

HybridOM: Ontology Matching using Hybrid Search

Marius-Horatiu Totoian^{1,*}, Anca Marginean¹, Philipp Blohm² and Mir Nawab Hussain²

¹Technical University of Cluj-Napoca, Cluj-Napoca, 400391, Romania

²msg systems ag, Ismaning (Munich), DE-85737, Germany

Abstract

Ontology matching targets identical concepts from different ontologies with the final purpose of interoperability and ontologies merging. The matching task is not restricted to ontologies, it is also relevant for knowledge graphs. Ontology matching solutions based on transformer-based embeddings, textual similarity, logical mapping, or Large Language Models (LLMs) are still facing problems, mainly due to the lack of uniform information about the concepts and lack of homogeneous semantic granularity along different ontologies. In this work, we present a framework that combines vector-based similarity and string-based similarity through hybrid searches. LLMs are used to generate descriptions for ontology concepts, hence the concepts' representation is enriched and the alignment process can benefit from both the knowledge captured by the initial ontologies and the extended LLM-generated textual descriptions. The proposed system, HybridOM, is an unsupervised approach independent of the ontologies' domain. HybridOM is evaluated within Bio-ML 2024 track for the task of concept matching. It achieves the highest values for F1-score and Recall for most of the ontology pairs while maintaining a balance between precision and recall. The proposed method has been adapted for industrial usage in a human capital management product called msg.ProfileMap.

Keywords

Ontology matching, hybrid search, Large Language Model, verbalization, vector database.

1. Presentation of the system

1.1. General Statement

Despite the long-standing goal of ontologies to provide a structured representation of domain knowledge, matching two ontologies from the same domain is not a straightforward task. The primary challenges in ontology matching arise from difficulties in creating comprehensive, exhaustive, and standardized representations of meaning.

Since ontologies are diverse, either in the naming schema, the granularity of properties and taxonomical structure, or the level of detailed textual descriptions, we combine different sources for meaning and different meaning representations can help the ontology matching. The considered sources for the meaning are: i) asserted axiom-based descriptions, ii) knowledge derived through symbolic reasoning, iii) textual descriptions from the ontology, and iv) textual descriptions generated with LLMs. The considered meaning representations are: i) axioms in the ontology, ii) plain text used in labels or concept descriptions, iii) text embeddings. The result is a solution for ontology mapping that is domain-agnostic.

In this work, we give a detailed description of our method together with the evaluation results obtained in 2024 BioML task, part of Ontology Alignment Evaluation Initiative. Elements of the current method were also tested in the context of msg.ProfileMap. msg.ProfileMap is a human capital management product that uses an ontology containing competences, job roles, and certificates. New entities can be entered into the ontology automatically by importing existing vocabularies of customers or by extracting them from CVs or job descriptions. Duplicates in the ontology would create problems when searching for employees or during statistical analyses. HybridOM's method supports these processes by identifying the corresponding ontology entities if existing and thus ensuring the quality of the ontology.

OM-2024: The 19th International Workshop on Ontology Matching collocated with the 23rd International Semantic Web Conference (ISWC 2024), November 11th, Baltimore, USA

*Corresponding author.

✉ totoian.em.marius@student.utcluj.ro (M. Totoian); anca.marginean@cs.utcluj.ro (A. Marginean);

Philipp.Blohm@msg.group (P. Blohm); Mir.Nawab.Hussain@msg.group (M. N. Hussain)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).




Figure 1: HybridOM Pipeline

1.2. Solution

Most of the ontology matching systems, such as OLaLa [1] and LLMs4OM [2], are using the retrieval methods to get the top-k candidates and then LLM to select the best one from the retrieved candidates. Differently, HybridOM finds the best candidate as a result of the retrieval step, while it uses LLM as an auxiliary method to generate descriptions involved in the retrieval step. The key elements of HybridOM are: the augmentation of the text-based descriptions of concepts in the ontology and ensembles of hybrid searches. The augmentation sources are the ontology axioms and LLM-generated text.

