

Truveta Mapper: A Zero-shot Ontology Alignment Framework

Mariyam Amir, Murchana Baruah, Mahsa Eslamialishah, Sina Ehsani,
Alireza Bahramali, Sadra Naddaf-Sh and Saman Zarandioon

Truveta, Bellevue, WA 98004, USA

Abstract

In this paper, a new perspective is suggested for unsupervised Ontology Matching (OM) or Ontology Alignment (OA) by treating it as a translation task. Ontologies are represented as graphs, and the translation is performed from a node in the source ontology graph to a path in the target ontology graph. The proposed framework, Truveta Mapper (TM), leverages a multi-task sequence-to-sequence transformer model to perform alignment across multiple ontologies in a zero-shot, unified and end-to-end manner. Multi-tasking enables the model to implicitly learn the relationship between different ontologies via transfer-learning without requiring any explicit cross-ontology manually labeled data. This also enables the formulated framework to outperform existing solutions for both runtime latency and alignment quality. The model is pre-trained and fine-tuned only on publicly available text corpus and inner-ontologies data. The proposed solution outperforms state-of-the-art approaches, Edit-Similarity, LogMap, AML, BERTMap, and the recently presented new OM frameworks in Ontology Alignment Evaluation Initiative (OAEI22), offers log-linear complexity, and overall makes the OM task efficient and more straightforward without much post-processing involving mapping extension or mapping repair. We are open sourcing our solution⁺.

Keywords

Ontology Alignment, Ontology Matching, multi-task sequence-to-sequence, contextual matching, transfer learning, zero-shot learning, attention-based models, deep learning

1. Introduction

Ontology Matching (OM) or Ontology Alignment (OA) is the process of finding correspondence between the entities of two ontologies. The purpose of this process is to unify data from different sources and reduce heterogeneity, making data more viable for research and development [1]. Classical state-of-the-art (SOTA) approaches on OM are based on non-contextual matching, where the model captures lexical similarity but fails to understand textual semantics. On the other hand, with contextual approaches, the objective is to match complex pairs which are lexically different but semantically similar and vice-versa.

Recently, a transformer-based contextual framework using BERT [2], has been proposed in [3], which showed promising results compared to other OM systems. In their approach the existing pre-trained BERT model was used which allowed to learn the textual semantics.

⁺<https://github.com/sadransh/ontology-matching-framework>

OM-2023: The 18th International Workshop on Ontology Matching collocated with the 22nd International Semantic Web Conference ISWC-2023 November 7th, 2023, Athens, Greece



© 2023 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)

However, it did not understand the ontology graph structure, which could significantly extend the capabilities of ontologies graph matching.

Motivated by the potential of the transformer models for understanding textual semantic context and overcoming the limitations in the existing methods, the present work proposes Truveta Mapper (TM), a novel zero-shot sequence-to-sequence multi-task transformer-based framework for OM, with the capability of learning both the graph-structure and textual semantics of the ontologies. The model is first pre-trained to learn the hierarchical graph structure of ontology and semantics of each class using Masked Language Modeling (MLM), then fine-tuned using class labels and synonyms as input and class hierarchical-ID as the output, capturing the structure of the ontology. As such, we treat OM as a translation task, where the source ontology class is translated to a path in the matching target ontology class in a zero-shot and multitask manner. Proposed approach is based on zero-shot learning and prediction, where “zero-shot learning” refers to the ability of the model to make source-to-target predictions without requiring manually labeled cross-ontologies matching pairs, and “zero-shot prediction” performs end-to-end mapping from the source to the target without the need of similarity calculation across the entire/subset target ontology or post-processing like extension/repair. With multi-task training, a single model is capable of matching different ontologies such as SNOMED to FMA, SNOMED to NCIT, and so on, by taking advantage of transfer learning.

In this work, empirical comparison is made with the state-of-the-art lexical matching approaches and the recent contextual models presented in [4, 5] on the Unified Medical Language System (UMLS) datasets as part of the Bio-ML track for OAEI 2022. The Ontology Alignment Evaluation Initiative (OAEI) organizes yearly campaigns on ontology matching tasks. Our solution surpasses state-of-the-art LogMap, AML models, Edit-similarity, and recently proposed BERTMap, AMD, LogMap-Lite, BERTMap-Lite, LSMatch, Matcha and ATMatcher.

