

Enriching Ontologies with Disjointness Axioms using Large Language Models

Elias Crum^{1,*}, Antonio De Santis^{2,*}, Manon Ovide^{3,*}, Jiaxin Pan^{4,*}, Alessia Pisu^{5,*},
Nicolas Lazzari⁶ and Sebastian Rudolph⁷

¹Ghent University, Belgium elias.crum@ugent.be

²Politecnico di Milano, Italy antonio.desantis@polimi.it

³University of Tours, France manon.ovide@univ-tours.fr

⁴University of Stuttgart, Germany jiaxin.pan@ki.uni-stuttgart.de

⁵University of Cagliari, Italy alessia.pisu96@unica.it

⁶University of Pisa and University of Bologna, Italy nicolas.lazzari3@unibo.it

⁷TU Dresden, Germany sebastian.rudolph@tu-dresden.de

Abstract

Ontologies often lack explicit disjointness declarations between classes, despite their usefulness for sophisticated reasoning and consistency checking in Knowledge Graphs. In this study, we explore the potential of Large Language Models (LLMs) to enrich ontologies by identifying and asserting class disjointness axioms. Our approach aims at leveraging the implicit knowledge embedded in LLMs, using prompt engineering to elicit this knowledge for classifying ontological disjointness. We validate our methodology on the DBpedia ontology, focusing on open-source LLMs. Our findings suggest that LLMs, when guided by effective prompt strategies, can reliably identify disjoint class relationships, thus streamlining the process of ontology completion without extensive manual input. For comprehensive disjointness enrichment, we propose a process that takes logical relationships between disjointness and subclass statements into account in order to maintain satisfiability and reduce the number of calls to the LLM. This work provides a foundation for future applications of LLMs in automated ontology enhancement and offers insights into optimizing LLM performance through strategic prompt design. Our code is publicly available on GitHub at <https://github.com/n28div/llm-disjointness>.

Keywords

Large Language Models, Disjointness Learning, Ontology Enrichment

1. Introduction

It is generally understood that complementing the factual (assertional) knowledge represented in Knowledge Graphs with ontological (terminological) information greatly advances the usefulness of the ensuing knowledge base in terms of querying and many other downstream tasks. This is because combining assertional information with terminological background knowledge allows for the derivation of a vast amount of implicit knowledge, which is not explicitly stated


KBC-LM'24: Knowledge Base Construction from Pre-trained Language Models workshop at ISWC 2024

This paper presents joint research that originated from the team project of “House Slytherin” at the 2024 International Semantic Web Summer School (ISWS) in Bertinoro, Italy.

*These authors contributed equally.



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

 CEUR Workshop Proceedings (CEUR-WS.org)

in the knowledge base but follows logically from it and thus can be taken into account for all kinds of knowledge management activities, including query answering.

The by far most widespread type of ontological information added to knowledge graphs is *taxonomic* in nature, that is, it is related to (i) putting the individual objects of interest into categories, usually referred to as *classes*, based on shared characteristics and (ii) establishing set-theoretic relationships between these classes. Among the diverse possible such taxonomic relationships, the *subclass/superclass* relationships – tightly connected to the linguistic hyponymy/hypernymy relationships of the corresponding class names – are the ones predominantly found across numerous ontologies today, typically forming sizeable conceptual hierarchies. As an example, the subclass/superclass relationship between the classes `Mammal` and `Vertebrate` implies that any object that belongs to (or, in more technical terms: is an *instance* of) the class `Mammal` also must belong to the class `Vertebrate`.

Another well-known basic type of taxonomic relationship between two classes is that of *disjointness*. Two classes are said to be *disjoint* if it is impossible that they have common instances, which, intuitively, means that the two classes cannot overlap, and membership in these two classes is mutually exclusive. For example, disjointness of the classes `Mammal` and `Fish` implies that any instance of `Mammal` must not be an instance of `Fish`. Given the symmetric nature of disjointness, this is logically equivalent to saying that any instance of `Fish` must not be an instance of `Mammal`. As opposed to subclass statements, which allow for inferring positive facts from other positive facts, disjointness statements enable the inference of *negated* facts. For example, given the fact that `Flipper` is an instance of `Mammal`, the above subclass relationship gives rise to the information that `Flipper` is an instance of `Vertebrate`, whereas the disjointness statement allows us to infer the information that `Flipper` must not be an instance of `Fish`. This fact makes disjointness information particularly valuable in the context of machine-learning approaches that rely on the presence of negative examples, such as Knowledge Graph Embedding.