Hybrid search is a special case of search that combines multiple semantic searches and string-based searches and matches the results into a single ranking output. HybridOM uses hybrid search to extract the best candidates for matching an external concept to concepts in the ontology. Using OpenSearch as a vector database, the data about concepts are stored as string-based and vector embedding fields.

HybridOM architecture consists of a pipeline (Figure 1). The first step is data preprocessing which i) identifies the essential raw data that can be used from ontology and ii) generates descriptions using LLM for each concept based on its medical domain. The ontology axioms are integrated through their verbalization based on the algorithm described in [3]. Once the textual information for each concept is prepared, embeddings are generated for specific fields within a subset of this text-based information using a language model. All the string and vector fields are ingested in OpenSearch. The last step is using hybrid searches that are customized based on i) the included fields and ii) their associated scores. In the end, a sorted list of the best candidates is retrieved. The best concept for matching is the first retrieved candidate. Each concept is mapped independently of the other concepts.

1.2.1. Data preprocessing

For preprocessing of ontology data, python package DeepOnto [4] is used to parse the ontology structure and retrieve useful information about concepts. DeepOnto provides an implementation for the verbalization algorithm [3] and HybridOM uses it for building concept's verbalization. Additionally, a biomedical description is generated for each concept in the ontologies by prompting LLaMa3 [5].

1.2.1.1. Raw Data extraction The most significant fields to be stored in the database are: *label*, *description*, and *synonyms*. Because not every concept has a description or synonyms, other fields are used to ensure a non-empty value. For some ontologies multiple properties can be used to describe or to create synonyms, in which case these fields are concatenated. Table 1 shows the HybridOM configuration of every field based on each ontology's properties. This step is done with the DeepOnto.

1.2.1.2. Verbalization Ontology axioms offer important information about the taxonomy of the ontology, but also a mathematical representation of the relations and properties between concepts. Bio-ML ontologies provide two most important axioms that HybridOM uses in verbalization composition: *SubClassOf* and *EquivalentClasses*. The article [3] describes an algorithm of translation from ontology axiom to a human-readable verbalization for LLM prompting. Using the mentioned algorithm, we build the verbalizations according to Table 2.

In case *EquivalentClasses* axiom is missing, the verbalization is represented just by the subclass axiom or parent concepts extracted using reasoning. In the case of no *SubClassOf* axiom, the verbalization is formed by parents found using the *Elk reasoner*. In case the reasoner does not find any parent, a hardcoded generic parent is used. The verbalizer and the reasoner are provided by DeepOnto package.

Table 1

OpenSearch fields based on ontology raw data. Operation + represents string concatenation using a delimiter, and x/y represents that in case of missing value of x , values of y will be used.

Ontology	Label	Description	Synonyms
OMIM	label	label	(hasExactSynonym + exactMatch) label
ORDO	label	definition label	alternative_term label
NCIT	label	(owl#P97 + owl#P325) label	owl#P90
SNOMED	label	definition prefLabel	altLabel prefLabel
DOID	label	IAO_0000115 label	(hasExactSynonym + hasRelatedSynonym + hasNarrowSynonym) label
FMA	label	definition label	(synonym + English_equivalent + slot_synonym + abbreviation + other_Latin_equivalents + non-English_equivalent) label

Table 2

Verbalization construction. The verbalization of *EquivalentClasses* axiom(\equiv) can be present in all cases below if the axiom exists. For a better table view, \equiv case is presented just in the first case.

Case	Verbalization
\sqsubseteq, \equiv axiom	<i><Concept></i> It is a type of <i><Verb_SubClassOf_{1, <i><Verb_SubClassOf_{2 ... It is an equivalent of <i><Verb_EquivalentClasses_{1, <i><Verb_EquivalentClasses_{2 ...}</i>}</i>}</i>}</i>
No \sqsubseteq axioms	<i><Concept></i> It is a type of <i><Parent_{1, <i><Parent_{2, ...}</i>}</i>
No Parents	<i><Concept></i> It is a type of General Bio Medical Entity.