The remainder of this paper is as follows. Section 2 reviews the recent SOTA-related works on OM/OA; Section 3 defines the problem statement; Section 4 elaborates on pre-training, fine-tuning, zero-shot learning, and predictions; Section 5 shows the overall model performance; and lastly, Section 6 provides conclusion and outlines our potential future works.

2. Related Work

OM classical approaches are primarily based on non-contextual matching. Related to that, some notable works in the field of OM include Edit-Similarity [6], LSMatch [7], LogMap [8], and AgreementMakerLight (AML) [9], among others. Edit-Similarity is a naïve lexical matching approach based on normalized edit similarity scores. LSMatch is another lexical matching approach based on string similarity match. LogMap and AML are two classical OM systems with leading performance in many equivalence matching tasks. These two approaches are based on lexical matching, mapping extension (adding new mappings for semantically related classes of the current mappings), and mapping repair (removing mappings that can lead to logical conflicts). However, these lexical approaches do not consider contextual semantics.

Several OM systems, such as OntoEmma [10], DeepAlignment [11], VeeAlign [12] leveraged dense word embeddings, in which words are projected into a vector. Word pairs with smaller Euclidean distances in the vector space have closer semantic meanings. Different techniques

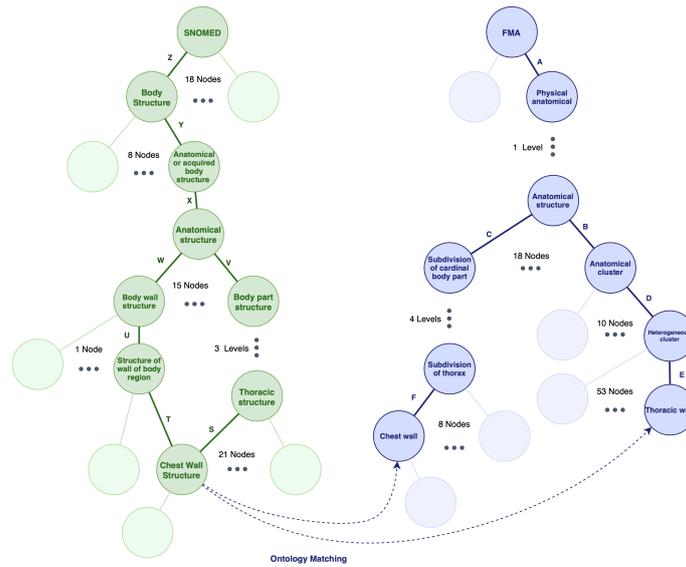


Figure 1: The equivalence matching between the SNOMED class ID 78904004 – “Chest Wall Structure” and two FMA concepts, “Wall of thorax” with ID fma10428 and “Chest wall” with ID fma50060. TM translates from the source node encoding “Chest Wall Structure” in the SNOMED graph to the highlighted path “A ...C ...F” (Chest Wall) and “A ...B ...E” (Thoracic Wall) in FMA ontology.

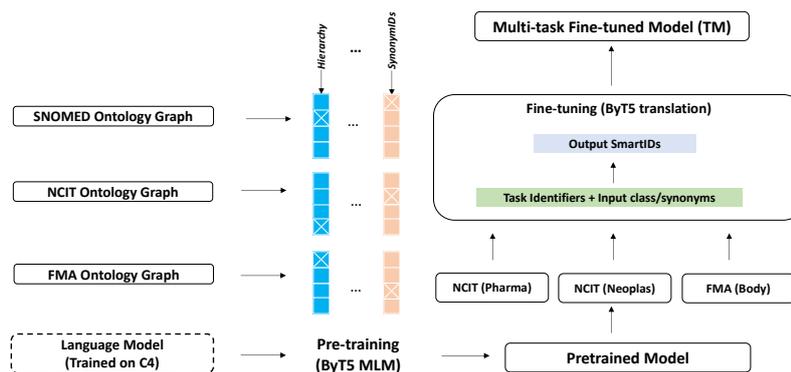


Figure 2: Training Architecture. Pre-training is done using MLM on the full ontology graphs and fine-tuning is done on downstream tasks, translating from the source ontology class descriptions (label and synonyms) to the target ontology node path (hierarchical-IDs). Pre-training and fine-tuning are done in a multi-task manner. Pre-training is performed on both source and target inner-ontologies, and fine-tuning is done on target subset ontologies.

are used to generate these embeddings. OntoEmma and [13] uses word2vec [14], which is trained on Wikipedia; [15] uses FastText [16]; LogMap-ML [17] uses OWL2Vec* [18], which is a word2vec model trained on corpora extracted from the ontology with different kinds of semantics; DeepAlignment uses refined word embeddings using counter-fitting; VeeAlign

proposes dual embeddings using class labels; [19] employs hyperbolic embeddings. These are primarily traditional non-contextual word embedding methods and do not consider contextual meaning. Some of these approaches, such as VeeAlign, are based on supervised training, which requires high-quality labeled mappings for training and can be challenging to obtain.

