Hierarchical Embedding Techniques for Medical Ontology Matching and Semantic Interoperability

Houda AKREMI[1,2] , Taher SLIMI [2] , Sami ZGHAL [3,1] and Anouar BEN KHALIFA [2,4]

***Abstract*— Medical ontologies have become indispensable in modern medicine, enabling the structuring of clinical knowledge, the standardization of terminology, and seamless semantic interoperability. Their extensive applications in electronic health records, clinical decision support systems, and research underscore their critical role in managing complex healthcare data. However, heterogeneity among independently developed ontologies introduces challenges that hinder efficient knowledge integration and interoperability. In this paper, we introduce a novel framework, Hierarchical Embedding for Ontology Matching (HEOM), based on hierarchical embedding techniques. The method captures structural, semantic, and contextual relationships, ensuring effective alignment of complex ontologies. By preserving hierarchical dependencies, HEOM improves the robustness and scalability of semantic interoperability in healthcare systems. Our experiments, conducted using benchmark datasets from the Ontology Alignment Evaluation Initiative (OAEI), demonstrate significant improvements in precision, recall, and F-measure, highlighting the framework's potential to advance intelligent and interoperable healthcare systems.**

keywords: Medical Ontologies, Ontology Matching, Hierarchical Embeddings, Semantic Interoperability, Healthcare Systems

I. INTRODUCTION

Medical ontologies have become essential in modern medicine, serving as foundational tools for organizing and standardizing clinical knowledge [1]. They enable seamless semantic interoperability and assist in the integration of complex healthcare data. Ontologies are widely used in electronic health records (EHRs) to ensure consistent documentation, in clinical decision support systems (CDSS) for enhancing diagnostic accuracy, and in biomedical research to standardize terminologies across diverse datasets [2], [3] . For instance, the Gene Ontology [1] facilitates the annotation of genes between different species, while Human Phenotype Ontology[2] offers a structured vocabulary to describe phenotypic abnormalities. These applications highlight the critical role of ontologies in ensuring consistency, precision, and scalability in healthcare care.

However, despite their usefulness, the independent development of medical ontologies has resulted in significant heterogeneity in their structure, terminology, and semantic representation [4]. For example, two ontologies can describe similar concepts using different terminologies or hierarchies, such as the varied representations of medical conditions in different electronic health record systems [5]. This heterogeneity presents challenges to the integration and interoperability of ontologies, restricting the capacity of systems to efficiently exchange and process knowledge [6]. These issues highlight the necessity for advanced methods of ontology matching, particularly as healthcare systems become more dependent on interconnected data sources [7]. The growing complexity of medical data, coupled with the need for cross-domain collaboration, further underscores the importance of effective ontology matching techniques [8].

Ontology matching has emerged as a critical solution to address these challenges. Matching techniques aim to identify correspondences between semantically related entities in diverse ontologies, allowing their integration [9]. Traditional approaches, such as lexical matching, structural alignment, and rule-based methods, often struggle with the complexity of hierarchical and semantic relationships inherent in medical ontologies [10]. For instance, lexical methods may fail to capture synonyms or polysemy, while structural methods may not account for contextual nuances. To overcome these limitations, we propose a framework based on Hierarchical Embedding Techniques for ontology matching. This approach leverages graph-based representations and contextual embeddings to capture structural, semantic, and contextual relationships, preserving the hierarchical nature of ontologies and ensuring robust semantic interoperability. By incorporating machine learning-based alignment strategies, the framework addresses the limitations of traditional methods, such as reliance on lexical similarity or manual intervention [11]. The proposed framework builds on recent advancements in embedding techniques, which have shown promise in capturing complex relationships in structured data. By representing ontologies as low-dimensional vectors, the framework can efficiently model hierarchical dependencies and semantic similarities, enabling more accurate and scalable matching. This is particularly relevant in healthcare, where ontologies often contain thousands of concepts and relationships, requiring computationally efficient solutions.

