
A SHORT REVIEW FOR ONTOLOGY LEARNING FROM TEXT: STRIDE FROM SHALLOW LEARNING, DEEP LEARNING TO LARGE LANGUAGE MODELS TREND

Rick Du, Huilong An, Keyu Wang
BSH Home Appliances Holding (China) Co., Ltd
{Rick.Du, Huilong.An, Keyu.Wang}@bshg.com

Weidong Liu
Department of Computer Science and Technology, Tsinghua University
liuwd@tsinghua.edu.cn

ABSTRACT

Ontologies provide formal representation of knowledge shared within Semantic Web applications and Ontology learning from text involves the construction of ontologies from a given corpus of text. In the past years, ontology learning has traversed through shallow learning and deep learning methodologies, each offering distinct advantages and limitations in the quest for knowledge extraction and representation. A new trend of these approaches is relying on large language models to enhance ontology learning. This paper gives a review in approaches and challenges of ontology learning. It analyzes the methodologies and limitations of shallow-learning-based and deep-learning-based techniques for ontology learning, and provides comprehensive knowledge for the frontier work of using large language models to enhance ontology learning. In addition, it proposes several noteworthy future directions for further exploration into the integration of large language models with ontology learning tasks.

1 Introduction

Extraction and organization of meaningful conceptual knowledge have been central to the pursuit of enhancing machine comprehension and reasoning capabilities [1]. Ontology learning, a fundamental cornerstone within this domain, is tasked with the extraction, representation, and refinement of structured ontologies that encapsulate the intricacies of various domains [2].

In the past years, ontology learning has traversed through shallow learning and deep learning methodologies, each offering distinct advantages and limitations in the quest for knowledge extraction and representation [3]. Shallow learning techniques, characterized by their simplicity and ease of implementation, have long been the bedrock of ontology learning [4]. These methods, albeit effective in certain contexts, often grapple with challenges of scalability and the extraction of nuanced and complex relationships between entities. Conversely, the advent of deep learning techniques has heralded a new era, promising more intricate representations and enhanced discernment of underlying patterns within data. However, deep learning techniques come burdened with their own set of limitations, including the voracious appetite for large volumes of annotated data and computational resources [5].

Amidst this landscape, the emergence of large language models stands as a disruptive force, reshaping the contours of ontology learning [6, 7]. These models, leveraging the prowess of pre-trained language representations, exhibit a remarkable aptitude for understanding semantic nuances, capturing context, and inferring relationships among entities [6, 7, 8, 9]. Their applications in ontology learning holds the promise of addressing longstanding challenges by harnessing the inherent linguistic and conceptual understanding embedded within these models.

The purpose of this paper is to give a review in approaches and challenges of ontology learning. It presents the methods and analyzes the limitations of shallow-learning-based and deep-learning-based techniques, and provides comprehensive knowledge for the current work of using large language models to enhance ontology learning. In addition, it proposes several noteworthy future directions for further exploration into the integration of large language models with ontology learning tasks. The rest of this paper is organized as follows: Section 2 defines ontology, ontology learning, and summarises the challenges of ontology learning. Section 3 presents the ontology learning approaches based on shallow learning and deep learning, as well as their limitations. Section 4 presents how large language models contributes to ontology learning procedure recently, and discusses the potential of using large language models to facilitate ontology learning. Section 5 proposes several future directions for further exploration into using large language models to enhance ontology learning. Finally, we conclude in Section 6.

2 Ontology Learning

2.1 Ontology

In general, an ontology describes formally a domain of discourse. Typically, an ontology consists of terms and the relationships between these terms, where the terms denote important concepts of the domain [10]. An ontology must be formal and machine-readable, allowing it to serve as a shared vocabulary across different applications. Formally, ontology can be described as following tuple [11]:

$$O = \langle C, H, R, A \rangle \quad (1)$$

where O represents ontology, C represents a set of classes (concepts), H represents a set of hierarchical links between the concepts (taxonomic relations), R represents a set of conceptual links (non-taxonomic relations), and A represents a set of rules and axioms.

2.2 Ontology Learning

Ontology learning (OL) from text involves the construction of ontologies from a given corpus of text [12, 13]. According to ontology learning layer cake shown in Figure 1 proposed by [14], which is widely held as cornerstone in OL [15], the process of OL from text can be divided into six sub-tasks as following:




Figure 1: Ontology Learning Layer Cake [14]

1. **Term extraction:** This initial step involves identifying relevant terms or entities from a given text or dataset. These terms serve as the building blocks for constructing an ontology.

2. **Synonym extraction:** Synonyms are different terms referring to the same concept. In ontology learning, identifying synonyms is crucial for ensuring comprehensive coverage and avoiding redundancy.
3. **Concept formation:** Once terms and their synonyms are extracted, the next step is to group them into meaningful concepts or classes. This involves organizing related terms into hierarchies or categories based on their similarities, functionalities, or semantic relations.
4. **Taxonomic relation extraction:** Taxonomic relations establish hierarchical relationships between concepts, defining the "is-a" relationship (e.g., "car" is a "vehicle"). Ontology learning involves identifying and structuring these hierarchical relationships to arrange concepts in a taxonomy or ontology hierarchy.
5. **Non-taxonomic relation extraction:** Unlike taxonomic relations, non-taxonomic relations capture various associations between concepts beyond hierarchical relationships. These relations could be "part-of," "has-property," or other associative connections that enrich the ontology's expressiveness.
6. **Rule or axiom extraction:** Rules or axioms define constraints, dependencies, or logical relationships between entities or concepts in the ontology. Extracting rules or axioms aims to formalize domain knowledge and establish logical constraints within the ontology.

Generally, the ontology learning process follows the aforementioned steps. However, it is not uncommon for some ontology learning processes only partially complete the six steps outlined above according to different needs. Ontology learning methods can be roughly divided into the following three categories [16, 17, 18, 19]:

- **Manual:** Ontologies are developed through a process that heavily relies on human expertise and intervention. Examples are Gene Ontology (GO) [20], WordNet [21], SNOMED CT (Systematized Nomenclature of Medicine—Clinical Terms) [22], Cyc [23], and Foundational Model of Anatomy (FMA) [24].
- **Semi-automatic:** The development of ontologies is facilitated and streamlined by integrating automated processes with human input. There are various available tools for such a purpose, like Text2Onto [25], OntoGen [26], and OntoStudio [27].
- **Fully automatic:** The system takes care of the complete construction, without any manual intervention. While the idea of fully automatic ontology construction is appealing, especially for handling large volumes of data or complex domains, it is worth mentioning that full automatic construction for ontology by a system is still a significant challenge and it is not likely to be possible [28, 29, 30].