Transformer-based models [20], thanks to their ability to learn textual contexts, obtained SOTA for several tasks in natural language processing such as machine translation [21, 22, 23], question answering [24], among others. Similarly, in the field of OM, recent developments have also shown the potential of using transformer-based frameworks [1, 3, 25]. Neutel and de Boer [1] employed contextual BERT embeddings to match two domain ontologies associated with occupations. Each sentence is embedded using BERT, and similarity is applied to get the scores for OM. More recently, [3] proposed BERTMap model, which is obtained by fine-tuning the already pre-trained BERT model for the binary classification task. The BERTMap model often outperformed non-contextual approaches such as LogMap, AML, and LogMap-ML. AMD [25] is another recent context-based matching approach that uses a BERT-based model to generate mappings and then filters these mappings using graph embedding techniques. Other related ontology matching systems that participated in OAEI 2022 [4] are LogMap-Lite, BERTMap-Lite, Matcha, and ATMatcher.

3. Methodology

3.1. Problem statement

Ontology Matching (OM) or Ontology Alignment (OA) is the process of finding correspondence between the entities/classes of two ontologies [5]. In this work, a new perspective is presented by treating OM as a translation task for equivalence matching and can be mathematically presented as $f(c_1, T)$, where function f gives the matching target ontology class $c_2 \in C_2$, given a source class $c_1 \in C_1$ and alignment task identifier as T . O_1 and O_2 as the source and target ontologies, with C_1 and C_2 being their respective named class sets. Since we are training a multi-task model, a unique identifier is used for each task.

The present work focuses on equivalence matching, where classes having the same semantic meaning in different ontologies are matched with each other. As shown in Figure 1, each ontology is presented in the form of a hierarchical graph structure with parent-child relation, where each class presents a node in the given ontology graph and target class $c_2 \in C_2$ is obtained as a path in the target ontology graph, for a given input node $c_1 \in C_1$ in the input ontology¹.

3.2. Ontologies

In this work, as a part of the OAEI 2022 Bio-ML track [4], we focus on three UMLS equivalence matching tasks, SNOMED to FMA (Body), SNOMED to NCIT (Neoplas), and SNOMED to NCIT (Pharm), in an unsupervised setting from [4], where the matching pairs between these ontologies are only divided into validation (10%) and testing (90%) sets, without any training data. Pharm, Neoplas, and Body are associated with the semantic types of “Pharmacologic

¹Note, each class is presented as a node in the ontology hierarchical graph-structure, as such, class and node are used interchangeably, as appropriate.

Substance”, “Neoplastic Process”, and “Body Part, Organ, or Organ Components” in UMLS, respectively. Based on these semantic types, subset ontologies are provided in [4], and are given as SNOMED (Body), SNOMED (Neoplas), SNOMED (Pharm), FMA (Body), NCIT (Neoplas) and NCIT (Pharm), where the first three are the source and last three are the target ontologies in our matching task. For each of the classes present in the given ontologies, class ID is provided along with its associated label and possible synonyms (class descriptions). For example, in Figure 1, for Snomed ID 78904004, the class label is “Chest Wall Structure,” and its synonyms are “Thoracic Wall” and “Chest Wall”.

4. Truveta Mapper (TM): Proposed approach for OM

Figure 2 demonstrates the training architecture, and more details are provided in the subsequent subsections, starting from the generation of hierarchical-ID to the final zero-shot predictions. We use ByT5 [26] as the base language model which is used for pre-training, fine-tuning and zero-shot predictions. ByT5 is token-free, is pre-trained on the mC4 dataset and supports multi-task training.

4.1. Hierarchical-ID generation

An ontology is represented in the form of a graph where each node represents a class, and the parent and child relations of the ontology serve as connections between classes. Based on this graph structure of each full ontology, hierarchical-IDs are generated for all the classes.