When specifying taxonomic relationships between classes in the course of the ontology design process, it should be kept in mind that they are not meant to reflect spurious relationships in the data currently available, but rather they are supposed to represent immutable background knowledge that continues to hold in different situations or at different points in time. For instance, although historically, no woman has served as US President, a woman may be elected as the US President in the future. Therefore, the corresponding classes `woman` and `USPresident` are not (ontologically) disjoint.¹ To reflect this situation more formally, one can employ the idea of *possible or conceivable worlds* (referred to as *interpretations* in model-theoretic terms), which e.g., include potential future or just hypothetical circumstances. Then, a certain taxonomic relationship (such as subclass or disjointness) between two classes holds if the corresponding set relation (such as subset or intersection-emptiness) holds between the sets of class instances *in every conceivable world (under every conceivable interpretation)*. Based on this, we will employ a very lightweight logical framework to give our arguments a formal underpinning: Stipulating a set \mathcal{F} of conceivable worlds, we define taxonomic relationships for this set. The goal of ontological knowledge modeling is to capture \mathcal{F} using a knowledge base \mathcal{K} whose statements rule out the inconceivable worlds so that only the conceivable ones remain as models of \mathcal{K} .

¹We might call them *materially disjoint* due to the absence of material evidence demonstrating their non-disjointness.

Definition 1. Fixing a vocabulary consisting of a set \mathbf{C} of class names and a set \mathbf{I} of individual names, an interpretation $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$ consists of a set Δ called the domain and a function $\cdot^{\mathcal{I}}$ mapping every class name $C \in \mathbf{C}$ to a subset $C^{\mathcal{I}} \subseteq \Delta$ and every individual name $i \in \mathbf{I}$ to an element $i^{\mathcal{I}} \in \Delta$.

Let \mathcal{F} be a set of interpretations, representing the conceivable worlds. Then, for an individual name $i \in \mathbf{I}$ and for concept names $C, D \in \mathbf{C}$ we call

- i an instance of C (written $i : C$) if every interpretation $\mathcal{I} \in \mathcal{F}$ satisfies $i^{\mathcal{I}} \in C^{\mathcal{I}}$,
- C a subclass of D (written $C \sqsubseteq D$) if every interpretation $\mathcal{I} \in \mathcal{F}$ satisfies $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$,
- C disjoint with D if every interpretation $\mathcal{I} \in \mathcal{F}$ satisfies $C^{\mathcal{I}} \cap D^{\mathcal{I}} = \emptyset$.
- C incoherent if every interpretation $\mathcal{I} \in \mathcal{F}$ satisfies $C^{\mathcal{I}} = \emptyset$.

As discussed above, ontologically dictated taxonomic relationships can be leveraged for sophisticated reasoning and consistency-checking tasks when reasoning over a knowledge graph. Yet, despite their usefulness, disjointness relationships are rarely explicitly recorded within an ontology. Research on 1,275 ontologies showed that only 97 of them include disjointness assertions [1]. Arguably, this can be explained by the fact that disjointness information is so self-evident from a human common-sense point of view, that human experts are often not aware that it is not logically “built-in” but needs to be explicitly specified. For this reason, semi-automated labeling of disjoint classes could be advantageous. Recent approaches [2, 3, 4] propose supervised and unsupervised models using various features in disjointness axioms. However, the generalizability of these methods is limited to their specific datasets and cannot be implemented on a large scale. Additionally, the sophisticated feature engineering required hinders their practical application. Therefore, a method that functions independently of feature design and dataset restrictions is highly desirable.

Given that (i) ontological class descriptions are often recorded as (or associated with) terms in natural language and (ii) LLMs have been found to possess wide linguistic and semantic working knowledge, we aim to assess the potential of LLMs to decide on the question which classes ought to be disjoint while assessing the impact of prompt engineering on classification validity. We hypothesize that through the use of prompt engineering, LLMs are to classify ontologically disjoint classes with high validity in both positive (two classes are ontologically disjoint), and negative (two classes are not ontologically disjoint), cases. We test our hypothesis on the DBpedia ontology² using LLMs. We propose a method that intertwines the LLM-based disjointness classification with basic logical inferencing to increase efficiency, maintain consistency, and minimize the number of calls to the LLM.

Thus, this paper is dedicated to answering the following main research questions:

- RQ1: Can LLMs help enrich ontologies with class disjointness axioms?
- RQ2: Which LLM prompts work better for disjointness discovery?
- RQ3: How can we exploit taxonomic relationships to reduce interaction with the LLM?

²<https://DBpedia.org/ontology/>

2. Related Work

Disjointness Learning Models for disjointness learning can be categorized into supervised and unsupervised approaches. In the unsupervised category, Schlobach [5] follows the *strong disjointness assumption* [6], which posits that children of a common parent in the subsumption hierarchy should be considered disjoint. They introduced a pinpointing algorithm to identify minimal sets of axioms that need revision to make an ontology coherent, thereby enriching appropriate disjointness statements. However, this approach neglects background knowledge, which could be beneficial in identifying disjoint classes. Rizzo et al. [4] proposes an unsupervised approach based on concept learning and inductive classification. This method employs a hierarchical conceptual clustering technique capable of providing intensional cluster descriptions and utilizes a novel form of semi-distances over individuals in an ontological knowledge base, incorporating available background knowledge. In the supervised category, Völker et al. [2, 3] gather syntactic and semantic evidence, such as positive and negative association rules as well as correlation coefficients, from various sources to establish a strong foundation for learning disjointness. However, their work exploits background knowledge and reasoning only to a limited extent. Subsequent work, the DL-Learner by Lehmann [7], uses Inductive Logic Programming (ILP) for learning class descriptions, including disjointness. Despite these advancements, disjointness learning with LLMs remains much underexplored.