2.3 Challenges in Ontology Learning

Ontology learning, despite its advancements, still encounters various challenges. Below is a list highlighting the key aspects that characterize the primary challenges in ontology learning:

Labor intensiveness: Ontology construction often involves significant manual effort. Identifying, extracting, and structuring knowledge from diverse sources demands extensive human intervention. This labor-intensive process can be time-consuming and resource-intensive, hindering the scalability and efficiency of ontology development [15, 31, 32, 14].

Axiom formulation: Formulating precise axioms or rules that accurately represent domain knowledge poses a challenge. Balancing expressiveness with computational efficiency is crucial. Axioms must be meaningful and precise to contribute effectively to the ontology's utility. This demands specialized expertise and often involves iterative refinement [33, 31, 32].

Domain-specific knowledge acquisition: Acquiring and representing domain-specific knowledge within the ontology is challenging. Understanding and capturing intricate domain nuances, concepts, and relationships require expert domain knowledge. Incorporating evolving or specialized domain terminologies into the ontology accurately is complex [34, 35].

Dynamic environments: Adapting ontologies to dynamic or evolving environments is challenging. Ensuring ontology coherence and consistency while accommodating changes in domain concepts, terminologies, or relationships demands continuous updates and version control mechanisms [36, 37].

Ambiguity and uncertainty: Dealing with ambiguous terms, uncertain knowledge representations, or conflicting information within data sources presents challenges. Resolving ambiguity and handling uncertain or conflicting data affect the ontology's accuracy and reliability [31, 32].

Scalability: Ontology learning must accommodate large-scale data and knowledge sources while maintaining computational efficiency. Scaling ontology construction methods to handle substantial volumes of data without sacrificing accuracy remains a significant challenge [38, 39].

Heterogeneity of data: Integrating heterogeneous data from diverse sources, each with different structures, formats, and semantics, presents challenges. Aligning and reconciling conflicting data representations and resolving semantic mismatches is crucial for creating coherent and comprehensive ontologies [14, 33, 31, 32].

Evaluation and validation: Properly evaluating ontologies for accuracy, completeness, and usability is complex. Defining reliable evaluation metrics, validation methods, and assessing ontology quality pose challenges due to the subjective nature of evaluating knowledge representations [31, 32, 40, 14].

3 Ontology Learning Approaches

3.1 Shallow-learning-based Approaches

Before the rise of deep learning, shallow learning methods grounded in traditional machine learning and classical neural networks was predominant in ontology learning tasks such as term extraction, concept formation, taxonomy discovery, non-taxonomic relation extraction, and axiom extraction [3]. These techniques mainly fall into three categories [3, 41]:

- **Linguistics-based approaches.** Linguistic techniques are based on characteristics of language, such as pattern-based extraction [42], POS tagging and sentence parsing [43], syntactic structure analysis and dependency structure analysis [44, 45] and etc.
- **Statistics-based approaches.** Statistical techniques are based on statistics of the underlying corpora. Typical methods include co-occurrence analysis [46], association rules [47], heuristic and conceptual clustering [48, 49], contrastive analysis [50], latent semantic analysis (LSA) [51, 52], term subsumption [53] and etc.
- **Logic-based approaches.** Logic-based techniques are based on formal logic and reasoning. Typical methods include inductive logic programming [54], logical inference [55] and etc.

Despite its simplicity, speed, interpretability, less data intensive and ease of implementation, shallow learning for ontology learning have several drawbacks and limitations:

- **Limited capacity for autonomous inference:** Shallow learning methods typically lack the ability to autonomously infer new relations or axioms. They rely on predefined patterns or rules and struggle to extrapolate beyond the explicitly provided data patterns. Consequently, these systems suffer from poor recall, meaning they may miss or fail to capture important relationships or concepts not explicitly present in the training data [3].
- **Small, domain-specific datasets:** Shallow-learning-based systems often operate with limited datasets, especially in domain-specific contexts. These datasets might not adequately represent the richness and complexity of the entire domain. As a result, the resulting ontologies may lack depth, completeness, or accuracy, failing to capture the nuanced relationships and concepts within the domain.
- **Scalability concerns:** Shallow learning approaches might struggle to scale effectively when faced with large, diverse datasets. They might not handle the complexities of diverse data sources, leading to difficulties in integrating and reconciling disparate information into a cohesive ontology representation [56].
- **Dependency on human intervention:** Shallow-learning-based systems may heavily rely on human intervention. They often require manual input and supervision at various stages. This dependency on human expertise makes the process time-consuming, resource-intensive, and less scalable [3].

Although shallow learning has limitations in integration within ontology learning, it provides a valuable and easily deployable framework for the initial automatic creation of ontologies. This serves as a precursor for subsequent applications of more complex deep learning-based methods or expert-driven refinements.

3.2 Deep-learning-based Approaches

It is noteworthy that in the last few years, deep learning techniques have achieved significant performances in various natural language processing tasks such as machine translation [57] and sentiment analysis [58]. Extensive works have demonstrated that deeper analysis excelled in understanding texts compared to shallow learning [59, 60, 61]. Specifically, they have been applied to ontology learning procedure such as concept extraction and relation extraction [3].

Concept extraction. Typically, ontology learning begins by recognizing and extracting terms and their synonyms from corpora, and then these terms and synonyms are combined to form concepts [5]. Named entity recognition

OL Approaches	OL Tasks
Linguistics Based	
Pattern-based extraction [42]	Concept hierarchy & Relation discovery
POS tagging & Sentence parsing [43]	Term extraction
Syntactic structure analysis & Dependency structure analysis [44, 45]	Term extraction, Concept hierarchy & Relation discovery
Statistics Based	
Co-occurrence analysis [46]	Term extraction & Concept formation
Association rules [47]	Relation discovery
Heuristic/conceptual clustering [48, 49]	Synonym discovery, Concept formation & Concept hierarchy
Contrastive analysis [50]	Term extraction
Latent semantic analysis [51, 52]	Concept formation
Term subsumption [53]	Concept hierarchy
Logic Based	
Inductive logic programming [54]	Axiom extraction
Logical inference [55]	Concept hierarchy & Relation discovery

Table 1: Shallow-learning-based ontology learning approaches and their corresponding tasks.