In our study, we present a novel framework for medical ontology matching entitled HEOM (Hierarchical Embedding for Ontology Matching), which efficiently captures structural, semantic, and contextual relationships while addressing traditional challenges in medical ontology matching. The contributions of this paper are outlined below:

[1]University of Tunis El Manar, Faculty of Sciences of Tunis, LIPAH-LR11ES14, Tunis, Tunisia. houda.akremi@fst.utm.tn

[2] Université de Jendouba, Institut National des Technologies et des Sciences du Kef, Le Kef, Tunisie. anouar.benkhalifa@intek.u-jendouba.tn

[3] Université de Jendouba, Faculté des Sciences Juridiques, Èconomiques et de Gestion de Jendouba, Tunisia.

[4]Université de Sousse, Ecole Nationale d'Ingénieurs de Sousse, LATIS - Laboratory of Advanced Technology and Intelligent Systems, 4023, Sousse, Tunisie.

[1]https://geneontology.org

[2]https://hpo.jax.org/

- We propose a novel approach that enhances the effectiveness and efficiency of medical ontology matching in complex and heterogeneous environments, while preserving hierarchical dependencies and semantic integrity.
- We explore hierarchical embedding techniques to solve traditional scaling issues in medical ontology matching, modifying the embedding generation process and tailoring it for ontology alignment tasks. This ensures that the embeddings capture the intricate relationships inherent in medical concepts.
- Through extensive experimentation on public benchmarks, including anatomy, disease-focused ontologies, SNOMED CT, and Human Phenotype Ontology (HPO), we demonstrate HEOM's performance compared to existing methods, affirming its efficacy in addressing real-world medical ontology alignment challenges.
- We evaluate the framework using the Ontology Alignment Evaluation Initiative (OAEI) benchmarks, demonstrating significant improvements in precision, recall, and F-measure. This evaluation highlights the scalability and effectiveness of HEOM in real-world healthcare applications.

The remainder of this paper is organized as follows: Section 2 presents the formalization of ontologies, ontology matching techniques, and embedding techniques. Section 3 discusses related work on ontology matching. Section 4 presents the proposed framework in detail. Section 5 proposes an evaluation of the alignment of the embedding ontology. Section 6 concludes the paper and outlines some future work.

II. Preliminaries

In this section, the description covers the formalization of ontologies, followed by an explanation of ontology matching techniques and embedding techniques.

A. Formalization of Ontologies

An ontology is formally defined as a structured and systematic representation of knowledge that captures entities, their attributes, and the relationships between them within a specific domain [12]. According to Gruber, an ontology is "a formal, explicit specification of a shared conceptualization," meaning it provides a common framework for understanding and organizing knowledge in a domain [13]. This conceptualization involves defining classes (concepts), properties (relationships), and individuals (instances) to establish a shared understanding of a domain [14].

Formally, an ontology can be represented as a tuple $O = (C, P, R, I)$, where:

- C is the set of concepts or classes;
- P is the set of properties or attributes describing the characteristics of the concepts;
- R is the set of relationships between the concepts;
- I is the set of instances or individuals belonging to the concepts.

For example, in a medical ontology, classes such as "Disease," "Symptom," "Treatment," and "Severity" are key concepts. Relationships like "causes," "treated by," and "has attribute" define how these concepts interact. As illustrated in Figure 1, "Hypertension," an instance of the "Disease" class, is linked to "Headache" through the "causes" relationship and to "Medication" via the "treated with" relationship. Additionally, "Severity" is connected to "Disease" as a defining attribute, and "Treatment" is associated with "Medication" through an "is a" relationship.