Large Language Models In recent years, Large Language Models (LLMs) have become state-of-the-art for Natural Language Processing and have also significantly impacted other fields such as knowledge engineering [8, 9, 10, 11]. LLMs rely on pre-training Transformer models [12] over large-scale unlabeled corpora. Pre-trained context-aware word representations achieve state-of-the-art performance on various downstream tasks and set the “pre-training and fine-tuning” learning paradigm. Early LLMs, such as BERT [13], utilized relatively small training corpora and required fine-tuning for specific downstream tasks. However, subsequent research demonstrated that scaling up both model size and dataset volume significantly enhances performance. GPT-3 [14], for instance, achieves competitive results through few-shot learning and in-context learning without parameter updates. GPT-3.5 further improves capabilities by incorporating reinforcement learning from human feedback (RLHF). The introduction of GPT-4 [15] marked a milestone by extending beyond text input to include multimodal signals. Meta AI introduced the collection of LLaMA models [16, 17] with four different sizes. Other notable LLMs, such as Claude, Gemini [18], and Mixtral [19], have also garnered significant attention.

Prompt Engineering Designing effective prompts for LLMs is essential for maximizing their potential. Key strategies in prompt engineering include zero-shot [20], few-shot [14], and chain-of-thought [21] prompting. *Zero-shot* [20] involves providing task descriptions to LLMs without any input-output examples, relying on the models’ pre-existing knowledge to generate responses. *Few-shot* [14] includes input-output examples, guiding the models’ generation process. *Chain-of-Thought (CoT)* [22] promotes coherent and step-by-step reasoning by decomposing a complex question into a series of simpler logical reasoning questions, mimicking human problem-solving processes. This method has been shown to significantly improve performance on reasoning tasks [22]. However, the need for multiple prompts makes this approach difficult to use at large

scales. With this in mind, Kojima et al. [23] proposed *Zero-shot-CoT* prompting. They found that by appending the phrase “Let’s think step by step.” to the end of a question, LLMs can generate a chain of thought that leads to more accurate answers w.r.t the vanilla zero-shot approach.

3. Resources

To effectively assess the ability of LLMs to support the assertion of disjointness axioms, we ideally require a reference ontology that includes a sized set of classes, to ensure diversity during the experiments and some disjoint classes in its description, preferably specified through a specific disjoint class property such as `owl:disjointWith`. These criteria maximize the generalizability of the approach and encourage its use for future studies.

Several ontologies can be identified for this task, from foundational ontologies, such as DOLCE³ or UFO⁴, to domain-specific ontologies, such as FoodOn⁵. Disjointness axioms from these ontologies, however, are not intuitive and require extensive common-sense reasoning and domain knowledge. For instance, DOLCE defines an `Event` to be disjoint from an `Object` while UFO does not. Both axioms are correct, as they deeply depend on their philosophical commitment to these abstract concepts. Similarly, the FoodOn ontology asserts that the *Arabia coffee plant*⁶, the plant used to produce black coffee, is disjoint with *Camellia sinensis*⁷, the plant used to produce black tea. In this case, deciding whether the two plants should be considered disjoint highly depends on the domain of the ontology. To avoid feeding the LLM with classes whose disjointness highly depends on the context or domain, we choose to avoid foundational and domain-specific ontologies for our initial experiments. Moreover, as our interaction with the LLM is based on natural language, we only consider ontologies that provide natural language labels for classes via labeling properties, such as `skos:prefLabel` or `rdfs:label`.

We ultimately decided to use the DBpedia ontology⁸ because of its general popularity and conformity with dataset minimal requirements. Since the DBpedia ontology is created through a crowdsourcing approach [24], the availability of disjointness axioms cannot be expected to be equally accurate across all classes, as it depends on the annotators’ expertise and diligence. This issue has been actively discussed within the DBpedia community⁹. The main drawback is the lack of a systematic approach in the creation of the taxonomy, which greatly impacts the consistency of the ontology when disjointness axioms are asserted. In particular, we found 23 explicit disjointness axioms in the DBpedia ontology. In Section 4 we show how exploiting automated reasoning techniques allows the creation of a larger pool of disjoint classes. In Table 1 a selection of disjointness axioms within the ontology is shown. Indeed, most of the disjointness axioms are *universally* known common-sense relations, such as disjointness between `dbo:Fish` and `dbo:Mammal` or `dbo:Agent` and `dbo:Place`.