1.2.1.3. Generating description using LLaMa In most cases, the ontologies do not include a meaningful textual description for each concept. String-based matching can't provide enough context to distinguish between similar concepts but with different names, such as Alcohol and Ethanol. We assume that the existence of a description could help in reducing ambiguity. Inspired by the way LLMs are used in RAG approaches, HybridOM proposes the usage of LLM-generated descriptions as a means to improve the ontology matching by providing useful information that might even extend the ontology's properties. The possible negative impact of using LLM-generated description is limited since hybrid searches combines the generated meaning with the one existing in the ontologies. In case the concept is sensitive (e.g. substance drug related), LLaMa doesn't generate a valid description. HybridOM detects this based on the text length and replaces the result with label.

LLaMa3 is an open-source large language model with impressive results that exceed GPT-3 in various benchmarks [5]. In order to balance the quality with the generation time we chose LLaMa3 8B. ¹

Because the medical domain of all ontologies is known, HybridOM uses specific prompts that generate simple definitions, chemical structures, symptoms, anatomy, etc. The medical domains were extracted from [6]. Since *Disease* and *Neoplas* categories are very similar, the prompt from *Disease* is used also for the ontology from *Neoplas* category. All used prompts are shown in Table 3.

1.2.2. Ingestion

The next step in the pipeline is the embedding generation for textual data extracted or generated from ontology. The fields to be embedded are: *verbalization*, *description*, and *LLM generated description*. While HybridOM can use any language model for the embeddings, gtr-t5-large² and BioLord-2023 [7] are used for Bio-ML task. BioLord-2023 is a language model trained using LLM for generating definitions in the training set and it is state-of-the-art among biomedical models regarding STS (Semantic Textual Similarity) [7]. Once the embeddings are built for each ontology, an *index* is created in OpenSearch for each ontology and the data is ingested in the database.

¹<https://ollama.com/library/llama3>

²<https://huggingface.co/sentence-transformers/gtr-t5-large>

Table 3
LLaMa prompts used to generate descriptions for ontology concepts

Category	Prompt
Pharm	<p>You are an expert from the healthcare and biomedical domain with extensive medical knowledge and practical experience. You will get a medical product as input. Describe the given product. Give complex, detailed, clear and understandable description. The answer must have the following structure:</p> <p>*Structure:</p> <p>Definition: [Detailed definition of the product]</p> <p>Chemical Structure: [Chemical structural components of the product based on chemical formula and chemical processes]</p> <p>Biological Structure: [Biological structural components of the product]</p> <p>Pharmacological Property: [The most important pharmacological property of the product]</p>
Disease	<p>You are an expert from the healthcare diseases and biomedical diseases domain with extensive medical knowledge and practical experience. You will get a medical disease as input. Describe the given disease. Give complex, detailed, clear, and understandable description. The answer must have the following structure:</p> <p>*Structure:</p> <p>Definition: [Detailed definition of the disease]</p> <p>Cause: [The most important biological causes of the disease]</p> <p>Symptoms: [Describe the symptoms of the disease over the biological human body]</p>
Body	<p>You are an expert from the healthcare and biomedical domain with extensive medical knowledge and practical experience. You will get an anatomical human body concept as input. Describe the given anatomical human body concept. Give complex, detailed, clear, and understandable description. The answer must have the following structure:</p> <p>*Structure:</p> <p>Definition: [Detailed definition of human body concept]</p> <p>Anatomy: [Describe the anatomy of the given concept regarding the human body]</p> <p>Function: [The main function of the concept in human body]</p>

1.2.3. Hybrid Search

Hybrid Search is the integration of semantic searches using Approximate k-NN search and keyword searches using BM25 algorithm. The search uses multiple keyword/vector similarity searches that run in parallel. The concepts from the source ontology (*SO*) are searched in all the concepts from the destination ontology (*DO*) based on the existing fields in the database. In the case of *NCIT-DOID* dataset, the *SO* is *NCIT* and the *DO* is *DOID*. All concepts are stored in OpenSearch indexes, one index per ontology. Based on a search configuration, for each search in hybrid search a score is set and the sum of all scores must be equal to 1.