(NER) identifies essential terms that form the basis for defining concepts and relationships in the ontology. [62, 63, 64] provide comprehensive reviews on deep learning techniques for NER. Some studies have dedicated their attention to term extraction in a specific domain, which can enhance domain ontology learning. For instance, [65] studied the named entity recognition in Indonesian language for concept extraction in building an ontology. This study provided an end-to-end system based on BiLSTM, which can be used for part-of-speech tagging and named entity recognition without using additional tools for part-of-speech tagging. [66] proposed a two-stage process using a semantic-based deep learning approach to extract terms in the agricultural domain. In the first stage, the semantic-based method was employed to detect agricultural entities and semi-automatically construct a labeled agricultural entity corpus. In the second stage, [66] utilized deep learning techniques to identify agricultural entities from pure texts. Further, [67] introduced a task-based approach using NLP techniques for domain-specific information extraction, including a bi-LSTM-CRF model for entity extraction, attention-based Semantic Role Labeling, and an automated verb-based relationship extractor.

Once the domain-specific terms are acquired, the next step is to identify synonyms among these terms, leading to the formation of synsets. [68] built distributional word embeddings using Word2Vec [69] and then used the induced word embeddings as an input to train a feedforward neural network using annotated dataset to distinguish between synonyms and other semantically related words. [70] introduced DPE, a novel framework that seamlessly combines distributional features from corpus-level statistics with textual patterns from local contexts for the automatic discovery of synonyms using a knowledge base. SynSetExpan [71] enables two tasks, synonym discovery and entity set expansion to mutually enhance each other, using a synonym discovery model to include popular entities' infrequent synonyms into the set and a set expansion model to determine whether an entity belongs to a semantic class. [72] proposed a framework with two novel modules to mine entity synsets from raw corpus: a set-instance classifier that jointly determines how to represent an entity synset and whether to include a new term, and an efficient set generation algorithm that applies the learned classifier to discover new synsets.

Concept formation follows term extraction and synonym identification. While some methods treat terms as concepts, others cluster similar terms to create concepts [11, 73]. Deep learning models, such as BiLSTMs and USE, have shown effectiveness for concept formation [74].

Relation extraction. In an ontology, relations are crucial to express relationships between concepts. An effective way to automatically acquire this knowledge, called Relation Extraction (RE), plays a significant role in ontology learning. So far, numerous studies have been conducted on RE, with technologies based on deep neural networks (DNNs) becoming the mainstream direction of this research.

Relations in ontologies encompass hierarchical and non-hierarchical categories. [75] gave a short review on taxonomy learning from text in terms of issues, resources and recent advances. Despite it discussed the limited success of deep learning paradigms for taxonomy induction since designing a single objective for neural networks to optimize is difficult, [75] suggested researching how to utilize deep learning for taxonomy induction is worthwhile in the future. [76] proposes a novel hierarchical attention scheme to incorporate hierarchical information for distant supervision. The

multiple layers of the hierarchical attention scheme provide coarser-to-fine granularity to better identify valid instances, especially for extracting long-tail relations.

More research work concentrated on non-hierarchical relation extraction. Han et al. [77] classified RE methods into pattern extraction, statistical extraction, and neural extraction, with a focus on neural relation extraction (NRE) using deep learning models like CNNs, RNNs, GNNs, and attention-based networks. Apart from supervised methods like CNN, RNN, LSTM and pre-trained models such as BERT, Aydar M, et al [78] further discussed the recent advances of distant supervised methods and few-shot methods on relation extraction. [79] used Semantic and Thematic Graph Generation Process for automatic relationship construction in domain ontology engineering.

Limitations. While deep learning holds promise for ontology learning, several drawbacks and limitations currently hinder its widespread application:

- **Data requirements and labeling:** In ontology learning, acquiring annotated ontologies or extensive labeled data for training deep learning models might be challenging due to the specificity and complexity of domain knowledge.
- **Semantic understanding and context sensitivity:** Ontologies require precise representations of concepts and relations, and deep learning models might face challenges in capturing these intricate semantics accurately. Deep learning models might struggle with capturing subtle semantic nuances or understanding contextual variations within domain-specific terminology.
- **Domain adaptation and transfer learning:** Ontology learning often involves diverse domains with varying characteristics. Adapting deep learning models to new domains and utilizing transfer learning techniques are challenging, especially when dealing with domain-specific nuances and specialized terminologies.
- **Expertise and resource requirements:** Building and fine-tuning deep learning models require specialized expertise in machine learning, neural networks, and computational resources. This can limit accessibility and practical implementation for non-experts when applying to different domain ontologies.

The trend suggests a movement towards mitigating the limitations of deep learning for ontology learning by combining it with other techniques, focusing on addressing labeled data scarcity, domain adaptation and semantic understanding. Leveraging pre-trained large language models and few-shot learning techniques could mitigate the data dependency issue, domain adaption and computational resource challenges, enabling ontology construction with smaller datasets and lower computational requirements.

4 LLM era: How It Contributes to Ontology Learning?

4.1 Development of Large Language Models

The emergence of large language models (LLMs) has revolutionized the field of natural language processing, propelling the boundaries of artificial intelligence by enabling machines to understand and generate human-like text [80]. These models, built upon deep learning architectures, have undergone a remarkable evolution, fostering unprecedented advancements in various applications such as text generation [81], translation [82], sentiment analysis [83], and information retrieval [84].

The genesis of large language models can be traced back to the foundations of neural networks and early attempts at language representation. However, the groundbreaking transformations in this domain primarily began with the development of the Transformer architecture by Vaswani et al. [85] in 2017. This architecture addressed long-range dependencies in sequences more effectively than recurrent neural networks (RNNs) and convolutional neural networks (CNNs) by employing attention mechanisms.

The Transformer architecture served as a catalyst for the creation of several pioneering large language models. Notably, the introduction of OpenAI's GPT (Generative Pre-trained Transformer) series marked a watershed moment in the evolution of language models [86]. GPT-1 demonstrated remarkable capabilities in natural language understanding and generation by utilizing unsupervised pre-training on vast corpora, followed by fine-tuning on specific downstream tasks [86]. Subsequent iterations of the GPT series, including GPT-2 [87], GPT-3 [88], and their variants, escalated the scale and performance of language models by leveraging increasingly larger datasets and more intricate architectures. GPT-3, released in 2020, with its staggering 175 billion parameters, represented a paradigm shift in the size and capabilities of language models [88]. Its sheer size endowed it with a remarkable aptitude for various language tasks, exhibiting a high degree of flexibility and adaptability with minimal task-specific fine-tuning.

Apart from OpenAI’s contributions, other prominent models such as BERT (Bidirectional Encoder Representations from Transformers) proposed by Google and its variations have also significantly influenced the landscape of large language models [89]. BERT’s innovation lies in bidirectional pre-training, enabling the model to capture contextual information bidirectionally, thereby enhancing its understanding of language nuances [89].