These are constructed by starting from the root node, separated by “-” at each hierarchy level, and traversing through each node in that level as shown in Figure 3. Following this method, a unique ID is generated for each path traversed. As such, for ontologies like SNOMED, where there are multiple paths between the root and any given class, there could be multiple IDs for that node. In such cases, the shortest ID is considered the hierarchical-ID of that node (highlighted in yellow in Figure 3), while the other path IDs are considered its synonymIDs.

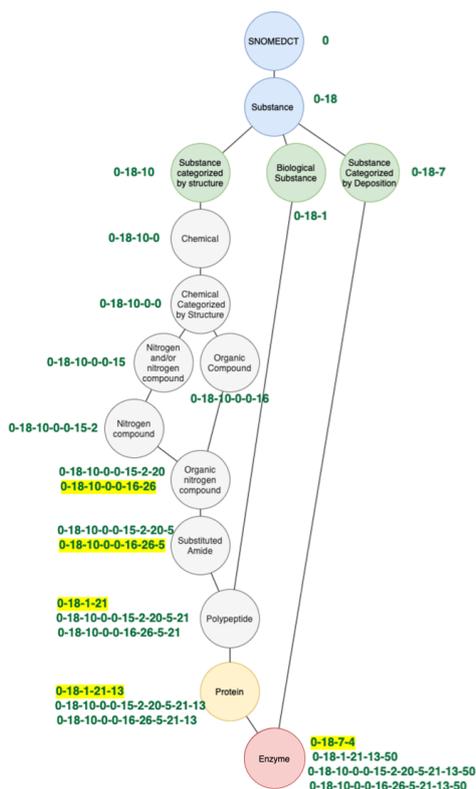


Figure 3: Hierarchical-IDs generation for the Enzyme concept in the SNOMED ontology. The shortest ID (highlighted) is chosen as a Hierarchical-ID, and others are SynonymIDs for this concept.

Pre-training tasks	Example
ID: Child & Parent	CP: 0-2-3-1-43-1 0-2-3-42-1
ID: SYN & SYN	0-2-3-1-43-1 0-2-3-42-1-0
ID & SYN	0-10-40-13-9 S0: Oromucosal solution for gargle
ID & FSN	0-9-18-6-10-28-2 F0: Coagulation or electrocoagulation of inner ear
FSN & SYN	F0: Coagulation or electrocoagulation of inner ear S0: electrocoagulation of ear
Fine-tuning tasks	Example (Source ontology FSN/SYN input, Target ontology ID output)
Pharm	F0: Ceprothin 2- , 5-5-6-4-93
Neoplas	F0: FAB M1 2- , 0-0-6-2-1-2-8-0-0-7
Body	F0: Hairs set 1- , 1-0-0-1-7-0-41

Table 1

Pre-training and fine-tuning tasks examples. Task identifiers are in bold. During pre-training, a random part of the input in the examples get masked. During fine-tuning, root node of the hierarchical-ID is used as a suffix to provide appropriate task identifier. Different identifiers are used for FMA and NCIT, and acronyms SYN and FSN corresponds to synonyms and Fully-Specified-Name (label) respectively.

Each node ID inherently captures the information of all its ancestors. This enables the model to trace from a broader class, starting from the root and getting more granular at each level, thus simplifying the translation task.

4.2. Pre-training

After generating the hierarchical-IDs, multi-task pre-training is done on the full ontologies (SNOMED, FMA, NCIT) using MLM by randomly masking the nodes, enabling the model to learn the hierarchy and semantics. For instance, “Structure of Forel’s H2 bundle” is represented as “1-1-0-0-0-4-1-1-0-0-0-7” and is masked as “1-1-0-0-0-0-[MASK]-1-0-0-0-7”. Furthermore, additional tasks are included in order for the model to learn the semantics of each class in the form of class-level synonyms, labels, and descriptions; class-level relations between child and parent nodes; and the relation between synonym-ID and hierarchical-ID, using separate identifiers for each task in the pre-training step (Figure 2). Task identifiers are added in the form of prefixes, to distinguish between different ontologies. For example, SNOMED ontology is prefixed as “F0:”, where “F” represents fully specified name (class label) and “0” indicates SNOMED Ontology. Similarly, FMA and NCIT are represented using “1” and “2” identifiers. Some representative examples are presented in Table 1, where similar tasks are defined for each ontology, with the objective of learning their hierarchical structure and semantics using MLM.