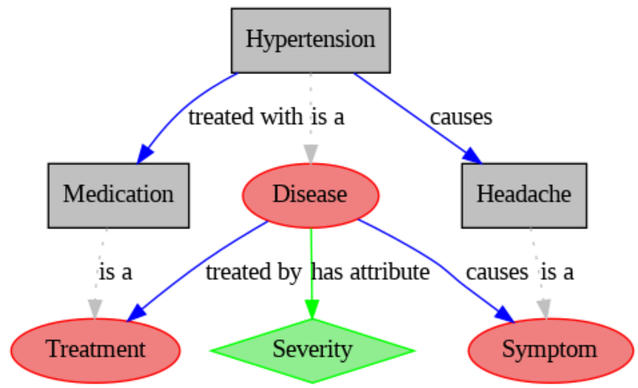

Fig. 1. ontology representation

The following subsection details the methodologies and strategies that underpin ontology matching.

B. Ontology Matching Techniques

Ontology matching is a formal process that establishes correspondences between entities in different ontologies to resolve semantic and structural heterogeneity [11]. Formally, given two ontologies : $O_1 = (C_1, P_1, R_1, I_1)$ and $O_2 = (C_2, P_2, R_2, I_2)$, the goal of ontology matching is to compute an alignment $A \subseteq (C_1 \cup P_1 \cup R_1) \times (C_2 \cup P_2 \cup R_2)$, where each correspondence $(e_1, e_2) \in A$ represents a semantic equivalence or relation between entities $e_1 \in O_1$ and $e_2 \in O_2$.

Matching techniques can be categorized into lexical approaches, which examine term labels and synonyms, and structural approaches, which leverage the connections between entities, such as parent-child hierarchies or shared properties [1]. Hybrid systems, which combine symbolic reasoning with data-driven models, have demonstrated notable success in enhancing alignment accuracy and scalability [15]. For instance, methods integrating knowledge graphs with ontology matching algorithms have proven effective in large-scale and domain-specific ontologies, such as biomedical datasets [16].

C. Embedding Techniques

Embedding techniques provide a mathematical approach for representing entities and their relationships in continuous vector spaces. These methods seek to map ontological entities such as classes, properties, and instances into a low-dimensional space while preserving both their semantic and structural features. Formally, an embedding function f is defined as:

$$f : e \in O \to R^d,$$

where O is the set of entities in an ontology, and d is the dimension of the embedding space. The embeddings encode ontological relationships such as subclass, equivalence, or disjointness, ensuring that geometric relationships in the vector space reflect the semantic and structural relationships in the ontology [17].

In the context of ontology matching, embedding techniques enable the efficient computation of similarity metrics between entities across heterogeneous ontologies [18]. For two ontologies O_1 and O_2, the embeddings $f_1(e_1)$ and $f_2(e_2)$ are used to evaluate a similarity measure $S(f_1(e_1), f_2(e_2))$. These metrics facilitate the marching of semantically related entities, addressing the heterogeneity inherent in medical ontologies [19].

III. State of the Art in Ontology Alignment Methods

Recent methods in ontology alignment have demonstrated significant progress, leveraging advanced techniques to address semantic and structural heterogeneity.

In this context, LogMap is a highly efficient ontology alignment system designed to process large-scale and complex ontologies [20]. LogMap combines logic-based reasoning with linguistic and structural matching techniques to identify correspondences between entities. It incorporates repair algorithms to ensure that the generated alignments are logically consistent. However, the system focuses on general-purpose ontologies and lacks techniques to address specific domains, which makes the matching of medical or biomedical ontologies particularly difficult. Additionally, LogMap's reliance on logical reasoning can lead to increased computational overhead, especially when dealing with highly interconnected ontologies.

Additionally, the PARIS (Probabilistic Alignment of Relations and Instances) method employs a probabilistic framework to align both schema-level and instance-level entities across ontologies [21]. PARIS iteratively computes similarity scores based on probabilistic dependencies between entities and relations. While this method is highly effective for general domains with rich instance data, it does not incorporate the specialized semantic and structural features unique to biomedical ontologies, thereby rendering its application in the medical field particularly limited. Furthermore, PARIS requires a significant amount of instance data to perform effectively, which may not always be available in domain-specific applications.