³<https://github.com/appliedontolab/DOLCE/blob/main/OWL/DOLCEbasic.owl>

<http://www.ontologydesignpatterns.org/ont/dul/DUL.owl>

⁴<https://nemo-ufes.github.io/gufo/>

⁵<https://foodon.org/>

⁶https://en.wikipedia.org/wiki/Coffea_arabica

⁷https://en.wikipedia.org/wiki/Camellia_sinensis

⁸<https://DBpedia.org/ontology/>, often referred to with the `dbo: namespace`, which we omit hereafter

⁹<https://github.com/DBpedia/ontology-tracker/issues/2>

Class A	Class B
http://DBpedia.org/ontology/Person	http://DBpedia.org/ontology/ProtohistoricalPeriod
http://DBpedia.org/ontology/Person	http://DBpedia.org/ontology/UnitOfWork
http://DBpedia.org/ontology/Agent	http://DBpedia.org/ontology/Place
http://DBpedia.org/ontology/Fish	http://DBpedia.org/ontology/Mammal
http://DBpedia.org/ontology/Event	http://DBpedia.org/ontology/Person

Table 1
Examples of pairs of classes explicitly specified as ontologically disjoint in the DBpedia ontology.

4. Proposed approach

We now describe our approach which, given a Knowledge Base, clarifies for *every* pair of named classes of that ontology if disjointness should hold between the two classes or not. At the core of the approach is prompting an LLM to exploit the semantic and linguistic “world knowledge” it has obtained from training on vast amounts of textual data. The two major underlying objectives of our approach are:

1. Ensuring that the resulting disjointness-enriched ontology is satisfiable (i.e., contradiction-free) for usability reasons since otherwise it would be unusable for any reasoning tasks, including ontology-supported querying.
2. Minimizing the number of interactions with the LLM for efficiency reasons and cost-awareness.

We propose to address both objectives using automated reasoning. More specifically, we continuously materialize all the (non-)disjointness information that follows logically from the original knowledge base plus the already acquired disjointness information. Thus, the LLM is only queried about the disjointness status of pairs of classes, when neither of the outcomes would result in an inconsistency. In this way, the derived information remains contradiction-free “by design” and, at the same time, the number of queries to the LLM is significantly reduced. Our approach relies on several logical correspondences, discussed in the following.

Proposition 1. *Let \mathcal{K} be a knowledge base and let C_1, C_2, D_1, D_2 be classes of \mathcal{K} such that the following statements follow from \mathcal{K} : (i) C_1 and C_2 are disjoint, (ii) D_1 is a subclass of C_1 , (iii) D_2 is a subclass of C_2 . Then \mathcal{K} also entails that D_1 and D_2 are disjoint.*

Proof. Consider an arbitrary model \mathcal{I} of \mathcal{K} . According to the assumptions and in view of Definition 1, we know that (i) $C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}} = \emptyset$, (ii) $D_1^{\mathcal{I}} \subseteq C_1^{\mathcal{I}}$, and (iii) $D_2^{\mathcal{I}} \subseteq C_2^{\mathcal{I}}$. We equivalently express (ii) and (iii) by (ii’) $D_1^{\mathcal{I}} = C_1^{\mathcal{I}} \cap D_1^{\mathcal{I}}$, and (iii’) $D_2^{\mathcal{I}} = C_2^{\mathcal{I}} \cap D_2^{\mathcal{I}}$. This allows us to infer $D_1^{\mathcal{I}} \cap D_2^{\mathcal{I}} = (C_1^{\mathcal{I}} \cap D_1^{\mathcal{I}}) \cap (C_2^{\mathcal{I}} \cap D_2^{\mathcal{I}}) = (C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}}) \cap (D_1^{\mathcal{I}} \cap D_2^{\mathcal{I}}) = \emptyset \cap (D_1^{\mathcal{I}} \cap D_2^{\mathcal{I}}) = \emptyset$. \square

We exploit this property to use subclass relationships from \mathcal{K} to deduce class disjointness statements from existing class disjointness statements. This way we avoid posing redundant disjointness queries to the underlying LLM.