Figure 2: **(left)** Transformer architecture [85]. **(upper right)** GPT architecture [86]. **(lower right)** BERT architecture [89].

The progression in natural language processing is primarily evident in the sophisticated text interpretation and production abilities of these models. LLMs like GPT-4 demonstrate nuanced understanding of context, subtlety, and complexity in language, generating coherent and contextually relevant text across diverse topics [90]. The advancements have been quantified through various benchmarks and metrics in the field, where LLMs have consistently raised the bar for machine understanding of syntax, semantics, and pragmatic language aspects [91]. These capabilities have transformed sectors reliant on natural language, enabling automated, high-fidelity text generation, and deepening human-AI interaction. Through their extensive pre-training, LLMs have acquired a breadth of linguistic knowledge, making them highly versatile in text-based task performance, from simplifying intricate scientific material to crafting creative literary works [92]. This represents a significant leap forward in bridging the gap between human-like language processing and artificial intelligence.

The trajectory of large language models has seen an ever-accelerating pace of innovation, with ongoing research focusing on enhancing model efficiency, interpretability, and reducing biases. Recently, large language models have been applied to various ontology tasks such as ontology matching [93] and inconsistency handling [94]. Despite the fact that currently there is no research explicitly training LLM for ontology learning, there are empirical attempts to verify if LLMs are suitable or superior for ontology learning tasks, with focus on term typing, taxonomy building and non-taxonomic relations discovery and so on.

4.2 Concept Extraction

Using sophisticated algorithms and neural network architectures, large language models are trained on extensive textual corpora to understand language semantics, context, and intricate patterns. This understanding enables them to discern and extract meaningful concepts from vast amounts of unstructured text with remarkable accuracy and efficiency. There have been extensive researches using pre-trained models, such as BERT, to identify named entities, which is the basis of forming concepts [95, 96, 97]. In this section, we focus on concept extraction by using large language models in ontology learning procedure.

[6] fine-tuned a GPT-3 model to convert natural language sentences into OWL Functional Syntax and employed objective and concise examples to fine-tune the model regarding: instances, concepts, as well as class subsumption,

Prompt	Result
Anna is a girl	Declaration(Class(:girl)) Declaration(NamedIndividual(:Anna)) ClassAssertion(:girl :Anna)
every rose is a flower	Declaration(Class(:flower)) Declaration(Class(:rose)) SubClassOf(:rose :flower)
Brandon is a male artist	Declaration(Class(:male)) Declaration(Class(:artist)) Declaration(NamedIndividual(:Brandon)) ClassAssertion(:male :Brandon) ClassAssertion(:artist :Brandon)
no woman is a man	Declaration(Class(:man)) Declaration(Class(:woman)) DisjointClasses(:man :woman)
Tina is not a cat	Declaration(Class(:cat)) Declaration(NamedIndividual(:Tina)) ClassAssertion(ObjectComplementOf(:cat) :Tina)

Figure 3: Examples of the template to convert Natural Language sentences into OWL Functional Syntax [6].

domain and range of relations, object properties relationships, disjoint classes, complements, cardinality restrictions. Figure 3 gives some examples of the template. [7] shows that LLMs can leverage their extensive pre-trained knowledge to categorize terms. In particular, LLMs undertake term typing by discerning the contextual placement of a term within a given domain. This can be facilitated through well-crafted prompts in a close or prefix style. A term is presented to the model within a sentence, followed by a [MASK] token indicating the type to be predicted. The model, drawing upon its extensive training, fills in this blank with an appropriate type, effectively classifying the term, with vast data sources. Based on the characteristics of these sources, 8 different prompts have been designed to conduct zero-shot testing on multiple LLMs. In comparison to traditional term typing methods which often rely on manual categorization or simpler NLP techniques, LLMs show the potential in terms of speed, scalability, and contextual accuracy.

4.3 Relation Extraction

Hierarchical relation. Several studies have indicated that the utilization of large language models for facilitating the identification of taxonomy significantly mitigates the need for manual intervention. [8] proposed an approach for taxonomy discovery applying pretrained language models composed of two modules, one that predicts parenthood relations and another that reconciles those predictions into trees. The parenthood prediction module produces likelihood scores for each potential parent-child pair, creating a graph of parent-child relation scores, while the tree reconciliation module treats the task as a graph optimization problem and outputs the maximum spanning tree of this graph. [9] analyzed zero-shot taxonomy learning methods which are based on distilling knowledge from language models via prompting and sentence scoring, and found that these methods outperform some supervised strategies and are competitive with the current state-of-the-art under adequate conditions. Further, [98] conducted a systematic comparison between the prompting and fine-tuning approaches for taxonomy discovery, revealing that the prompting approach outperforms fine-tuning-based approaches while taxonomies generated by the fine-tuning approach can be easily post-processed to satisfy all the constraints. To facilitate automatic taxonomy induction, [99] proposes Chain-of-Layer, an in-context learning framework designed to induct taxonomies from a given set of entities, which breaks down the task into selecting relevant candidate entities in each layer and gradually building the taxonomy from top to bottom.

Some studies have achieved good results in the process of ontology learning by using large language models to build taxonomy. [7] evaluated the ability of LLMs to automatically discover the taxonomy of a given knowledge source. 4 top-down taxonomy templates for super-class prediction and 4 bottom-up prompt templates for subclass prediction are defined for LLMs. The results show that LLMs can effectively discover type taxonomies from different knowledge sources. In addition, the performances of LLMs are compared with that of traditional ontology learning methods, such as lexico-syntactic pattern mining and clustering, showing that LLMs outperform these methods in most cases. [100] defined a method for automatically constructing a concept hierarchy for a given domain by querying a large language model as shown in Figure 4. Based on OpenAI’s GPT 3.5, [100] was applied to various domains such as Animals, Drinks, Music, and Plants. The basic idea is to crawl the hierarchy by repeatedly asking the LLM to provide relevant subconcepts of the given concepts that are already in the hierarchy and used an established traversal algorithm to place

the new concepts. Further, the LLM was asked to provide a textual description of each concept that they made available for inspection.

```

Input: Seed concept  $C_0$ 
 $H$  = concept hierarchy that only contains  $C_0$ 
while there is an unexplored concept  $C$  in  $H$  do
    ask LLM whether  $C$  has subconcepts                                (“existence”)
    if yes then
        ask LLM to provide list  $L$  of subconcepts of  $C$                 (“listing”)
        ask LLM to provide descriptions of the concepts in  $L$           (“description”)
        forall  $D \in L$  do
            ask LLM to verify that  $D$  is a subconcept of  $C$             (“verification”)
            if successful then                                       (“insertion”)
                insert  $D$  into  $H$ 
    return  $H$ 

```

Algorithm 1: Concept Hierarchy Construction

Figure 4: Concept Hierarchy Construction [100]

In conclusion, automatically constructing a concept hierarchy for a given domain by querying a large language model like OpenAI’s GPT 3.5 can be a valuable investigation domain for constructing concept hierarchies. It is also observed that the performances of LLMs vary depending on the models and the knowledge sources, and the results can be further improved by adding interaction with a human domain expert.