Based on the tasks stated in Table 1, we generate the pre-training dataset which has 2,406,456 instances constituting SNOMED, NCIT, and FMA ontologies. The model is trained for 3 epochs, with an increasing masking percentage linearly over time, starting at 10% and increasing to 35% in the final batch. The pre-training is done on 8 V100 32GB Nvidia GPUs with a batch size of 20, using a learning rate of 1e-3 with linear decay scheduler and AdamW optimizer.

4.3. Fine-tuning

The fine-tuning step aims to train the model on the downstream OM tasks. Only target subset ontologies, i.e., NCIT (Pharm), NCIT (Neoplas), and FMA (Body), are used for fine-tuning. The

training data of each target sub-ontologies is augmented using the exact matches present in the labels and synonyms of the other subset ontologies. We are also taking advantage of older ontology versions to add more synonyms to each target label. This expands the training corpus, enriches the data with minimal processing, and helps to perform more comprehensive learning. After the data augmentation for all the target sub-ontologies, fine-tuning is performed only on these target sub-ontologies corpora. Training data is generated for each class in the target ontologies, where the input is the class label, synonyms, and descriptions, and output is the corresponding node hierarchical-ID (generated in Section 4.1), using a separate identifier for each task. Similar to our pre-training approach, multi-task fine-tuning is performed on the downstream OM tasks. Some examples are presented in Table 1.

Based on the fine-tuning tasks described in Table 1, we generate the fine-tuning training data which has 462,789 samples from Pharm, Neoplas, and Body subsets. Using 8 Nvidia V100 32GB GPUs with a batch size of 20, the fine-tuning took around 21 epochs. For the fine-tuning, a learning rate of 1e-3 with linear decay scheduler and warm-up of 1.5 epoch using AdamW optimizer with eps of 1e-8 and weight decay of 1e-2 is used.

4.4. Zero-shot Predictions

TM is a multi-task model with the capability to translate between multiple ontologies from the input source class labels/synonyms to target hierarchical-IDs. Thus, given a source term, the model predicts the potential candidate in the target ontology graph. For confidence scoring, two approaches have been adopted here: (i) Greedy search score: Scores are generated based on greedy search with softmax probabilities using temperature scaling. This is a naive way to compute the confidence directly from the model prediction. (ii) Using embeddings: This is a sophisticated method proposed to make the TM predictions more robust and improve model precision, by leveraging semantic similarity using embeddings of source terms and predicted target candidates. Using the same model, the embeddings are generated for the target candidate and the similarity score is obtained between the source term and predicted target term embeddings. Scores are generated across the source and predicted class labels and synonyms, all of which are also augmented by singularization. The maximum generated score is considered as the similarity score. As such, the proposed model takes advantage of both graph search and semantic matching. Mathematically, similarity score S is given as:

$$S = \begin{cases} 1.0, & \text{if } \Omega(c_1) \cap \Omega(c_2) \neq \emptyset \\ \max(\text{Sim}(\Omega(c_1), \Omega(c_2))), & \text{otherwise} \end{cases} \quad (1)$$

where c_2 is the predicted class for c_1 , $\Omega(c_1)$ and $\Omega(c_2)$ are sets of labels and synonyms for c_1 and c_2 , respectively, and $\max(\text{Sim}(\Omega(c_1), \Omega(c_2)))$ selects the maximum cosine similarity score across all the labels and synonyms of c_1 (source) and c_2 (predicted). If an exact match is available between the labels and synonyms of source and target classes, we assign a maximum similarity score, since embedding similarity will also give a similar result. The source and the target candidates are considered valid mapping pairs if their similarity score exceeds a selected threshold for both the approaches.

One of the main advantages of our proposed TM is that it reduces the time complexity to log-linear as opposed to the naive solution of search that results in quadratic complexity². Given an input term with a specified task identifier, TM is able to predict the best possible match from the target ontology with $O(\log(n))$ complexity, where n corresponds to the number of nodes in the target ontology graph (same as the number of classes). Overall, TM reduces the time-complexity to $O(n\log(n))$, noting that a single search in a tree structure with n nodes can be performed in $O(\log(n))$ time.

