

Language Agnostic Ontology Extension Framework

Abstract—Ontologies are powerful structures used to define the schema and organization of knowledge. They provide a framework for precisely and explicitly organizing concepts, entities, attributes, and relationships within a domain, enabling effective data management, knowledge sharing, and intelligent decision-making. Knowledge graphs, on the other hand, store data in a graphical form, facilitating semantically rich and interconnected data analysis, management, and exploration. Nodes represent entities, edges depict relationships between nodes, and attributes showcase node properties. Combining the strengths of ontology, which offers schema and concept-level modeling, with the data richness of knowledge graphs, one can unlock powerful reasoning and inference capabilities that closely resemble facts and truths. The symbiotic relationship between knowledge graphs and ontology, wherein ontology serves as a blueprint for representing knowledge in the graphs, creates a robust foundation for knowledge representation. Traditionally, ontologies have been created by experts or skilled teams with domain knowledge, aiming to build comprehensive ontologies supporting holistic data representation. However, these domain-specific ontologies require manual intervention for updates and remain language-specific, making it challenging to transfer them across different language formats. To address this challenge, a crucial component is an automatic generation framework for language-independent ontologies. Such a capability would provide a data-driven approach for creating necessary ontologies. The process involves taking a text corpus as input and subsequently constructing a knowledge graph through named entity recognition. This graph is then transformed into a concept-level modeling graph, where an ID-based approach is utilized to represent the concepts. The ID-based concepts can be extended to facilitate merging and updating the ontology with new relevant information. Concurrently, this approach can also be expanded to achieve language independence, enabling the conversion of the ontology to any required language. Expressing these ontologies in mathematical terms is essential to achieve language-agnosticism, ensuring completeness, relevance, and independence from specific domains, thereby making them applicable across various linguistic contexts.

keywords—Ontology, Knowledge graphs, Named entity recognition, RDF, language-agnostic

I. INTRODUCTION

Ontologies have proven to be an effective and robust schema framework for knowledge representation. This has, in turn, influenced the schema of knowledge graphs, providing them with a powerful blueprint that specifies entities and their relations. Knowledge graphs, as graphical data structures, excel in storing and representing information through interconnected nodes and edges, effectively representing real-world

entities and their relationships. The intricate interconnections offered by knowledge graphs enable efficient data analysis and facilitate the discovery of previously unknown relationships, patterns, and insights. The data stored in knowledge graphs can originate from structured data sources such as datasets or can be extracted from unstructured text. The latter necessitates a considerable amount of preprocessing and natural language processing to extract knowledge. In the scope of this paper, we concentrate on the generation of knowledge graphs from raw text corpora.

Raw text data is diverse and consists of unorganized strings of characters, lacking discernible structure or semantic meaning. Therefore, to extract meaningful information from raw text input, advanced natural language processing techniques and intelligent knowledge acquisition methodologies are essential. The foundational structure of knowledge graphs relies on triples, which comprise subject-predicate-object entities and form the fundamental unit in knowledge graph construction. By transforming text input into these triples through the use of pretrained pipelines and large language models, and performing named-entity recognition, one can establish the groundwork for supporting advanced reasoning, inference, and data exploration. This, in turn, aids in the process of ontology generation.

Once the knowledge graph representation is complete, the next steps involve labeling the entities representing subjects and objects. The objective is to generalize the nodes and extract essential relations for concept-level modeling of the data, leading to ontology generation. This process can be facilitated through mechanisms like text annotation, which allows for a data-driven approach to ontology generation with human intelligence intervention.

The resulting ontology, however, remains language-specific. To achieve language-agnosticism, the next step involves representing the entities and relationships in the ontology using specific mathematical representations. These mathematical representations are derived from the core ontology generated in the language of the input text corpus, essentially capturing the concepts in a mathematical form. Subsequently, these mathematical representations are mapped to their corresponding meanings in different languages using various translation APIs like Google Translate. As a result, the base ontology is stored in a mathematical format, containing entities and relations represented in mathematical representations. This

2) *Language agnostic ontology generation*: The ontology created so far is specific to the language of the text corpus and, as a result, lacks language independence. To achieve language agnosticism, we can represent the ontology in a mathematical format that can be easily translated into different languages. Each mathematical representation corresponds to a concept, which can be converted into various languages using a translation API, such as Google Translation API.

To begin, we assign unique numbers to represent each concept in the ontology, ensuring that these numbers are mapped back to the original concepts. By doing so, we can translate these numbers into any desired language by invoking a custom-designed function that accepts a language code for the target language. The translated results can then be stored as attributes of nodes in the graph, reducing time complexity by avoiding repeated API calls for already translated language ontologies.