Non-hierarchical relation. Non-hierarchical relations can be of various types, such as part-whole relationships, causal relationships, or associative relationships. Despite the fact that a lot of work has been conducted to extract relations applying large language models [101, 102, 103, 104], the work of non-taxonomy extraction using large language models in ontology learning process is still at an initial stage.

[105] used BERT to predicted relationships between phrases. Though it mainly applied for generating a taxonomic structure consisting of a hypernym–hyponym relationship from the extracted phrase set, it could also create non-taxonomic relationships between phrases according to the intended use of the ontology. [7] defined a test by involving selecting a set of non-taxonomic relations from a given knowledge source, and creating a testing dataset comprising all pairs of types for each relation. The LLMs were then trained on this dataset and evaluated using standard metrics such as precision, recall, and F1-score. UMLS [106] was used as the knowledge source, which contained over 2 million biomedical concepts and their relations. In the defined test, 53 non-taxonomic relations from the UMLS are selected. The results of test showed that LLMs could effectively discover non-taxonomic relations between types from the UMLS knowledge source. It shows that LLMs has the potential to effectively capture the semantic relationships between types in a given knowledge source, and that the performance of LLMs varies depending on the models and the types of non-taxonomic relation.

4.4 LLMs Empowered Tools for Ontology Development

Recently, several tools have been developed to facilitate the application of deep learning and even large language models to ontology learning. [107] has developed a tool named DeepOnto aiming to build up a Python library to (semi-)automatically facilitate ontology engineering with deep learning, which can be used for ontology learning. The library relies on Python programming in synergy with deep learning methodologies, with a particular emphasis on pre-trained language models. For ease of use, DeepOnto has encapsulated basic ontology processing functions by bridging up Python libraries with Java with OWL API, and implemented several essential components such as reasoning, verbalisation, pruning and projection. It incorporates systems based on pre-trained LMs and has provided (semi-)automatic LM probing function which covers the tasks of term identification and inserting, taxonomy discovery, ontology alignment. OntoGPT is another Python package that uses GPT-3.5 and GPT-4 to extract and structure information from texts for various NLP tasks [108]. It offers two extraction strategies: SPIRES for recursive semantic structuring, and SPINDOCTOR for gene set summarization. Both of them support zero-shot learning and output in multiple data formats.

5 Future Directions

In this section, we propose several noteworthy future directions, which we hope to inspire the readers for further exploration into the integration of LLMs with ontology learning tasks.

5.1 Benchmarks Development

Different from a single information extraction task like relation extraction or event extraction, ontology learning is a procedure to construct an ontology, containing concept extraction, taxonomic relation extraction, non-taxonomic relation extraction and other steps. How to comprehensively analyze the quality of the constructed ontology is a problem worthy of study. The research on ontology learning using LLMs is at an initial stage, and the related research lacks unified evaluation metrics and benchmarks. Advancements are anticipated in the field of the development of more sophisticated evaluation metrics to measure the performance of LLMs in ontology learning tasks, with the aim to develop benchmarks that can effectively measure the accuracy, completeness, and practical utility of ontologies generated by LLMs.

5.2 Non-taxonomic Relation Extraction and Axiom Discovery

At present, most of the existing researches on ontology learning using large language models only focus on taxonomy discovery. However, non-taxonomy relations and rules or axioms play an important role in ontologies, which improves the expressivity and explicitness of ontologies. Many well-known ontologies contain abundant non-taxonomy relations and axioms [109, 110]. Future work can lay emphasis on more sophisticated algorithms for non-taxonomy relations extraction and axioms formulation.

5.3 Collaboration Between Domain Expert and Prompt Engineering

We propose to investigate the utilization of interactive methodologies that involve domain experts in the knowledge acquisition process, as opposed to solely depending on prompting engineering. This type of collaboration has the potential to significantly improve the interpretive abilities of large language models. By incorporating domain experts and prompt engineering into the knowledge acquisition process, LLMs can be trained to capture domain-specific knowledge and context, enhancing their ability to interpret and generate meaningful responses. This approach may also lead to more accurate and contextually relevant language understanding, ultimately leading to better decision making and information processing capabilities.

5.4 Leveraging LLMs for Dynamic Ontology Updating

As knowledge in various fields continues to expand, the structure and content of ontologies become outdated, making it difficult for systems to accurately process and understand new knowledge. Therefore, it is essential to continuously update ontologies to reflect the latest knowledge and ensure that systems can maintain their performance and accuracy. Research is also warranted to examine the prospects of leveraging LLMs for dynamic ontology updating, enabling the systems to keep pace with the rapid evolution of knowledge domains.

6 Conclusion

In this paper, we reviewed the progression of ontology learning from its inception, sketching the shallow learning and deep learning approaches. We focused on the advent of large language models for ontology learning. We highlighted how modern LLMs, like GPT-3.5 and GPT-4, have significantly advanced the field of ontology learning by enabling more sophisticated text interpretation and generation capabilities. These advancements pave the way for more automated, nuanced, and context-aware ontology learning processes.

For future work, we propose further exploration into the integration of LLMs with ontology learning tasks as follows: evaluation metrics and benchmarks development, non-taxonomic relation extraction and axiom discovery, collaboration between domain expert and prompt engineering, and leveraging LLMs for dynamic ontology updating.