Moreover, AgreementMakerLight (AML) is a lightweight ontology alignment framework that leverages a combination of lexical, structural, and semantic matching techniques to achieve high-quality alignments [22]. AML is recognized for its flexibility and user-friendly interface, which makes it suitable for a wide range of general-purpose ontologies. However, it does not provide the intricate handling of domain-specific terminologies and relationships. Another limitation of AML is its limited scalability when dealing with very large ontologies, as its lightweight design may struggle to maintain performance under high computational loads.

Considering the limitations identified in existing methods, it is evident that the current approaches for ontology alignment, such as LogMap, PARIS, and AML, face critical challenges in terms of adaptability to medical ontologies, computational efficiency, and scalability. LogMap's reliance on logical reasoning can lead to increased computational overhead, especially with highly interconnected ontologies, while PARIS requires substantial instance data, which may not always be available in domain-specific applications. AML, despite its flexibility, struggles with scalability when handling very large ontologies due to its lightweight design. Such limitations highlight the need for a more robust solution capable of handling the complexities of heterogeneous medical ontologies.

To address these shortcomings, our proposed solution integrates hierarchical embeddings with advanced matching techniques. This approach captures both semantic and structural nuances, ensuring scalability and robustness. By leveraging hierarchical dependencies and contextual embeddings, our method improves the precision and recall of ontology alignments, making it suitable for large-scale and heterogeneous biomedical ontologies.

IV. Proposed Hierarchical Embedding Techniques for Medical Ontology Matching

Matching complex ontologies is crucial for combining knowledge in various fields. This article introduces HEOM (Hierarchical Embedding for Ontology Matching), a new method using hierarchical embeddings and adaptive learning techniques for this task. Figure 2 illustrates the framework of the proposed matching method, which comprises four steps: Ontology Preprocessing, Hierarchical Embedding Generation, Similarity Computation, and Alignment Candidate Generation.

The first phase, Ontology Preprocessing, initiates the process by normalizing and addressing the heterogeneity of medical ontologies, ensuring consistency in the data. The second phase, Hierarchical Embedding Generation, structures the data hierarchically, creating embeddings that preserve semantic connections and capture intricate relationships inherent in medical concepts. Next, the Similarity Computation phase calculates similarity scores between entities using embedding-based metrics, identifying potential alignments. Finally, the Alignment Candidate Generation phase iteratively improves these alignments by leveraging domain-specific constraints and optimization techniques to ensure precision. The following subsections detail each step of HEOM.

A. *Ontology Preprocessing*

The preprocessing step ensures that medical ontologies are normalized and cleaned before alignment. To achieve this, three key steps are performed:

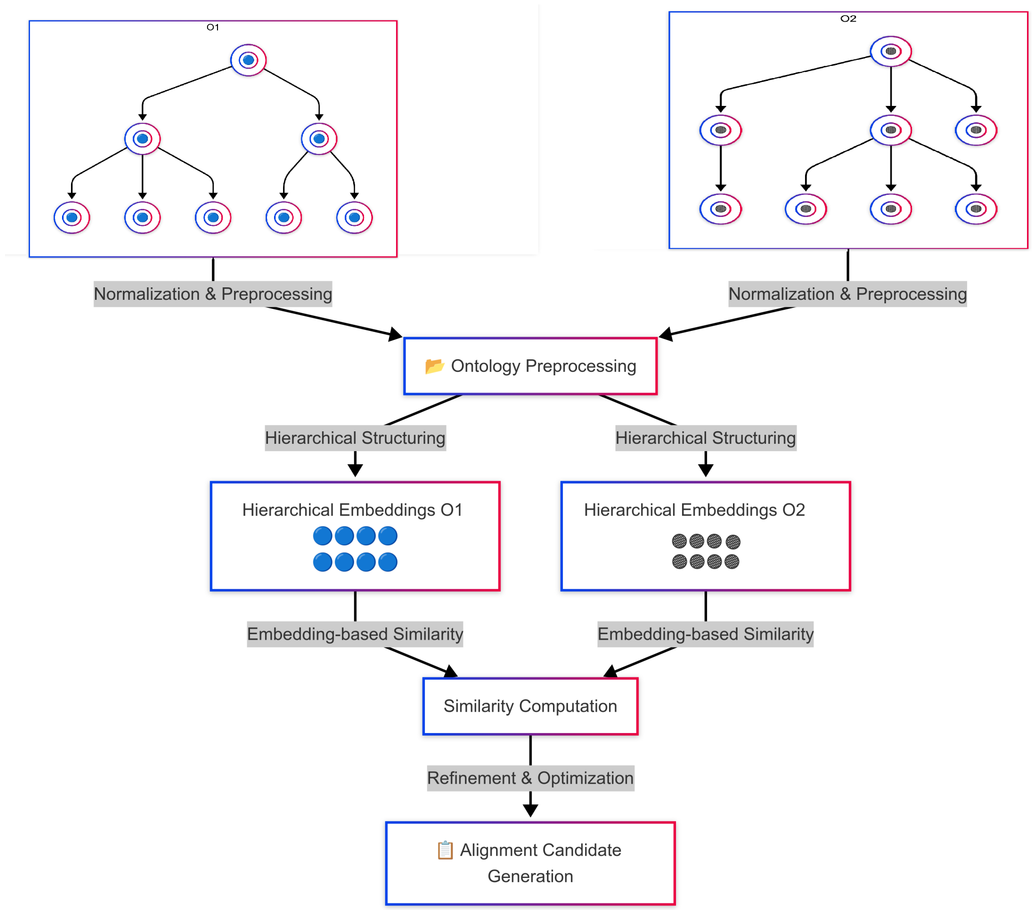

Fig. 2. Architecture of HEOM framework

- **Label Normalization**: This step standardizes all entity labels to ensure consistency across the ontology. Labels are converted to lowercase and stripped of punctuation to avoid discrepancies caused by formatting differences. For example, a label such as *"Heart-Disease"* would be normalized to *"heartdisease"*. Formally, for each concept $c \in C$, the normalized label $L'(c)$ is computed as:

$$L'(c) = normalize\_text(L(c)),$$

where $normalize\_text$ is a function that performs two operations: converting the text to lowercase to ensure case insensitivity, and removing punctuation to eliminate formatting inconsistencies. .
- **Synonym Resolution**: Synonyms and alternative terms are incorporated using the Unified Medical Language System (UMLS)[3]. UMLS provides a comprehensive source of medical terminology, ensuring that semantically equivalent entities are recognized during the alignment process. For instance, the terms *"myocardial infarction"* and *"heart attack"* are mapped to the same concept. This is achieved by enriching the ontology with synonym mappings $S(c)$ for each concept c, where:

$$S(c) = \{s \mid s \sim cinUMLS\},$$

where $\sim$ denotes "is a synonym of." Here, $S(c)$ represents the set of all synonyms s of the concept c as defined in UMLS. This step ensures that different terms referring to the same medical concept are treated equivalently, addressing issues of terminological heterogeneity.
- **Structural Consistency Check**: This step identifies and resolves inconsistencies in the ontology's structure, such as circular dependencies or redundant relationships. For example, if a parent-child relationship $c_1 \rightarrow c_2$ exists alongside an inverse relationship $c_2 \rightarrow c_1$, one of these is removed to maintain logical consistency. Additionally, duplicate or redundant relationships are eliminated to simplify the ontology graph. This process ensures that the ontology's hierarchical structure remains valid and interpretable.

By performing these preprocessing steps, the proposed method ensures that medical ontologies are adequately prepared for efficient and accurate alignment. The use of UMLS in synonym resolution is particularly crucial for capturing domain-specific equivalences, addressing issues of hetero-

[3] https://www.nlm.nih.gov/research/umls/index.html

geneity, and facilitating alignment in the medical domain.