5. Results

5.1. Evaluation criteria

Commonly used metrics for evaluating OM systems [5]: Precision (P), Recall (R), and F-score are used as the global evaluation metrics. Mathematically,

$$P = \frac{|M_{out} \cap M_{ref}|}{|M_{out}|}, \quad R = \frac{|M_{out} \cap M_{ref}|}{|M_{ref}|}, \quad F_\beta = (1 + \beta^2) \frac{P \cdot R}{\beta^2 \cdot P + R} \quad (2)$$

where, M_{ref} are the reference mappings, consisting of matching pairs, $m = (c, c')$, such that c and c' are two classes from the to-be-aligned ontologies, and M_{out} are the mappings computed by OM systems and $\beta = 1$.

Local evaluation metrics, *Hits@K* and Mean Reciprocal Rank (*MRR*), introduced in [5] are also used for current evaluation and can be represented as:

$$Hits@K = \frac{|\{m \in M_{ref} | Rank(m) \leq K\}|}{|M_{ref}|}, \quad MRR = \frac{\sum_{m \in M_{ref}} Rank(m)^{-1}}{|M_{ref}|} \quad (3)$$

where $Rank(m)$ returns the ranking position of m among $M_m \cup \{m\}$ according to their scores, M_m represents a set of negative mappings pairs for each of the source term c in M_{ref} , such that $(c, c'_i) \in M_m$ with $i \in \{1, 2, \dots, 100\}$ and c'_i are the 100 negative output candidates from target ontologies for each of the source terms c in M_{ref} . As such, the Hits and MRR would be different for different selected 100 samples. We have published the results of our model based on the provided M_m set in [5] for a fair comparison. To provide a more robust measure of local metrics, we are reporting overall accuracy as well, although this is not provided for any of the other models. Accuracy here can be mathematically presented as:

$$Accuracy = \frac{|\{m \in M_{ref} | f(c, T) = c'\}|}{|M_{ref}|} \quad (4)$$

²Note that BERTMap reduces the time complexity from $O(n^2)$ in traditional approaches to $O(kn)$, where $k \ll n$ with an additional preprocessing step by considering only a small portion of target subset ontology classes with at least one subword token common to the source class candidate, which adds dependency on the tokenizer and could be error prone since some semantically matching cases with lexical variations could get filtered out in this process. Contrary to that, such limitation does not exist in TM since it performs matching from source to target without reducing the target corpora size.

Task	Precision	Recall	F-score	MRR	Hit@1	Accuracy
TM(Ours) ¹	0.947	0.738	0.830	0.960	0.942	0.801
TM(Ours) ²	0.960	0.720	0.823			
Edit-Similarity*	0.976	0.660	0.787	0.895	0.869	NA
LogMap*	0.702	0.581	0.636	0.545	0.330	NA
AML*	0.841	0.776	0.807	NA	NA	NA
BERTMap*	0.997	0.639	0.773	0.954	0.930	NA
LogMap-Lite**	0.967	0.543	0.695	NA	NA	NA
AMD**	0.890	0.704	0.786	NA	NA	NA
BERTMap-Lite**	0.976	0.660	0.787	0.895	0.869	NA
Matcha**	0.875	0.594	0.707	NA	NA	NA
ATMatcher**	0.264	0.226	0.244	NA	NA	NA
LSMatch**	0.809	0.072	0.132	NA	NA	NA

Table 2

Result for equivalence matching – SNOMED (Body) to FMA (Body).

where $m = (c, c')$ represents matching pairs in the M_{ref} set, and $f(c, T)$ refers to the target candidate predicted by the model, given an input term c and appropriate task identifier T .

Baselines. Results are compared with the SOTA approaches: Edit-Similarity, LogMap, AML, BERTMap [5], and recently published results in [4]. To be consistent, evaluation for P, R, F-score, Hit@1, and MRR is done using [6] library.

5.2. Prediction Results

Prediction results are shown in Tables 2–4, for the three equivalence OM tasks, from SNOMED to FMA (Body), SNOMED to NCIT (Pharm), and SNOMED to NCIT (Neoplas). The results demonstrate precision, recall, F-score, Hit@1, MRR, and accuracy for TM and baseline approaches presented in He et al. [5] and OAEI [4] on the test data for the unsupervised setting. In the given tables, superscripts^{1,2} are based on our proposed TM, where the former is based on embedding similarity score and later is based on greedy search score, superscript* results are based on He et al. [5] and we used the same evaluation metrics for TM, and superscript** correspond to OAEI [4] published results. The highest numbers for each of these metrics are highlighted in the tables to emphasize which model is outperforming others in each category.

