXiv preprint arXiv:2109.07401 (2021).
- [42] L. Wu, A. Fisch, S. Chopra, K. Adams, A. Bordes, J. Weston, Starspace: Embed all the things!, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 32, 2018.
- [43] J. H. Lau, T. Baldwin, An empirical evaluation of doc2vec with practical insights into document embedding generation, arXiv preprint arXiv:1607.05368 (2016).
- [44] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in

vector space, arXiv preprint arXiv:1301.3781 (2013).

- [45] Y. Yu, X. Si, C. Hu, J. Zhang, A review of recurrent neural networks: Lstm cells and network architectures, *Neural computation* 31 (2019) 1235–1270.
- [46] D. Cer, Y. Yang, S.-y. Kong, N. Hua, N. Limtiaco, R. S. John, N. Constant, M. Guajardo-Cespedes, S. Yuan, C. Tar, et al., Universal sentence encoder, arXiv preprint arXiv:1803.11175 (2018).
- [47] Z. Sun, Z.-H. Deng, J.-Y. Nie, J. Tang, Rotate: Knowledge graph embedding by relational rotation in complex space, arXiv preprint arXiv:1902.10197 (2019).
- [48] Y. He, J. Chen, D. Antonyrajah, I. Horrocks, Biomedical ontology alignment with BERT, in: *Proceedings of the 16th Workshop on Ontology Matching*, 2021, pp. 1–12.
- [49] T. Kenter, A. Borisov, M. De Rijke, Siamese cbow: Optimizing word embeddings for sentence representations, arXiv preprint arXiv:1606.04640 (2016).
- [50] X. Lu, Y. Tsao, S. Matsuda, C. Hori, Speech enhancement based on deep denoising autoencoder., in: *Interspeech*, volume 2013, 2013, pp. 436–440.
- [51] B. Ganter, R. Wille, *Formal concept analysis: mathematical foundations*, Springer Science & Business Media, 2012.
- [52] R. Sennrich, B. Haddow, A. Birch, Neural machine translation of rare words with subword units, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers, The Association for Computer Linguistics*, 2016. URL: <https://doi.org/10.18653/v1/p16-1162>.
- [53] A. Conneau, H. Schwenk, L. Barrault, Y. LeCun, Very deep convolutional networks for text classification, in: *Proceedings of the 15th EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 1: Long Papers*, 2017, pp. 1107–1116. URL: <https://doi.org/10.18653/v1/e17-1104>.
- [54] X. Wang, H. Pham, P. Arthur, G. Neubig, Multilingual neural machine translation with soft decoupled encoding, arXiv preprint arXiv:1902.03499 (2019).
- [55] M. Kulmanov, F. Z. Smaili, X. Gao, R. Hoehndorf, Semantic similarity and machine learning with ontologies, *Briefings Bioinform.* 22 (2021). URL: <https://doi.org/10.1093/bib/bbaa199>.
- [56] T. Zhao, X. Zhang, S. Wang, Graphsmote: Imbalanced node classification on graphs with graph neural networks, in: L. Lewin-Eytan, D. Carmel, E. Yom-Tov, E. Agichtein, E. Gabrilovich (Eds.), *WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, March 8-12, 2021, ACM*, 2021, pp. 833–841. URL: <https://doi.org/10.1145/3437963.3441720>. doi:10.1145/3437963.3441720.
- [57] A. Solimando, E. Jiménez-Ruiz, G. Guerrini, Detecting and correcting conservativity principle violations in ontology-to-ontology mappings, in: *International Semantic Web Conference*, Springer, 2014, pp. 1–16.
- [58] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, P. Abbeel, Infogan: Interpretable representation learning by information maximizing generative adversarial nets, *Advances in Neural Information Processing Systems* 29 (2016).
- [59] J. An, S. Cho, Variational autoencoder based anomaly detection using reconstruction probability, *Special Lecture on IE 2* (2015) 1–18.
- [60] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, A. Lerchner, beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework, in: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.