Figure 2 shows the way we use the hybrid search. For keyword/semantic search, one field from *SO* is searched in multiple fields from *DO*. Based on the concept's id and the field names used in the keyword and semantic search, the concept is identified in the database and its fields' values are retrieved. These values are searched in several fields from *DO*.

In the given example (Figure 2) the *SO* keyword field is *label* and the *SO* semantic search field is *description_llm_emb* (embedding of LLM generated description). Two *SO* values are retrieved based on the concept's id by the *Data Module*: the string *Ethanol* and the stored embedding of the LLM description. Based on the given search configuration, the string *Ethanol* is searched in the *DO* fields *label* and *synonyms*, and two k-NN searches are running in parallel for pairs: ***description_llm_emb (SO) - description_emb (DO)*** and ***description_llm_emb (SO) - verbalization_emb (DO)***. In the end the final score is a weighted score based on the configured weights for each search.

The final ranking is calculated by running multiple hybrid searches. The score for ontology matching is the sum of all scores returned by hybrid searches divided by the number of searches.

1.2.3.1. Keyword search Lucene Engine from OpenSearch is used for the keyword search. HybridOM search configuration defines one field from *SO* (*keyword_field*) whose value is to be searched in one

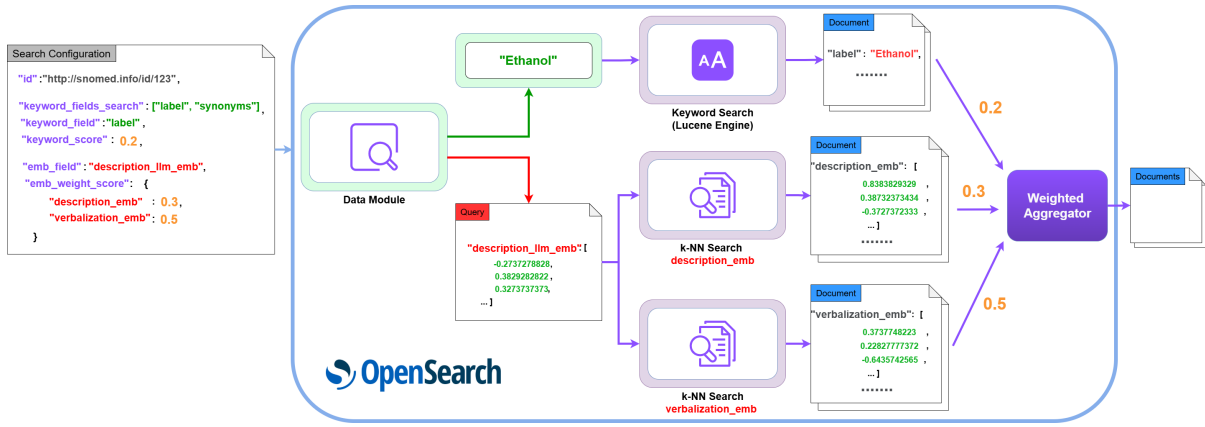


Figure 2: Hybrid Search using OpenSearch

```

"keyword_fields_search" : [⟨dst1⟩^⟨weight1⟩, ...]
"keyword_field" : ⟨src⟩
"keyword_score" : ⟨score⟩
"emb_field" : ⟨src_emb⟩
"emb_weight_score" : {⟨dst_emb1⟩ : ⟨score1⟩,
                    ⟨dst_emb2⟩ : ⟨score2⟩, ...}

```

Figure 3: Keyword/k-NN search configuration template. *dst* - destination field name, *src* - source field name, *src/dst_emb* - source/destination vector field, $\langle dst_1 \rangle^x$ - the importance of finding the field $\langle src \rangle$ in $\langle dst_1 \rangle$ is x .

or more *DO* fields (*keyword_fields_search*) (Figure 3). Only one search is executed and if the searched value is found in multiple *DO* fields the score increases. In the search configuration, it can also be configured the importance of finding a searched value in a certain *DO* keyword field (Figure 3).