The overall results illustrate that TM is outperforming all the baselines for all three OM tasks in F-score, Hit@1, and MRR. A high threshold is selected to generate the most confident cross-ontology matching pairs. Note that a single unified model is trained and leveraged here to predict all the results in the form of a source class to target hierarchical-IDs, using appropriate task identifiers.

There are two TM results presented in the given tables, and both are based on different scoring schemes. TM² is based on greedy search scores and TM¹ is based on a new and more robust prediction scheme using embeddings described in Subsection 4.4, taking advantage of both graph search and semantic similarity. It can be seen that both of our methods surpass SOTA for all the tasks, but TM¹ is more robust and has significant improvements as compared

Task	Precision	Recall	F-score	MRR	Hit@1	Accuracy
TM(Ours) ¹	0.972	0.929	0.950	0.987	0.982	0.946
TM(Ours) ²	0.977	0.872	0.922			
Edit-Similarity*	0.979	0.432	0.600	0.836	0.760	NA
LogMap*	0.915	0.612	0.733	0.820	0.695	NA
AML*	0.940	0.615	0.743	NA	NA	NA
BERTMap*	0.966	0.606	0.745	0.919	0.876	NA
LogMap-Lite**	0.995	0.598	0.747	NA	NA	NA
AMD**	0.962	0.745	0.840	NA	NA	NA
BERTMap-Lite**	0.979	0.432	0.600	0.836	0.760	NA
Matcha**	0.941	0.613	0.742	NA	NA	NA
ATMatcher**	0.937	0.566	0.706	NA	NA	NA
LSMatch**	0.982	0.551	0.706	NA	NA	NA

Table 3

Results for equivalence matching – SNOMED (Pharm) to NCIT (Pharm).

Task	Precision	Recall	F-score	MRR	Hit@1	Accuracy
TM(Ours) ¹	0.809	0.795	0.802	0.962	0.944	0.802
TM(Ours) ²	0.812	0.773	0.792			
Edit-Similarity*	0.815	0.709	0.759	0.900	0.876	NA
LogMap*	0.823	0.547	0.657	0.824	0.747	NA
AML*	0.747	0.554	0.636	NA	NA	NA
BERTMap*	0.655	0.777	0.711	0.960	0.939	NA
LogMap-Lite**	0.947	0.520	0.671	NA	NA	NA
AMD**	0.836	0.534	0.652	NA	NA	NA
BERTMap-Lite**	0.815	0.709	0.759	0.900	0.876	NA
Matcha**	0.754	0.564	0.645	NA	NA	NA
ATMatcher**	0.866	0.284	0.428	NA	NA	NA
LSMatch**	0.902	0.238	0.377	NA	NA	NA

Table 4

Result for equivalence matching – SNOMED (Neoplas) to NCIT (Neoplas).

to any of the existing methods. To be precise, 2.3% improvement over the second best result (AML) in Body, 11.0% improvement for Pharm (as compared to AMD), and 4.3% improvement for Neoplas as compared to BertMap-Lite and Edit-Similarity, is seen for TM¹ in the F-score. It should be noted that even without TM, none of these methods are SOTA in all the tasks.

For generating local metrics for Hit@1 and MRR, TM is used to generate the embedding similarity score of input terms in the test set and their corresponding candidates in $M_m \cup \{m\}$ set. We are also outperforming all existing SOTA methods based on MRR and Hit@1. Additionally, we are reporting accuracy metric, which is consistent, and more representative of the model performance. For this metric, the TM predictions are obtained across the entire target ontology without using any smaller subset of negative samples from the test set, while reducing the time complexity from quadratic to log-linear.

6. Conclusions and Discussions

This work presents a new approach to OM by treating the OM process as a translation task and performing multi-task pre-training, fine-tuning, and predictions in a zero-shot, unified and end-to-end manner. The proposed approach takes advantage of transfer learning across different ontologies and does not require manual annotations for training. Additionally, the trained model understands the semantics of the text as well as the structure of the ontologies. We show that our proposed method outperforms Edit-Similarity, LogMap, AML, BERTMap, and the recently proposed OM frameworks in the OM22 conference [4] in all the tasks.