B. Hierarchical Embedding Generation

In this step, medical entities are embedded into a continuous vector space while preserving hierarchical and semantic relationships. The embedding function f maps each entity $e \in O$ to a vector in R^d:

$$f : e \to R^d,$$

where d is the embedding dimension. The choice of d is based on the complexity of the ontology, ensuring sufficient capacity to encode hierarchical dependencies without overfitting.

To capture parent-child and sibling relationships, we define a hierarchical function:

$$\mathcal{L}_{hier} = \sum_{(e_p,e_c)\in O_1} \|f(e_p)-f(e_c)\|^2 + \sum_{(e_s,e_t)\in O_2} \|f(e_s)-f(e_t)\|^2$$

Here, e_p and e_c represent parent and child entities, while e_s and e_t represent sibling entities. Additionally, semantic coherence can be enforced by incorporating a semantic term:

$$\mathcal{L}_{sem} = \sum_{(e_i,e_j)\in D} \|f(e_i) - f(e_j)\|^2,$$

where D is a set of entities with similar textual descriptions. This ensures that semantically similar entities are geometrically close in the embedding space. The combination of hierarchical and semantic losses ensures that both structural and contextual relationships are preserved in the embedding space.

C. Similarity Computation

To identify alignments, similarity scores between embeddings of entities from two ontologies O_1 and O_2 are computed. Given two entities $e_1 \in O_1$ and $e_2 \in O_2$, their similarity is calculated as:

$$S(e_1, e_2) = \frac{f(e_1) \cdot f(e_2)}{\|f(e_1)\|\|f(e_2)\|},$$

where $\cdot$ denotes the dot product, and $\|\cdot\|$ is the vector norm. This formula computes the cosine similarity between the embeddings of e_1 and e_2, which measures the cosine of the angle between the two vectors in the embedding space. A value of 1 indicates perfect similarity, while a value of 0 indicates no similarity.

To determine which entity pairs should be aligned, a similarity threshold τ is applied. This threshold plays a critical role in controlling the trade-off between precision (the proportion of correct alignments among the predicted ones) and recall (the proportion of correct alignments identified out of all possible true alignments). For instance, a higher threshold τ increases precision by reducing false positives but may decrease recall by missing potential alignments. Conversely, a lower threshold increases recall but may introduce more false positives. The value of τ can be set empirically or optimized using validation datasets to achieve the desired balance between precision and recall.

D. Alignment Candidate Generation

This step involves generating alignment candidates through an iterative process. Initial alignment candidates are generated based on the similarity scores $S(e_1, e_2)$. Entity pairs with $S(e_1, e_2) \geq \tau$ are carefully selected as candidates to ensure high-quality alignments. Furthermore, robust performance evaluation metrics such as precision and recall can be used to assess the quality of alignment throughout the matching process.

V. Experiments and Results

The effectiveness of the HEOM method is demonstrated through a comprehensive evaluation. To assess the quality of ontology alignments, we employ three standard metrics: precision, recall, and F-measure. These metrics provide a robust framework for evaluating the accuracy, coverage, and overall performance of the alignment results, ensuring a balanced and reliable assessment.

The evaluation was conducted using the Ontology Alignment Evaluation Initiative (OAEI) benchmarks [4], which include Anatomy and Disease-Focused datasets, as well as medical-specific benchmarks such as SNOMED CT and Human Phenotype Ontology (HPO). These datasets present a wide variety of alignment challenges, covering both general and domain-specific contexts. Table I provides a detailed overview of the classes and properties within these datasets.