Fig. 5: Language agnostic ontology for sample 1 and sample 2 respectively

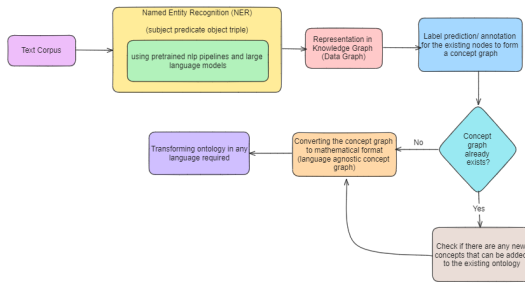


Fig. 6: Flow diagram of language agnostic ontology generation

3) *Updation of existing ontology*: As the world around us is constantly transforming, there is a need for continuous updating in how we structurally represent knowledge. This is particularly relevant for existing ontologies, as new text corpora may contain information that current ontologies are unable to represent, causing them to fall short. Hence, ensuring that ontologies remain fresh and up-to-date becomes essential to better represent recent and new knowledge. In this context, it involves comparing the already existing ontology with the newly created one to see if there is something that can be added to the original one. This is done by taking the edges and documents stored in ArangoDB for both the vertices and

performing the join operation to add the new nodes and edges in the existing ontology. Thereafter, a new updated ontology is formed. This join operation is performed using libraries like pandas and AQL queries to retrieve data in our case. For instance, let's consider the two graphs depicted in Figure 3 and Figure 4. Since they are not the same, the ontologies are compared, and a new graph, shown in Figure 5, is formed, incorporating the existing ontologies. By employing graph and join queries in updating ontologies, we can ensure that these knowledge representations remain relevant, accurate, and comprehensive in the face of evolving information and text corpora.




Fig. 7: Updated ontology with sample text 1 (existing ontology) and sample text 2 (new ontology)

B. Mathematical Formulation

We have modeled the ontologies as a set of triplets (subject, predicate, object). Let's say we already have ontology A in our system and we have learned about ontology B based on the new corpus, in this section we merge these two ontologies to create a final ontology. The merging is done using the ID-based approach depicted in Algorithm 1. This approach basically assigns a unique ID to concepts that exist in the ontology in turn generating a mathematically represented ontology.

The next step involves merging the ontologies represented in terms of IDs with is done as explained in Algorithm 2.

Ontology 1 (A):

$$\begin{aligned}
 A_1^s, A_1^P, A_1^O &\rightarrow A_1^I \\
 A_2^s, A_2^P, A_2^O &\rightarrow A_2^I \\
 &\vdots \\
 &\vdots \\
 &\vdots \\
 A_n^s, A_n^P, A_n^O &\rightarrow A_n^I
 \end{aligned}$$

Where A_x^s, A_x^P, A_x^O and A_x^I represent the Subject, Predicate, Object and Instance (Triple) in Ontology A

Ontology 2 (B):

$$\begin{aligned}
 B_1^s, B_1^P, B_1^O &\rightarrow B_1^I \\
 B_2^s, B_2^P, B_2^O &\rightarrow B_2^I \\
 &\vdots \\
 &\vdots \\
 &\vdots \\
 B_n^s, B_n^P, B_n^O &\rightarrow B_n^I
 \end{aligned}$$

Where B_x^s , B_x^P , B_x^O and B_x^I represent the Subject, Predicate, Object and Instance (Triple) in Ontology B

$$S_A = \{ A_1^I, A_2^I, \dots, A_n^I \}$$

$$S_B = \{ B_1^I, B_2^I, \dots, B_n^I \}$$

where S_A and S_B are set of instances from ontology A and B respectively.

Let the updated graph be C

$$C = S_A \cup S_B$$

Algorithm 1 Create ID based Ontology (Ontology O)

- 1: Initialize an empty Graph O_{gen} .
 - 2: **for** each node n in graph O **do**
 - 3: $ID \leftarrow$ IdentityService(n) {Call IdentityService to get the ID for the node.}
 - 4: Add node with ID ID to graph O_{gen} .
 - 5: **end for**
 - 6: **for** each edge e in graph O **do**
 - 7: $ID \leftarrow$ IdentityService(e) {Call IdentityService to get the ID for the edge.}
 - 8: $source \leftarrow$ GetSourceNodeID(e) {Get the ID of the source node for the edge.}
 - 9: $target \leftarrow$ GetTargetNodeID(e) {Get the ID of the target node for the edge.}
 - 10: Add edge with ID ID , source $source$, and target $target$ to graph O_{gen} .
 - 11: **end for**
 - 12: **Return** Graph O_{gen} containing nodes and edges with their respective IDs. =0
-