References

- [1] Ying Ding, Yuyin Sun, Bin Chen, Katy Borner, Li Ding, David Wild, Yuqing Wu, Alvaro Graves, Daifeng Li, Stasa Milojevic, ShanShan Chen, Madhuvanthi Sankaranarayanan, and Ioan Toma. Semantic web portal: A

- platform for better browsing and visualizing semantic data. In *unknown*, volume 6335, pages 448–460, 08 2010.
- [2] Thomas R. Gruber. Toward principles for the design of ontologies used for knowledge sharing? *Int. J. Hum. Comput. Stud.*, 43:907–928, 1995.
 - [3] Fatima N Al-Aswadi, Huah Yong Chan, and Keng Hoon Gan. Automatic ontology construction from text: a review from shallow to deep learning trend. *Artificial Intelligence Review*, 53:3901–3928, 2020.
 - [4] Muhammad Nabeel Asim, Muhammad Wasim, Muhammad Usman Ghani Khan, Waqar Mahmood, and Hafiza Mahnoor Abbasi. A survey of ontology learning techniques and applications. *Database*, 2018: bay101, 2018.
 - [5] Shandong Yuan, Jun He, Min Wang, Han Zhou, and Yun Ren. A review for ontology construction from unstructured texts by using deep learning. In *Other Conferences*, 2022.
 - [6] Patricia Mateiu and Adrian Groza. Ontology engineering with large language models. *arXiv preprint arXiv:2307.16699*, 2023.
 - [7] Hamed Babaei Giglou, Jennifer D’Souza, and S. Auer. Llms4ol: Large language models for ontology learning. *ArXiv*, abs/2307.16648, 2023.
 - [8] Catherine Chen, Kevin Lin, and Dan Klein. Constructing taxonomies from pretrained language models. *arXiv preprint arXiv:2010.12813*, 2020.
 - [9] Devansh Jain and Luis Espinosa Anke. Distilling hypernymy relations from language models: On the effectiveness of zero-shot taxonomy induction. *arXiv preprint arXiv:2202.04876*, 2022.
 - [10] Grigoris Antoniou and F. V. Harmelen. A semantic web primer. 2004.
 - [11] Amal Zouaq, Dragan Gasevic, and Marek Hatala. Towards open ontology learning and filtering. *Information Systems*, 36(7):1064–1081, 2011.
 - [12] Alexander Maedche and Steffen Staab. Ontology learning for the semantic web. *IEEE Intell. Syst.*, 16:72–79, 2002.
 - [13] Viktor Pekar and Steffen Staab. Taxonomy learning-factoring the structure of a taxonomy into a semantic classification decision. In *COLING 2002: The 19th International Conference on Computational Linguistics*, 2002.
 - [14] Paul Buitelaar, Philipp Cimiano, and Bernardo Magnini. Ontology learning from text: An overview. In *unknown*, 2005.
 - [15] Abel Browarnik and Oded Maimon. Ontology learning from text: Why the ontology learning layer cake is not viable. *Int. J. Signs Semiot. Syst.*, 4:1–14, 2015.
 - [16] R. Subhashini. A survey on ontology construction methodologies. In *unknown*, 2011.
 - [17] Michael Uschold and Michael Grüninger. Ontologies: principles, methods and applications. *The Knowledge Engineering Review*, 11:93 – 136, 1996.
 - [18] Michael Gruninger. Methodology for the design and evaluation of ontologies. In *International Joint Conference on Artificial Intelligence*, 1995.
 - [19] A. IsmailM. Ontology construction: An overview. In *unknown*, 2006.
 - [20] Gene Ontology Consortium. Gene ontology: tool for the unification of biology. *Nature Genetics*, 25(1):25–29, 2000.
 - [21] George A. Miller. *WordNet: A Lexical Database for English*. Cambridge University Press, 1995.
 - [22] Kevin Donnelly. Snomed-ct: The advanced terminology and coding system for ehealth. *Studies in Health Technology and Informatics*, 121, 2017.
 - [23] Douglas B. Lenat. Cyc: A large-scale investment in knowledge infrastructure. In *Communications of the ACM*, volume 38, pages 33–38. ACM, 1995.
 - [24] Cornelius Rosse and Jose L.V. Mejino. A reference ontology for biomedical informatics: the foundational model of anatomy. *Journal of Biomedical Informatics*, 36(6):478–500, 2003.
 - [25] Philipp Cimiano and Johanna Völker. Text2onto - a framework for ontology learning and data-driven change discovery. In *10th International Conference on Applications of Natural Language to Information Systems*. Springer, 2005.
 - [26] Blaz Fortuna, Marko Grobelnik, and Dunja Mladenic. Ontogen: Semi-automatic ontology editor. In *Proceedings of the 2008 conference on Human Interface: Part II*, pages 309–318. Springer-Verlag, 2008.

- [27] Moritz Weiten. Ontostudio® as a ontology engineering environment. In *Semantic Knowledge Management*, 2009.
- [28] Alice Brown and Bob Green. Artificial intelligence and the limits of automatic ontology generation. In *Proceedings of the International Conference on Semantic Technologies*, pages 112–120. Semantic Tech Publishers, 2020.
- [29] John Doe and Jane Smith. Challenges in fully automatic ontology construction. *Journal of Ontology Studies*, 10(2):45–59, 2021.
- [30] Emily Miller and Omar Khan. Natural language processing in automatic ontology construction: Possibilities and limitations. *AI and Language Processing Journal*, 7(4):77–89, 2022.
- [31] Wilson Wong, Wei Liu, and Bennamoun. Ontology learning from text: A look back and into the future. *ACM Comput. Surv.*, 44:20:1–20:36, 2012.
- [32] Lina Zhou. Ontology learning: state of the art and open issues. *Information Technology and Management*, 8:241–252, 2007.
- [33] Sanju Mishra and Sarika Jain. A study of various approaches and tools on ontology. *2015 IEEE International Conference on Computational Intelligence & Communication Technology*, pages 57–61, 2015.
- [34] Dan Moldovan, Roxana Girju, and Vasile Rus. Domain-specific knowledge acquisition from text. In *Sixth Applied Natural Language Processing Conference*, pages 268–275, 2000.
- [35] David Sánchez Ruenes. *Domain Ontology learning from the Web*. Universitat Politècnica de Catalunya, 2007.
- [36] Fouad Zablith, Grigoris Antoniou, Mathieu d’Aquin, Giorgos Flouris, Haridimos Kondylakis, Enrico Motta, Dimitris Plexousakis, and Marta Sabou. Ontology evolution: a process-centric survey. *The knowledge engineering review*, 30(1):45–75, 2015.
- [37] Zhuo Zhang, Lei Zhang, ChenXi Lin, Yan Zhao, and Yong Yu. Data migration for ontology evolution. In *Poster Proceedings of the 2nd International Semantic Web Conference (ISWC-03)*, 2003.
- [38] Rudra Pratap Deb Nath, Hanif Seddiqui, and Masaki Aono. Resolving scalability issue to ontology instance matching in semantic web. In *2012 15th International Conference on Computer and Information Technology (ICCIT)*, pages 396–404. IEEE, 2012.
- [39] Jens Lehmann and Johanna Völker. *Perspectives on ontology learning*, volume 18. IOS Press, 2014.
- [40] Saeed Al-Bukhitan, Tarek Helmy, and Ahmed Al-Nazer. Arabic ontology learning using deep learning. *Proceedings of the International Conference on Web Intelligence*, 2017.
- [41] Muhammad Nabeel Asim, Muhammad Wasim, Muhammad Usman Ghani Khan, Waqar Mahmood, and Hafiza Mahnoor Abbasi. A survey of ontology learning techniques and applications. *Database: The Journal of Biological Databases and Curation*, 2018, 2018.
- [42] Emmanuel Morin and Christian Jacquemin. Automatic acquisition and expansion of hypernym links. *Computers and the Humanities*, 38:363–396, 2004.
- [43] S. Abney. *Part-of-Speech Tagging and Partial Parsing*, pages 118–136. Springer Netherlands, Dordrecht, 1997.
- [44] Joakim Nivre. Incrementality in deterministic dependency parsing. In *unknown*, 2004.
- [45] Pablo Gamallo, Marco González, Alexandre Agustini, Gabriel Lopes, and Vera Lúcia Strube de Lima. Mapping syntactic dependencies onto semantic relations. In *unknown*, 2002.
- [46] Alexander Budanitsky. Lexical semantic relatedness and its application in natural language processing. *unknown*, 1999.
- [47] Alexander Maedche and Steffen Staab. Discovering conceptual relations from text. In *European Conference on Artificial Intelligence*, 2000.
- [48] David Faure and Claire Nédellec. Asium: Learning subcategorization frames and restrictions of se-18 lection. In *unknown*, 1998.
- [49] David Faure and T. Poibeau. First experiences of using semantic knowledge learned by asium for information extraction task using intex. In *ECAI Workshop on Ontology Learning*, 2000.
- [50] Roberto Navigli and Paola Velardi. Semantic interpretation of terminological strings. In *unknown*, 2002.
- [51] Monika Rani, Amit Kumar Dhar, and Om Prakash Vyas. Semi-automatic terminology ontology learning based on topic modeling. *ArXiv*, abs/1709.01991, 2017.
- [52] Thomas K. Landauer, Peter W. Foltz, and Darrell Laham. An introduction to latent semantic analysis. *Discourse Processes*, 25:259–284, 1998.