For the keyword search, OpenSearch returns a score that is not limit-bounded. Consequently, HybridOM uses a custom $[0, 1]$ normalization implemented with OpenSearch sigmoid function.

1.2.3.2. Approximate k-NN search OpenSearch uses an approximate k-NN search for retrieving the most relevant vectors using cosine similarity as a metric. The default value of *k* is 100. The configuration template for the fields and weights of a semantic search is given in the right part of Figure 3.

1.2.3.3. Normalization A simple hybrid search aggregates the results from each search and computes the final score based on the weights from the configuration. Each search retrieves a fixed number of candidates (configured). The usage of normalization in this context transforms each score for the retrieved candidate using min_max normalization method. The best candidate will have the score **1** and the difference in ranking will be stronger marked among candidates by the normalization process.

1.2.3.4. Dual hybrid search The current flow uses hybrid searches to find the concepts from *SO* in *DO*. To reverse the search, after the candidates for a concept are retrieved, a predefined number of candidates are searched backward to discover their matching candidates from *SO*. The same search configuration is used for the backward searching (*DO* to *SO*). In the case of the second searching flow, the mapping concept is known (the one from *SO*) and if it is found on the reverse search, the score will be taken into consideration too. The final score of the matching is calculated by both flow-directional hybrid searches, both of them having an equal contribution. Table 8 presents the increase in performance obtained by using both flow-directional searches.

Table 4

HybridOM results on Bio-ML track 2024

Dataset	Precision	Recall	F-score	MRR	Hits@1	Hits@2	Hits@5	Hits@10
OMIM-ORDO	0.690	0.679	0.685	0.849	0.792	0.874	0.920	0.934
NCIT-DOID	0.924	0.913	0.918	0.952	0.928	0.960	0.981	0.990
SNOMED-FMA	0.870	0.722	0.790	0.907	0.861	0.931	0.964	0.969
SNOMED-NCIT Pharm	0.916	0.889	0.902	0.964	0.936	0.985	0.994	0.995
SNOMED-NCIT Neoplas	0.807	0.710	0.755	0.911	0.870	0.934	0.961	0.962

Table 5

HybridOM search configuration used in the competition. The last search focuses on a generic equally weighted semantic search using field *description_llm_emb* and the keyword search for field *synonyms*.

Semantic			Text		
Src	Dst	Score	Src	Dst	Score
description_emb	description_emb	0.5	label	label ^{0.7} synonyms ^{0.3}	0.3
	description_llm_emb	0.15			
	verbalisation_emb	0.05			
description_llm_emb	description_emb	0.15	label	label ^{0.7} synonyms ^{0.3}	0.3
	description_llm_emb	0.5			
	verbalisation_emb	0.05			
verbalisation_emb	description_emb	0.05	synonyms	label ^{0.7} synonyms ^{0.3}	0.3
	description_llm_emb	0.15			
	verbalisation_emb	0.5			
description_llm_emb	description_emb	0.2	synonyms	label ^{0.7} synonyms ^{0.3}	0.3
	description_llm_emb	0.3			
	verbalisation_emb	0.2			

2. Results

HybridOM is evaluated on all datasets from Bio-ML track for unsupervised tasks (Table 4). For a fair evaluation, each dataset was evaluated using the same configuration. The system can be adapted by changing its parameters per track to obtain better results. The best HybridOM configuration uses BioLord-2023 [7] as an embedding language model, dual hybrid search, and normalization. The assignment of a higher weight is made for the pair *source_field - destination_field* with the same name (Table 5). The weights were chosen empirically, providing a general configuration of weights for each hybrid search, the difference is manifested by which pair receives a higher score (Table 5). For the final configuration, the number of retrieved candidates per search is 10, and only the first 5 candidates found for each concept are used in the backward search for the dual hybrid search process. All evaluations are made only for the concepts that have field *use_in_alignment* set to true.