Algorithm 2 Merge O_i and O_j

- 1: Initialize an empty map $visited$ to store if a predicate has already been added or not.
 - 2: Initialize an empty Ontology/Graph O .
 - 3: **for** each predicate p in graph O_i **do**
 - 4: Add predicate p to graph O .
 - 5: Mark predicate p as visited in map $visited$.
 - 6: **end for**
 - 7: **for** each predicate p in graph O_j **do**
 - 8: **if** p is not already present in map $visited$ **then**
 - 8: Add edge p to graph O .
 - 8: Mark predicate p as visited in map $visited$.
 - 9: **end if**
 - 10: **end for**
 - 11: **Output:** The graph O representing the merge of O_i and O_j . =0
-

IV. RESULTS and Discussion

The data-driven approach described above enables language agnostic ontology generation using deep learning models. The initial step involves utilizing the input text corpus as

the foundation for building the ontology. By structuring the unstructured data through knowledge graph representations, the system extracts entities and relationships via natural language processing (NLP), specifically named entity recognition and dependency parsing. These extracted entities are then transformed into triples to be represented in a knowledge graph format, as illustrated in Figure 1 and Figure 2, showcasing two distinct text prompts in the knowledge graph format. To ensure the relevancy and factual accuracy of the knowledge graphs, a crucial subsequent step involves cleaning the graphs, eliminating unnecessary nodes and irrelevant information. Following the cleaning process, the nodes are categorized and annotated using annotation models, resulting in the formation of a final ontology. This entire process is automated, deriving the ontology directly from the unstructured data corpus, as depicted in Figure 3 and Figure 4. Once the ontology is established, the next challenge is to update it with new information and extend its scope. For instance, consider Figure 3 as the original ontology and Figure 4 as the new ontology obtained. Since both ontologies have a common node 'company,' it implies the existence of fresh data about companies that can be incorporated into the existing ontology. To achieve this, the system verifies that the two ontology graphs are different. Subsequently, the system appends the extra nodes from the new ontology to the original one, resulting in an updated ontology (Figure 5) that encompasses the newly acquired information. This newly acquired ontology is language agnostic, allowing it to be transformed into any required language using translation APIs. The function made for translating the ontology into different languages takes as the input the language name and the language code.

```
translate_ontology("English", "en")
translate_ontology("Japanese", "ja")
translate_ontology("tamil", "ta")
translate_ontology("hindi", "hi")
```

Figure 8 displays the base ontology mathematically represented and translated into different languages such as English, Hindi, Tamil, and Japanese. This is achieved by providing a mathematical representation for each concept in the original ontology, establishing a framework for generating a language agnostic ontology.

The research successfully addresses the need for an automated mechanism to generate ontologies from unstructured text data without requiring human intervention. By leveraging deep learning models, and NLP techniques, the system autonomously processes raw text data, creates knowledge graphs, annotates nodes, and generates a comprehensive ontology that is language agnostic. Moreover, the system's ability to accommodate new information enhances the ontology's relevance and applicability over time.

V. CONCLUSION

In conclusion, the paper aims to automate the generation of language agnostic ontologies from raw, unstructured text corpora and also update existing ontologies. Ontologies serve