Proposition 2. *Let \mathcal{K} be a knowledge base and let C_1, C_2, C be classes of \mathcal{K} such that the following statements follow from \mathcal{K} : (i) C_1 and C_2 are disjoint, (ii) C is a subclass of C_1 , (iii) C is a subclass of C_2 . Then \mathcal{K} also entails that C is incoherent.*

Proof. Consider an arbitrary model \mathcal{I} of \mathcal{K} . According to the assumptions and in view of Definition 1, we know that (i) $\mathcal{C}_1^{\mathcal{I}} \cap \mathcal{C}_2^{\mathcal{I}} = \emptyset$, (ii) $\mathcal{C}^{\mathcal{I}} \subseteq \mathcal{C}_1^{\mathcal{I}}$, and (iii) $\mathcal{C}^{\mathcal{I}} \subseteq \mathcal{C}_2^{\mathcal{I}}$. We equivalently express (ii) and (iii) by (ii') $\mathcal{C}^{\mathcal{I}} = \mathcal{C}_1^{\mathcal{I}} \cap \mathcal{C}^{\mathcal{I}}$, and (iii') $\mathcal{C}^{\mathcal{I}} = \mathcal{C}_2^{\mathcal{I}} \cap \mathcal{C}^{\mathcal{I}}$. This allows us to infer $\mathcal{C}^{\mathcal{I}} = \mathcal{C}^{\mathcal{I}} \cap \mathcal{C}^{\mathcal{I}} = (\mathcal{C}_1^{\mathcal{I}} \cap \mathcal{C}^{\mathcal{I}}) \cap (\mathcal{C}_2^{\mathcal{I}} \cap \mathcal{C}^{\mathcal{I}}) = (\mathcal{C}_1^{\mathcal{I}} \cap \mathcal{C}_2^{\mathcal{I}}) \cap \mathcal{C}^{\mathcal{I}} = \emptyset \cap \mathcal{C}^{\mathcal{I}} = \emptyset$. \square

We exploited this property indirectly under the assumption that any named class C in the considered ontology is supposed to have instances – which seems to be a reasonable assumption since, otherwise, the definition of the class appears to be meaningless. In that case, any two classes that have a common subclass must be not disjoint.

Proposition 3. *Let K be a knowledge base, let C_1, C_2 be classes and let e be an individual of \mathcal{K} that such that the following statements follow from \mathcal{K} : (i) C_1 and C_2 are disjoint, (ii) e is an instance of C_1 , (iii) e is an instance of C_2 . Then \mathcal{K} is unsatisfiable.*

Proof. Suppose \mathcal{I} is a model of \mathcal{K} . According to the assumptions and in view of Definition 1, we know that (i) $\mathcal{C}_1^{\mathcal{I}} \cap \mathcal{C}_2^{\mathcal{I}} = \emptyset$, (ii) $e^{\mathcal{I}} \in \mathcal{C}_1^{\mathcal{I}}$, and (iii) $e^{\mathcal{I}} \in \mathcal{C}_2^{\mathcal{I}}$. Then, combining (ii) and (iii) we obtain $e^{\mathcal{I}} \in \mathcal{C}_1^{\mathcal{I}} \cap \mathcal{C}_2^{\mathcal{I}}$ and applying (i) yields $e^{\mathcal{I}} \in \emptyset$ which is a contradictory statement. Thus \mathcal{K} cannot have any models, which means it is unsatisfiable. \square

Again, this property can be exploited by noting that any two classes having common instances must not be disjoint. These considerations lead to the proposed methodology, detailed in Algorithm 1, which achieves the above-mentioned objective of producing an enriched knowledge base that is guaranteed to be contradiction-free, provided that the original knowledge base is.

Algorithm 2 achieves the objective of reducing the number of interactions with the LLM and maintaining satisfiability as new disjointness information is added through the LLM. The aim of producing an output that accurately reflects taxonomic relationships crucially depends on the quality and accuracy of the LLM’s responses. This, in turn, is influenced by both the LLM itself and the chosen prompting strategy. We focus on these issues in Section 5.

The last steps of Algorithm 2 (lines 14 and 15) are optional, but highly recommended, as they remove logically redundant statements from the disjointness-enriched knowledge base $\mathcal{K} \cup \mathcal{D}$. This yields a knowledge base that is logically equivalent but typically much smaller in size and hence both easier to process algorithmically and to scrutinize and maintain manually. Also, this “pruning step” is not computationally expensive, as it only requires $|\mathcal{D}^*|$ calls to a reasoner.

5. Experiments

In this section, we experiment with the approach proposed in Section 4 on the classes extracted from the DBpedia ontology. In particular, by relying on Algorithm 1, we obtain the list \mathcal{L} related to the DBpedia ontology. We have that $|\mathcal{L}| = 1148$, with 370 pairs labeled as disjoint and 778 pairs labeled as not *unknown*. In Table 2, we provide some examples of classes in \mathcal{L} .

Note that the list \mathcal{L} assumes that the ontology designers carefully produced a taxonomy that is intended to also reflect disjointness between classes. As shown in Section 3, however, this is not the case. The design of the taxonomy of DBpedia is structured such that disjoint axioms might result in unwanted inconsistencies. For this reason, we employ multiple metrics to evaluate the LLMs’ performances, each measuring a different behavior of the model. For

Algorithm 1 Determine the pair of disjoint classes derivable from \mathcal{K}

Input A knowledge base \mathcal{K} containing a class hierarchy. The knowledge base might or might not contain defined disjointness axioms and/or more complex axiomatizations.