Our approach provides several advantages: (1) It reduces the time complexity to log-linear during inference, (2) It is based on zero-shot prediction, without requiring much post-processing and does not employ mapping extension or mapping repair in contrast to the other methods, (3) It does not require any manual labeled cross-ontologies matching pairs due to zero-shot learning, (4) One unified framework is used as a result of multi-tasking, which makes it easier to productionize these large transformer-based models, (5) It is robust toward different tokenization schemes as it uses byte level tokenization, (6) It learns complete ontologies graphs, using the hierarchical-IDs which provides a more natural path for translation, and would be significantly helpful for subsumption mappings.

In the future, we will pre-train the starting checkpoint with more domain-related corpus (e.g., PubMed, MIMIC-III, clinical notes) instead of the mC4 dataset. Another interesting track can be ensemble learning of existing SOTA models with TM.

References

- [1] S. Neutel, M. H. de Boer, Towards Automatic Ontology Alignment using BERT., in: AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering, 2021.
- [2] 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).
- [3] Y. He, J. Chen, D. Antonyrajah, I. Horrocks, BERTMap: a BERT-based ontology alignment system, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 36, 2022, pp. 5684–5691.
- [4] OAEI, <https://www.cs.ox.ac.uk/isg/projects/ConCur/oei/2022/> Accessed: 2022-08-08 (2022).
- [5] Y. He, J. Chen, H. Dong, E. Jiménez-Ruiz, A. Hadian, I. Horrocks, Machine Learning-Friendly Biomedical Datasets for Equivalence and Subsumption Ontology Matching, arXiv preprint arXiv:2205.03447 (2022).
- [6] DeepOnto, <https://github.com/KRR-Oxford/DeepOnto> Accessed: 2022-10-14 (2022).
- [7] A. Sharma, A. Patel, S. Jain, LSMatch and LSMatch-Multilingual Results for OAEI (2022).
- [8] E. Jiménez-Ruiz, B. Cuenca Grau, Logmap: Logic-based and scalable ontology matching, in: International Semantic Web Conference, Springer, 2011, pp. 273–288.
- [9] D. Faria, C. Pesquita, E. Santos, M. Palmonari, I. F. Cruz, F. M. Couto, The agreementmakerlight ontology matching system, in: OTM Confederated International Conferences” On the Move to Meaningful Internet Systems”, Springer, 2013, pp. 527–541.

- [10] 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).
- [11] 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, Volume 1 (Long Papers), 2018, pp. 787–798.
- [12] V. Iyer, A. Agarwal, H. Kumar, VeeAlign: a supervised deep learning approach to ontology alignment., in: OM@ ISWC, 2020, pp. 216–224.
- [13] 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.
- [14] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, arXiv preprint arXiv:1301.3781 (2013).
- [15] 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.
- [16] P. Bojanowski, E. Grave, A. Joulin, T. Mikolov, Enriching word vectors with subword information, Transactions of the association for computational linguistics 5 (2017) 135–146.
- [17] J. Chen, E. Jiménez-Ruiz, I. Horrocks, D. Antonyrajah, A. Hadian, J. Lee, Augmenting ontology alignment by semantic embedding and distant supervision, in: European Semantic Web Conference, Springer, 2021, pp. 392–408.
- [18] 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.
- [19] P. Kolyvakis, Approaching Ontology Alignment through Representation Learning to Bridge the Semantic Gap in Engineering Applications, Technical Report, EPFL, 2020.
- [20] 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).
- [21] M. Johnson, M. Schuster, Q. V. Le, et al., Google’s multilingual neural machine translation system: Enabling zero-shot translation, Transactions of the Association for Computational Linguistics 5 (2017) 339–351.
- [22] W. Xu, M. Carpuat, Editor: an edit-based transformer with repositioning for neural machine translation with soft lexical constraints, Transactions of the Association for Computational Linguistics 9 (2021) 311–328.
- [23] H.-I. Liu, W.-L. Chen, Re-transformer: a self-attention based model for machine translation, Procedia Computer Science 189 (2021) 3–10.
- [24] C. Clark, K. Lee, M.-W. Chang, T. Kwiatkowski, M. Collins, K. Toutanova, BoolQ: Exploring the surprising difficulty of natural yes/no questions, arXiv preprint arXiv:1905.10044 (2019).
- [25] Z. Wang, AMD Results for OAEI 2022 (2022).
- [26] L. Xue, A. Barua, N. Constant, R. Al-Rfou, S. Narang, M. Kale, A. Roberts, C. Raffel, Byt5: Towards a token-free future with pre-trained byte-to-byte models, Transactions of the Association for Computational Linguistics 10 (2022) 291–306.