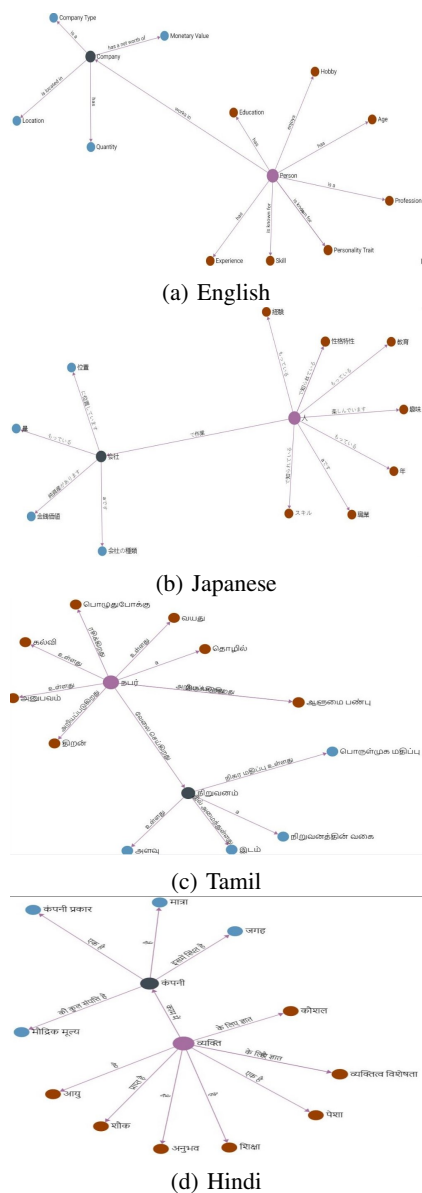


Fig. 8: Output depicting Language Agnostic domain-independent ontology

as powerful schema frameworks, enabling efficient, uniform, and comprehensive knowledge representation. When utilized as blueprints for building knowledge graphs, they enhance the structure, semantics, and utility, making knowledge graphs potent tools for representation, data integration, reasoning, and exploration across diverse domains and applications. However, developing ontologies typically requires manual effort and a skilled team of experts with domain and ontology knowledge. This challenge is addressed in the paper by proposing a data-driven approach to ontology generation. Through natural language processing, the input text containing data is processed to extract concepts and create the ontology graph. This approach ensures robust, efficient, and accurate schema generation without human intervention. Moreover, the created

ontologies can be updated as new text prompts may introduce novel information not present in the existing ontology. This is done by various mechanisms such as graph queries and data manipulation libraries like pandas. To tackle the problem of ontologies being language-specific, mathematical representations of concepts were used to make the ontologies easy to transform into various languages using translation APIs. This allowed the ontology to not only be domain-specific but also provided a way for them to be represented in a language-agnostic way. Overall, the automatic ontology generation reduces the need for manual curation, enhancing the overall robustness and efficiency of the process.

FUTURE WORK

In the future, research aims to process even more complex sentences as text input. These sentences may exhibit a richer, more human-like vocabulary with diverse and vast semantic relationships that current language models may not fully capture or comprehend. To achieve this, the focus lies on selecting models that can process and analyze sentences at a deeper level. These models will be trained on a massive corpus of diverse texts, enabling them to learn intricate relationships between words, grammar, semantics, and pragmatics within a language. This approach allows for scope in terms of validating and filtering sentences effectively. Furthermore, to enhance the accuracy of the approach, a mechanism for error control, thresholding, and validation will be implemented. Reference ontologies and knowledge graphs will play a crucial role in achieving this, as they will aid in cleaning the knowledge graphs and ontologies generated. By incorporating these enhancements, the research aims to create more accurate, precise, and versatile ontologies that are domain-independent and can be automatically generated.

REFERENCES

- [1] Elnagar, S., Yoon, V., Thomas, M. A. (2020). An Automatic Ontology Generation Framework with An Organizational Perspective. Proceedings of the 53rd Hawaii International Conference on System Sciences. Retrieved from <https://doi.org/10.48550/arXiv.2201.05910>
- [2] O. Shatalov and N. Ryabova, "Named Entity Recognition Problem for Long Entities in English Texts," 2021 IEEE 16th International Conference on Computer Sciences and Information Technologies (CSIT), Lviv, Ukraine, 2021, pp. 76-79, doi: 10.1109/CSIT52700.2021.9648768.
- [3] G. Gu, Y. Nam, K. Park, Z. Galil, G. F. Italiano and W. -S. Han, "Efficient Graph Isomorphism Query Processing using Degree Sequences and Color-Label Distributions," 2022 IEEE 38th International Conference on Data Engineering (ICDE), Kuala Lumpur, Malaysia, 2022, pp. 872-884, doi: 10.1109/ICDE53745.2022.00070.
- [4] Y. Chen and X. Xing, "Constructing Dynamic Knowledge Graph Based on Ontology Modeling and Neo4j Graph Database," 2022 5th International Conference on Artificial Intelligence and Big Data (ICAIBD), Chengdu, China, 2022, pp. 522-525, doi: 10.1109/ICAIBD55127.2022.9820199.
- [5] Lingfei Wu; Yu Chen; Kai Shen; Xiaojie Guo; Hanning Gao; Shucheng Li; Jian Pei; Bo Long, Graph Neural Networks for Natural Language Processing: A Survey , now, 2023.
- [6] A. H. Osman and O. M. Barukub, "Graph-Based Text Representation and Matching: A Review of the State of the Art and Future Challenges," in IEEE Access, vol. 8, pp. 87562-87583, 2020, doi: 10.1109/ACCESS.2020.2993191.

- [7] J. Jetschni and V. G. Meister, "Schema engineering for enterprise knowledge graphs: A reflecting survey and case study," 2017 Eighth International Conference on Intelligent Computing and Information Systems (ICICIS), Cairo, Egypt, 2017, pp. 271-277, doi: 10.1109/IN-TELCIS.2017.8260074.