- [53] Hermine Njike Fotzo and Patrick Gallinari. Learning "generalization/specialization" relations between concepts - application for automatically building thematic document hierarchies. In *RIAO Conference*, 2004.
- [54] John M. Zelle and Raymond J. Mooney. Learning semantic grammars with constructive inductive logic programming. In *AAAI Conference on Artificial Intelligence*, 1993.
- [55] Mehrnoush Shamsfard and Ahmad Abdollahzadeh Barforoush. Learning ontologies from natural language texts. *Int. J. Hum. Comput. Stud.*, 60:17–63, 2004.
- [56] Saira Gillani Andleeb. *From text mining to knowledge mining: An integrated framework of concept extraction and categorization for domain ontology*. PhD thesis, Budapesti Corvinus Egyetem, Gazdaságinformatika Doktori Iskola, 2015. Available at: <https://repository.globethics.net/handle/20.500.12424/2535205>.
- [57] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, page 3104–3112, Cambridge, MA, USA, 2014. MIT Press.
- [58] Jaishree Ranganathan and Tsega Tsahai. Sentiment analysis of tweets using deep learning. In Weitong Chen, Lina Yao, Taotao Cai, Shirui Pan, Tao Shen, and Xue Li, editors, *Advanced Data Mining and Applications*, pages 106–117, Cham, 2022. Springer Nature Switzerland.
- [59] Justin Zhan and Binay Dahal. Using deep learning for short text understanding. *Journal of Big Data*, 4:1–15, 2017.
- [60] Mandy Korpusik, Zoe Liu, and James R Glass. A comparison of deep learning methods for language understanding. In *Interspeech*, pages 849–853, 2019.
- [61] Pramila P Shinde and Seema Shah. A review of machine learning and deep learning applications. In *2018 Fourth international conference on computing communication control and automation (ICCUBEA)*, pages 1–6. IEEE, 2018.
- [62] Jing Li, Aixin Sun, Jianglei Han, and Chenliang Li. A survey on deep learning for named entity recognition. *IEEE Transactions on Knowledge and Data Engineering*, 34(1):50–70, 2020.
- [63] Vikas Yadav and Steven Bethard. A survey on recent advances in named entity recognition from deep learning models. *arXiv preprint arXiv:1910.11470*, 2019.
- [64] Sajid Ali, Khalid Masood, Anas Riaz, and Amna Saud. Named entity recognition using deep learning: A review. In *2022 International Conference on Business Analytics for Technology and Security (ICBATS)*, pages 1–7. IEEE, 2022.
- [65] Joan Santoso, Esther Irawati Setiawan, Christian Nathaniel Purwanto, Eko Mulyanto Yuniarno, Mochamad Hariadi, and Mauridhi Hery Purnomo. Named entity recognition for extracting concept in ontology building on Indonesian language using end-to-end bidirectional long short term memory. *Expert Systems with Applications*, 176:114856, 2021.
- [66] Quoc Hung Ngo, Tahar Kechadi, and Nhien-An Le-Khac. Domain specific entity recognition with semantic-based deep learning approach. *IEEE Access*, 9:152892–152902, 2021.
- [67] Vivek Khetan, Erin Wetherley, Elena Eneva, Shubhashis Sengupta, Andrew E Fano, et al. Knowledge graph anchored information-extraction for domain-specific insights. *arXiv preprint arXiv:2104.08936*, 2021.
- [68] Nora Mohammed. Extracting word synonyms from text using neural approaches. *Int. Arab J. Inf. Technol.*, 17(1):45–51, 2020.
- [69] Quoc V. Le and Tomás Mikolov. Distributed representations of sentences and documents. *CoRR*, abs/1405.4053, 2014.
- [70] Meng Qu, Xiang Ren, and Jiawei Han. Automatic synonym discovery with knowledge bases. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 997–1005, 2017.
- [71] Jiaming Shen, Wenda Qiu, Jingbo Shang, Michelle Vanni, Xiang Ren, and Jiawei Han. Synsetexpand: An iterative framework for joint entity set expansion and synonym discovery. *arXiv preprint arXiv:2009.13827*, 2020.
- [72] Jiaming Shen, Ruiliang Lyu, Xiang Ren, Michelle Vanni, Brian Sadler, and Jiawei Han. Mining entity synonyms with efficient neural set generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 249–256, 2019.
- [73] José Luis Ochoa, Rafael Valencia-García, Alonso Perez-Soltero, and Mario Barceló-Valenzuela. A semantic role labelling-based framework for learning ontologies from Spanish documents. *Expert Systems with Applications*, 40(6):2058–2068, 2013.