According to the values for *Hits@K* metric from Table 4, the HybridOM system seems to be struggling to discover the true matching concept between the first two candidates in ranking. For the *Hits@K* and *MRR* metrics a system should provide for each list of possible candidates, a ranking or a score for each candidate. The HybridOM provides a matching ranking but with a variable length of concepts based on the search result. Some candidates used for evaluation of *Hits@K* and *MRR* metrics are not present in the HybridOM final ranking, as a result, the score for them will be 0.0. If the ground truth candidate is not found in the final HybridOM ranking (does not appear on any rank), the concept will have a value of **-1.0** such that we are still able to build the evaluation file. This value is assigned with awareness of the ground truth, but with the scope of not increasing the metrics score; the candidate will have the lowest score among the possible candidates from the file. HybridOM returns a ranking of top candidates for ontology matching, and concepts that are not present in the ranking will be considered automated with a score **0.0** for matching.

The benefits of using BioLord-2023 as the main embedding model, compared to a general domain language model, are outlined in Table 6. The only exception is the dataset *SNOMED-NCIT Pharm*, where the performance is higher using the GTR model.

Table 6Comparison of results: HybridOM with *BioLord-2023* and HybridOM with *GTR-5-Large*

Dataset	BioLord-2023			GTR-5-Large		
	Precision	Recall	F-score	Precision	Recall	F-score
OMIM-ORDO	0.690	0.679	0.685	0.681	0.671	0.676
NCIT-DOID	0.924	0.913	0.918	0.899	0.888	0.893
SNOMED-FMA	0.870	0.722	0.790	0.775	0.633	0.697
SNOMED-NCIT Pharm	0.916	0.889	0.902	0.943	0.915	0.929
SNOMED-NCIT Neoplas	0.807	0.710	0.755	0.738	0.643	0.687

Table 7

Comparison of results: simple and normalized hybrid search

Dataset	Normalization			Simple		
	Precision	Recall	F-score	Precision	Recall	F-score
OMIM-ORDO	0.690	0.679	0.685	0.602	0.594	0.598
NCIT-DOID	0.924	0.913	0.918	0.865	0.854	0.859
SNOMED-FMA	0.870	0.722	0.790	0.809	0.669	0.732
SNOMED-NCIT Pharm	0.916	0.889	0.902	0.888	0.861	0.874
SNOMED-NCIT Neoplas	0.807	0.710	0.755	0.779	0.687	0.730

Table 8

Comparison of results: one-direction hybrid search and dual hybrid search

Dataset	Dual Hybrid Search			One-Direction Hybrid Search		
	Precision	Recall	F-score	Precision	Recall	F-score
OMIM-ORDO	0.690	0.679	0.685	0.654	0.644	0.649
NCIT-DOID	0.924	0.913	0.918	0.905	0.894	0.899
SNOMED-FMA	0.870	0.722	0.790	0.866	0.724	0.789
SNOMED-NCIT Pharm	0.916	0.889	0.902	0.900	0.873	0.886
SNOMED-NCIT Neoplas	0.807	0.710	0.755	0.806	0.713	0.757

Normalization has an important role in hybrid search performance in comparison with a simple approach (Table 7). It increases the score distance between candidates in the local search ranking. The combined local rankings facilitate the creation of a more realistic final matching ranking of candidates, achieving better results for all datasets.

Nevertheless, the backward hybrid search for the retrieved candidate provides a double checking for the initial matching and it increases the performance (Table 8). For some cases, the one-direction hybrid search exceeds the dual hybrid search, but the difference is insignificant.

3. General Comments

3.1. Impact of results

The results obtained of HybridOM in unsupervised Bio-ML track 2024 stand out especially for Recall and F-score metrics. Taking three *1st places* and two *2nd places* for F-score metric and three *1st places*, one *2nd place* and one *3rd place* for Recall metric according to unsupervised track on all datasets.