Output A list \mathcal{L} of pair of classes such that disjointness statements logically following from \mathcal{K} are explicitly asserted.

```
1: Create a list  $\mathcal{L}$  of all pairs of classes  $(C_1, C_2) \mid C_1$  is lexicographically smaller than  $C_2$ .
2: Label all entries in  $\mathcal{L}$  as “unknown”.
3: for all  $D_1$  disjoint with  $D_2$  in  $\mathcal{K}$  do
4:   for all  $(C_1, C_2) \in \mathcal{L}$  do
5:     if  $(C_1 \sqsubseteq D_1 \wedge C_2 \sqsubseteq D_2)$  or  $(C_1 \sqsubseteq D_2 \wedge C_2 \sqsubseteq D_1)$  then
6:        $(C_1, C_2) \leftarrow$  “disjoint”
7:     end if
8:   end for
9: end for
10: for all  $(C_1, C_2) \in \mathcal{L} \mid (C_1, C_2) =$  “unknown” do
11:   if  $C_1$  and  $C_2$  have joint subclasses then
12:      $(C_1, C_2) \leftarrow$  “not disjoint”
13:   end if
14: end for
15: for all  $(C_1, C_2) \in \mathcal{L} \mid (C_1, C_2) =$  “unknown” do
16:   Query  $K$  for joint instances of  $C_1$  and  $C_2$ .
17:   if  $\exists e \in \mathbf{I} \mid e : C_1 \wedge e : C_2$  then
18:      $(C_1, C_2) \leftarrow$  “not disjoint”
19:   end if
20: end for
```

Table 2

Examples of disjointness between classes derived from Propositions 1, 2, and 3.

A	B	Disjoint	Reason
dbo:Chancellor	dbo:Species	×	Proposition 1
dbo:LatterDaySaint	dbo:Religious	×	Proposition 1
dbo:Person	dbo:Race	✓	Proposition 2
dbo:AcademicConference	dbo:Person	✓	Proposition 2

all metrics, a higher score indicates better performances, with 1 being the maximum score. In particular, disjoint recall (DR) measures how much the LLM aligns with humans by measuring the amount of *true* disjointness axioms that have been identified by the LLM. This measure provides an evaluation of the reliability of the prompt. Non-disjoint F1 (NDF1) measures the F1 score between the non-disjoint couples in L and the ones identified by the LLM. This provides a measure of how conservative the LLM is on its answers – i.e. how much the LLM acknowledges the open-world assumption. The F1 metrics measure the end-to-end performances of the model. The symmetric consistency metric (SC) measures how much the answers provided by the LLM

Algorithm 2 Determine the set of disjointness statements \mathcal{D} consistent with \mathcal{K}

Input A list \mathcal{L} containing pairs of classes labeled as “unknown”, “disjoint” or “not disjoint”; a prompt P for disjointness classification, with $LLM_P : \mathbf{C} \times \mathbf{C} \rightarrow \{\text{“disjoint”}, \text{“not disjoint”}\}$ the function that queries an LLM for disjointness of two classes using prompt P

Output A set \mathcal{D} of class disjointness axioms, such that all valid disjointness statements logically follow from $\mathcal{K} \cup \mathcal{D}$ and no invalid disjointness statements follow from it.

```
1: while  $\exists (D_1, D_2) \in \mathcal{L} \mid (D_1, D_2) = \text{“unknown”}$  do
2:   Select  $(D_1, D_2) \in \mathcal{L} \mid (D_1, D_2) = \text{“unknown”}$ 
3:    $d \leftarrow LLM_P(D_1, D_2)$ 
4:   if  $d = \text{“disjoint”}$  then
5:     for all  $(C_1, C_2) \in \mathcal{L} \mid (C_1 \sqsubseteq D_1 \wedge C_2 \sqsubseteq D_2) \vee (C_2 \sqsubseteq D_1 \wedge C_1 \sqsubseteq D_2)$  do
6:        $(C_1, C_2) \leftarrow \text{“disjoint”}$ 
7:     end for
8:   else
9:     for all  $(C_1, C_2) \in \mathcal{L} \mid (D_1 \sqsubseteq C_1 \wedge D_2 \sqsubseteq C_2) \vee (D_2 \sqsubseteq C_1 \wedge D_1 \sqsubseteq C_2)$  do
10:       $(C_1, C_2) \leftarrow \text{“not disjoint”}$ 
11:    end for
12:   end if
13: end while
14:  $\mathcal{D}^* \leftarrow \{\text{DisjointClasses}(C_1, C_2) \mid (C_1, C_2) \in \mathcal{L} \wedge (C_1, C_2) = \text{“disjoint”}\}$ 
15: Determine the minimal subset  $\mathcal{D}$  of  $\mathcal{D}^*$  such that  $\mathcal{K} \cup \mathcal{D}$  entails  $\mathcal{D}^*$ 
```

respect the symmetric property of the disjointness axiom – i.e. if A is disjoint from B then B is disjoint from A . Finally, we measure the overall accuracy of each model.