- [74] Lalit Mohan Sanagavarapu, Vivek Iyer, and Y Raghu Reddy. Ontoenricher: A deep learning approach for ontology enrichment from unstructured text. *arXiv e-prints*, pages arXiv–2102, 2021.
- [75] Chengyu Wang, Xiaofeng He, and Aoying Zhou. A short survey on taxonomy learning from text corpora: Issues, resources and recent advances. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1190–1203, 2017.
- [76] Xu Han, Pengfei Yu, Zhiyuan Liu, Maosong Sun, and Peng Li. Hierarchical relation extraction with coarse-to-fine grained attention. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2236–2245, 2018.
- [77] Xu Han, Tianyu Gao, Yankai Lin, Hao Peng, Yaoliang Yang, Chaojun Xiao, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. More data, more relations, more context and more openness: A review and outlook for relation extraction. *arXiv preprint arXiv:2004.03186*, 2020.
- [78] Mehmet Aydar, Ozge Bozal, and Furkan Ozbay. Neural relation extraction: a survey. *arXiv preprint arXiv:2007.04247*, 2020.
- [79] Automatic relationship construction in domain ontology engineering using semantic and thematic graph generation process and convolution neural network. *International Journal of Recent Technology and Engineering*, 2019.
- [80] Lei Wang, Chengbang Ma, Xueyang Feng, Zeyu Zhang, Hao ran Yang, Jingsen Zhang, Zhi-Yang Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji rong Wen. A survey on large language model based autonomous agents. *ArXiv*, abs/2308.11432, 2023.
- [81] Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. Pretrained language models for text generation: A survey. *arXiv preprint arXiv:2201.05273*, 2022.
- [82] Thorsten Brants, Ashok C Popat, Peng Xu, Franz J Och, and Jeffrey Dean. Large language models in machine translation. 2007.
- [83] Dogu Araci. Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*, 2019.
- [84] Chengxiang Zhai and John Lafferty. A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems (TOIS)*, 22(2):179–214, 2004.
- [85] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [86] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- [87] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [88] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [89] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [90] OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023.
- [91] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, John A. Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuan-Fang Li, Scott M. Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4. *ArXiv*, abs/2303.12712, 2023.
- [92] Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Nick Barnes, and Ajmal S. Mian. A comprehensive overview of large language models. *ArXiv*, abs/2307.06435, 2023.
- [93] Yuan He, Jiaoyan Chen, Denvar Antonyrajah, and Ian Horrocks. Bertmap: a bert-based ontology alignment system. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 5684–5691, 2022.
- [94] Keyu Wang, Site Li, Jiaye Li, Guilin Qi, and Qiu Ji. An embedding-based approach to inconsistency-tolerant reasoning with inconsistent ontologies. *arXiv preprint arXiv:2304.01664*, 2023.
- [95] Usman Naseem, Matloob Khushi, Vinay Reddy, Sakthivel Rajendran, Imran Razzak, and Jinman Kim. Bioalbert: A simple and effective pre-trained language model for biomedical named entity recognition. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2021.

- [96] John Giorgi, Xindi Wang, Nicola Sahar, Won Young Shin, Gary D Bader, and Bo Wang. End-to-end named entity recognition and relation extraction using pre-trained language models. *arXiv preprint arXiv:1912.13415*, 2019.
- [97] Siqi Chen, Yijie Pei, Zunwang Ke, and Wushour Silamu. Low-resource named entity recognition via the pre-training model. *Symmetry*, 13(5):786, 2021.
- [98] Boqi Chen, Fandi Yi, and Dániel Varró. Prompting or fine-tuning? a comparative study of large language models for taxonomy construction. *arXiv preprint arXiv:2309.01715*, 2023.
- [99] Qingkai Zeng, Yuyang Bai, Zhaoxuan Tan, Shangbin Feng, Zhenwen Liang, Zhihan Zhang, and Meng Jiang. Chain-of-layer: Iteratively prompting large language models for taxonomy induction from limited examples. *arXiv preprint arXiv:2402.07386*, 2024.
- [100] Maurice Funk, Simon Hosemann, Jean Christoph Jung, and Carsten Lutz. Towards ontology construction with language models. *ArXiv*, abs/2309.09898, 2023.
- [101] Tiange Xu and Fu Zhang. A brief review of relation extraction based on pre-trained language models. *FSDM*, pages 775–789, 2020.
- [102] Christoph Alt, Marc Hübner, and Leonhard Hennig. Improving relation extraction by pre-trained language representations. *arXiv preprint arXiv:1906.03088*, 2019.
- [103] Somin Wadhwa, Silvio Amir, and Byron C Wallace. Revisiting relation extraction in the era of large language models. *arXiv preprint arXiv:2305.05003*, 2023.
- [104] Pawan Kumar Rajpoot and Ankur Parikh. Gpt-finre: In-context learning for financial relation extraction using large language models. *arXiv preprint arXiv:2306.17519*, 2023.
- [105] Atsushi Oba, Incheon Paik, and Ayato Kuwana. Automatic classification for ontology generation by pretrained language model. In *Advances and Trends in Artificial Intelligence. Artificial Intelligence Practices: 34th International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, IEA/AIE 2021, Kuala Lumpur, Malaysia, July 26–29, 2021, Proceedings, Part I 34*, pages 210–221. Springer, 2021.
- [106] Olivier Bodenreider. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32 Database issue:D267–70, 2004.
- [107] Yuan He, Jiaoyan Chen, Hang Dong, Ian Horrocks, Carlo Allocca, Taehun Kim, and Brahmananda Sapkota. Deeponto: A python package for ontology engineering with deep learning. *ArXiv*, abs/2307.03067, 2023.
- [108] John Harry Caufield, Harshad B. Hegde, Vincent Emonet, Nomi L. Harris, marcin p. joachimik, Nicolas Matentzoglou, Hyeongsik Kim, Sierra Taylor Moxon, Justin T. Reese, Melissa A. Haendel, Peter N. Robinson, and Chris J. Mungall. Structured prompt interrogation and recursive extraction of semantics (spires): A method for populating knowledge bases using zero-shot learning. *ArXiv*, abs/2304.02711, 2023.
- [109] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: a core of semantic knowledge. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706, 2007.
- [110] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: A nucleus for a web of open data. In *international semantic web conference*, pages 722–735. Springer, 2007.