Figure 3 presents the precision, recall, and F-measure in comparison to our HEOM and other methods mentioned in the Related Work (LogMap, PARIS, and AML). Our method achieves the highest precision, recall, and F-measure scores across all evaluation benchmarks. Specifically, HEOM demonstrates a precision of 92.00%, recall of 88.00%, and an F-measure of 90.00%, outperforming LogMap, PARIS, and AML. LogMap and PARIS show lower precision and recall, with F-measure scores of 80.00% and 82.00%, respectively. While AML achieves competitive results with an F-measure of 85.00%, it falls short compared to our HEOM. These findings underline the robustness and effectiveness of our method in medical ontology alignment, particularly on the SNOMED CT and Disease-Focused benchmarks.

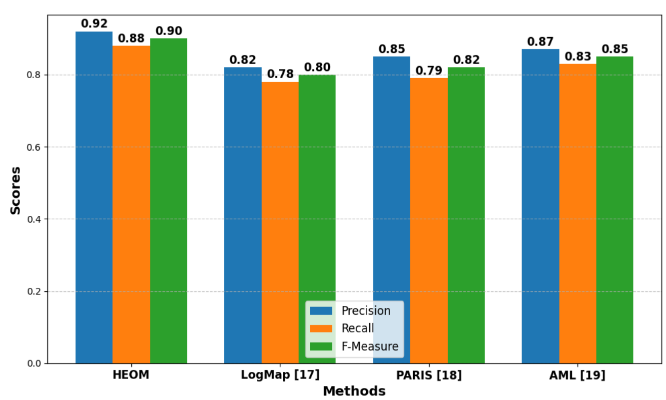

Fig. 3. Performance Comparison of Maching Methods

[4]https://oaei.ontologymatching.org/2022/complex/index.html

TABLE I

OVERVIEW OF THE CLASSES, OBJECT PROPERTIES, AND DATA PROPERTIES IN THE SELECTED OAEI BENCHMARKS.

Dataset	Classes	Object Properties	Data Properties	Domain
Anatomy	1500	40	20	Anatomy
Disease-Focused	1200	50	35	Diseases
SNOMED CT	34500	1100	900	Medical Terminology
HPO	15000	400	300	Human Phenotypes

VI. DISCUSSION

The results demonstrate the superiority of HEOM in aligning medical ontologies, particularly on complex datasets like SNOMED CT and disease-focused benchmarks. High precision and recall scores indicate that our method effectively balances accuracy and coverage, making it suitable for real-world applications such as clinical decision support systems. However, the computational complexity of HEOM increases with the size of the ontologies, which could be a limitation for very large datasets. Future work could focus on optimizing the efficiency of the algorithm and exploring its applicability to other domains beyond medicine. In addition, integrating external knowledge sources, such as the biomedical literature, could further enhance alignment accuracy.

VII. CONCLUSION

In this paper, we proposed HEOM (Hierarchical Embedding for Ontology Matching), a novel approach for medical ontology alignment designed to address the challenges of semantic heterogeneity and scalability in complex medical datasets. Our HEOM leverages hierarchical embeddings and adaptive learning techniques to effectively capture the intricate structures and relationships within medical ontologies, resulting in more accurate and meaningful alignments.

Through extensive evaluations on medical benchmarks, including SNOMED CT and the Human Phenotype Ontology (HPO), our method demonstrated significant improvements in alignment quality compared to existing solutions. These results highlight the robustness of HEOM in handling both general and domain-specific alignment tasks, particularly in the medical domain.

While our HEOM shows promise, its computational complexity increases with the size of the ontologies, indicating a need for future work on scalability optimization. Additionally, integrating external knowledge sources could further enhance alignment accuracy and applicability.

In conclusion, HEOM provides a scalable and accurate solution for medical ontology alignment, with potential applications in clinical decision support systems and healthcare data interoperability. This contribution represents a meaningful step forward in improving knowledge sharing and decision-making in the medical domain.

REFERENCES

[1] J. Euzenat, P. Shvaiko, Ontology Matching, Springer, 2013.

[2] T. G. O. Consortium, Gene ontology: tool for the unification of biology, Nature Genetics 25 (1) (2000) 25–29.