Prompting We adopt different prompting strategies: a naive approach, where the LLM has to autonomously understand the task, a task description approach, where the disjointness task is described and a few-shot approach that extends the task description by also providing some positive and negative examples. For each prompt, we frame the problem as a question-answering (QA) task, where the LLM has to answer positively or negatively to classify two classes as disjoint. To identify the best QA approach, we identify two prompts: (i) the LLM has to answer positively to classify two classes as disjoint and (ii) the LLM has to answer negatively. Table 3 describes the prompt templates we used. When possible, we rely on the instruction format of each LLM and use the *Prompting Strategy* template to instruct the LLM while we use the *QA Strategy* as a query to the instructed LLM.

5.1. Experimental setup

We perform our experiments on publicly available LLMs, to ensure full reproducibility of the experiments. For each LLM, we set the sampling temperature to 0, to reduce the randomness of the result. Moreover, we only rely on *small* LLMs – i.e. LLMs with approximately 8 billion of parameters. Through the use of proper optimization techniques, it is possible to run these models on consumer-level devices without the need for specialized hardware. We perform

Table 3

Prompting strategy templates. A and B are natural language labels of classes from the ontology.

Prompting Strategy	Template
Naive	Answer only “yes” or “no”.
Zero-shot Task Description	This is a question about ontological disjointness, answer only with “yes” or “no”
Few-shot Task Description	This is a question about ontological disjointness, answer only with “yes” or “no”. Examples of disjoint are: “person” and “file system”, “tower” and “person”, “place” and “agent”, “continent” and “sea”, “baseball league” and “bowling league”, “planet” and “star”. Examples of not disjoint are: “basketball player” and “baseball player”, “means of transportation” and “reptile”, “garden” and “historic place”, “president” and “beauty queen”, “castle” and “prison”.
QA Strategy	Template
Positive	Is the class A disjoint from B ?
Negative	Can a A be a B ?

our experiments on a selection of the current state-of-the-art models, including Mistral 0.3 7B [25], Gemma 2 9B [26], LLama 3 8B¹⁰, and Qwen 2 7B [27]¹¹. All experiments are run on 8-bit quantized models on an RTX3090 with 24GB of RAM. We experiment with each combination of the prompts of Table 3.

5.2. Results

The overall results are shown in Table 4. In general, LLMs achieve promising results in disjointness detection. Notably, the best prompting technique is not providing few-shot examples, but rather providing the LLM with little to no description of the task. Indeed, it has been observed how few-shot prompting is more effective when in-context learning is required, while zero-shot prompting is more effective when the implicit knowledge of the LLM should be exploited [28]. Nonetheless, further research on few-shot prompting for disjointness classification should be performed, as lower performances can also be attributed to the amount and nature of the examples we provide in the prompt. We manually select examples that are likely to provide meaningful disjointness instances. However, a more complex approach could be employed, such as exploiting Retrieval Augmented Generation (RAG) techniques to provide examples that are more likely to be relevant for the classes used as input. Different heuristics can be used to measure the relevance of other classes, such as word embeddings or knowledge graph embeddings. Interestingly, framing the problem as a negative QA task – i.e. asking whether an individual of a class can also be an instance of another class – consistently outperforms the positive QA prompt. This could be attributed to the fact that using the negative approach is

¹⁰<https://llama.meta.com/>

¹¹Due to their closed-source nature and high costs, we reserve the exploration of GPT-3.5 and GPT-4 for future work.

Table 4

Performance on disjointness detection for LLMs and prompt strategies. The best results for each prompt are underlined, while the best results overall in a metric are in **bold**.