In comparison with other models, HybridOM has a weaker precision score among all datasets. For the SNOMED-FMA, BERTMap [8] exceeds the HybridOM on the precision metric with a large difference, but when it comes to F-score, both models have the same score taking the first place in ranking. A crucial fact about HybridOM is that the scores of the three metrics (P, R, F-score) in each dataset are very similar and the difference is not large. In the OMIM-ORDO dataset the large discrepancy between precision and recall scores can be seen in competition ranking with a difference of at least *0.15* on all participants, while the HybridOM all three metrics have close scores, with a difference at most of *0.012*.

Using simple semantic and keyword searches, HybridOM achieves consistency among the scores for all metrics (P, R, F-score). The proposed method has a simple yet effective approach of integrating multiple searches such that the limitation of one search is alleviated by the strengths of the others. Integration of LLM for generating description, the usage of the verbalizations from ontology's axioms, and raw data from ontology are the core data used by HybridOM.

3.2. Improvements

HybridOM can be further developed by adopting new methods for extracting or generating data for ontology concepts. Different prompts can be applied along with other LLMs to provide more domain-specific information. The usage of different embedding models can increase the retrieval performance. New ways of combining different kinds of searches besides keyword or semantic-based should offer a large perspective for ontology matching. Enriching the verbalization semantics using other axioms and relations might also be beneficial. An automated way of calculating the search weights based on ontology data, searching fields, or statistics, might be a key improvement for developing a highly adaptable and reusable ontology matching system agnostic to ontology data variance and domain.

4. Conclusion

HybridOM proposes a domain-agnostic solution for ontology matching. It tackles the diversity of ontologies with a diversity of meaning sources, meaning representation, and a combination of hybrid searches. Explicit structured knowledge captured in the ontology is combined with text-based description generated by LLM. The role of LLM is not decisive since no decision is taken based only on LLM, yet, LLM generated descriptions extend the meaning associated to each concept. The impact of wrongly generated descriptions is reduced in HybridOM due to the way hybrid searches are configured.

HybridOM obtained very good and balanced values for precision and recall for all the ontology pairs in the unsupervised 2024 Bio-ML track: for three of them, it obtained the best F1 scores, respectively second best for two of them.

References

- [1] S. Hertling, H. Paulheim, OLaLa: Ontology Matching with Large Language Models, in: Proceedings of the 12th Knowledge Capture Conference 2023, K-CAP '23, ACM, 2023.
- [2] H. B. Giglou, J. D'Souza, F. Engel, S. Auer, Llm4om: Matching ontologies with large language models, 2024. URL: <https://arxiv.org/abs/2404.10317>. arXiv: 2404.10317.
- [3] Y. He, J. Chen, E. Jimenez-Ruiz, H. Dong, I. Horrocks, Language model analysis for ontology subsumption inference, in: A. Rogers, J. Boyd-Graber, N. Okazaki (Eds.), Findings of the Association for Computational Linguistics: ACL 2023, ACL, Toronto, Canada, 2023.
- [4] Y. He, J. Chen, H. Dong, I. Horrocks, C. Allosca, T. Kim, B. Sapkota, DeepOnto: A Python package for ontology engineering with deep learning, Semantic Web 15 (2024) 1991–2004.
- [5] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, G. Lample, LLaMA: Open and Efficient Foundation Language Models, 2023. arXiv: 2302.13971.
- [6] 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: International Semantic Web Conference, Springer, 2022, pp. 575–591.
- [7] F. Remy, K. Demuyne, T. Demeester, BioLORD-2023: semantic textual representations fusing large language models and clinical knowledge graph insights, Journal of the American Medical Informatics Association 31 (2024) 1844–1855. URL: <http://dx.doi.org/10.1093/jamia/ocae029>.
- [8] Y. He, J. Chen, D. Antonyrajah, I. Horrocks, BERTMap: A BERT-based Ontology Alignment System, in: Proceedings of 36th AAAI Conference on Artificial Intelligence 2022 (AAAI 2022), 2022.