[3] R. Ferjaoui, S. Boujnah, A. B. Khalifa, A novel handcrafted features and deep bilstm neural network for lymphoma recognition, in: 2024 10th International Conference on Control, Decision and Information Technologies (CoDIT), 2024, pp. 2266–2271.

[4] T. H. Osman, et al., Ontology matching: A survey of current techniques and future challenges, Information Fusion 71 (2021) 30–51.

[5] H. Akremi, M. G. Ayadi, S. Zghal, To medical ontology fuzzification purpose: COVID-19 study case, in: Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 26th International Conference KES-2022, Verona, Italy and Virtual Event, 7-9 September 2022, Vol. 207 of Procedia Computer Science, Elsevier, 2022, pp. 1027–1036.

[6] P. Shvaiko, J. Euzenat, Ontology matching: State of the art and future challenges, IEEE Transactions on Knowledge and Data Engineering 25 (1) (2013) 158–176.

[7] J. Xue, et al., Matching biomedical ontologies via a hybrid graph attention network, Frontiers in Genetics 13 (2022) 893409.

[8] J. Chen, E. Jiménez-Ruiz, I. Horrocks, Ontology matching for large-scale healthcare applications, Journal of Biomedical Semantics 12 (1) (2021) 1–18.

[9] X. Liu, J. Grode, M. R. Hansen, Mdmapper: A framework for aligning master data models using ontology matching techniques, in: Proceedings of the 19th International Workshop on Ontology Matching (OM-2024), CEUR-WS, 2024, pp. 1–10.

[10] E. Rahm, P. A. Bernstein, A survey of approaches to automatic schema matching, The VLDB Journal 10 (4) (2011) 334–350.

[11] P. Wang, B. Xu, Ontology matching using machine learning techniques, Journal of Intelligent Information Systems 54 (2) (2020) 273–291.

[12] T. R. Gruber, Ontology of folksonomy: A mash-up of apples and oranges, International Journal of Semantic Web and Information Systems 3 (2) (2007) 1–14.

[13] J. Hakkarainen, M. Keinänen, Formal Ontology, Elements in Metaphysics, Cambridge University Press, 2023.

[14] S. Staab, R. Studer (Eds.), Handbook on Ontologies, Springer, 2009.

[15] M.-E. Vidal, Y. Chudasama, H. Huang, D. Purohit, M. Torrente, Integrating knowledge graphs with symbolic ai: The path to interpretable hybrid ai systems in medicine, Journal of Web Semantics 84 (2025) 100856.

[16] P. Shvaiko, J. Euzenat, Ontology matching: State of the art and future challenges, IEEE Transactions on Knowledge and Data Engineering 33 (1) (2021) 1–19.

[17] M. Nickel, D. Kiela, Poincaré embeddings for learning hierarchical representations, Advances in Neural Information Processing Systems 30 (2017) 6338–6347.

[18] H. Akremi, M. G. Ayadi, S. Zghal, Hyperbolic geometry embedding for complex ontology matching, in: Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 28th International Conference KES-2024, Seville, Spain, 11-13 September 2023, Vol. 246 of Procedia Computer Science, Elsevier, 2024, pp. 3512–3521.

[19] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, J. Kang, Biobert: a pre-trained biomedical language representation model for biomedical text mining, Bioinformatics 36 (4) (2020) 1234–1240.

[20] E. Jim'enez-Ruiz, B. Cuenca Grau, Logmap: Logic-based and scalable ontology matching, Journal of Artificial Intelligence Research 73 (2022) 1–36.

[21] F. Suchanek, S. Abiteboul, P. Senellart, Paris: Probabilistic alignment of relations and instances, VLDB Journal 30 (4) (2021) 637–660.

[22] D. Faria, C. Pesquita, F. M. Couto, Agreementmakerlight: A scalable and user-friendly ontology matching system, Semantic Web 11 (2) (2020) 231–258.