Prompt	QA	LLM	DR	NDF1	F1	SC	Accuracy
Naive	Positive	Gemma 2	0.99	0.26	<u>0.53</u>	0.89	0.42
		LLama 3	0.19	0.63	0.17	0.65	0.37
		Mistral 0.3	1.00	0.03	0.49	0.98	0.33
		Qwen 2	0.00	0.99	0.01	0.96	<u>0.66</u>
	Negative	Gemma 2	0.71	<u>0.91</u>	0.69	0.85	0.79
		LLama 3	0.90	0.86	<u>0.74</u>	<u>0.89</u>	<u>0.80</u>
		Mistral 0.3	0.85	0.81	0.68	0.81	0.74
		Qwen 2	<u>0.92</u>	0.80	0.70	0.84	0.74
Task description	Positive	Gemma 2	0.99	0.35	<u>0.55</u>	0.86	0.47
		LLama 3	0.86	0.08	0.45	0.90	0.31
		Mistral 0.3	1.00	0.20	0.52	<u>0.91</u>	0.40
		Qwen 2	0.04	<u>0.82</u>	0.05	0.68	<u>0.49</u>
	Negative	Gemma 2	<u>0.97</u>	0.83	0.75	0.84	<u>0.79</u>
		LLama 3	0.98	0.78	0.71	<u>0.90</u>	0.75
		Mistral 0.3	0.85	0.76	0.64	0.76	0.69
		Qwen 2	0.96	0.61	0.61	0.72	0.60
Few shot	Positive	Gemma 2	0.90	0.49	<u>0.54</u>	0.83	0.51
		LLama 3	0.76	0.30	0.44	0.72	0.36
		Mistral 0.3	<u>0.95</u>	0.20	0.50	<u>0.85</u>	0.38
		Qwen 2	0.05	<u>0.87</u>	0.07	0.74	<u>0.54</u>
	Negative	Gemma 2	0.85	<u>0.89</u>	0.75	<u>0.81</u>	0.82
		LLama 3	<u>0.99</u>	0.54	0.59	0.77	0.57
		Mistral 0.3	0.74	0.86	0.65	0.79	0.75
		Qwen 2	0.98	0.37	0.54	0.71	0.47
Average among all prompts	Gemma 2	0.90	0.62	0.63	0.85	0.75	
	LLama 3	0.78	0.53	0.52	0.81	0.47	
	Mistral 0.3	0.90	0.48	0.58	0.85	0.49	
	Qwen 2	0.49	0.74	0.33	0.78	0.36	

more consistent with natural language questions. LLMs can actively exploit their pre-training phase, which generally includes a fine-tuning phase to solve QA tasks akin to our negative prompt. On average, Gemma 2 performs better than the other LLMs. However, depending on the requirements, other LLMs might be better suited. For instance, Mistral 0.3 is better aligned with human judgment, since it has a higher recall on disjointness axioms.

5.3. Disjointness on DBpedia

Given the results of Table 4, we consider Gemma 2 with task description prompt and a negative QA strategy as the most effective way of producing disjointness axioms among the methods

tested. We execute Algorithm 2 on the whole DBpedia ontology. We rely on a straightforward random selection for the pair (D_1, D_2) (line 3). In total, the algorithm takes $21589.75s \approx 6h$ to execute. Note that given the random selection, we are not able to exploit parallelism and query the LLM with single prompts. However, a selection strategy that enables parallel selection would greatly enhance the performance of the algorithm. In total, we find 510,600 disjointness axioms, which results in $\approx 98\%$ of the classes participating in at least one disjointness axiom. The number of axioms can be greatly reduced by relying on the “pruning” operation of Algorithm 2 (line 15). In the case of the DBpedia ontology, the number of resulting axioms is 170,122 – a reduction of $\approx 66\%$.

Table 5
Example of disjointness judgments retrieved by Algorithm 2

Class A	Class B	Disjoint
dbo:GeneLocation	dbo:HumanGene	×
dbo:VideogamesLeague	dbo:Website	×
dbo:InformationAppliance	dbo:MobilePhone	×
dbo:Engineer	dbo:Embryology	×
dbo:Identifier	dbo:District	×
dbo:Mosque	dbo:Museum	✓
dbo:MeanOfTransportation	dbo:Swimmer	✓
dbo:WikimediaTemplate	dbo:WomensTennisAssociationTournament	✓
dbo:Racecourse	dbo:Area	✓
dbo:PlayboyPlaymate	dbo:Camera	✓

For illustration and discussion purposes, Table 5 shows a non-representative selection of particularly discussion-worthy positive and negative disjointness statements retrieved via Algorithm 2. We observe that for some class relationships, including both common-sense and domain-specific classes, our approach resulted in the “conservative” misclassification of classes as non-disjoint, meaning the LLM classified the classes as non-disjoint despite the classes actually being disjoint. Examples include `dbo:VideogamesLeague` and `dbo:Website`, `dbo:GeneLocation` and `dbo:HumanGene`, and `dbo:Identifier` and `dbo:District`. Conversely, we also observed “aggressive” misclassification where our approach classified classes as disjoint despite them being really non-disjoint. Straightforward examples include `dbo:PlayboyPlaymate` and `dbo:Camera` or `dbo:WikimediaTemplate` and `dbo:WomensTennisAssociationTournament`, with a more complicated example being `dbo:Mosque` and `dbo:Museum`. The latter being disproven by a counter-example, the famous Mosque Hagia Sophia in Turkey¹². To address these misclassification instances, we suspect that providing more contextual information in the prompt may improve classification accuracy, especially for domain-specific scenarios. Also, future work could be done to assess how, through prompt design, the approach could encourage more “aggressive” or “conservative” disjointness classifications in scenarios where relationships are more uncertain.

¹²<https://muze.gen.tr/muze-detay/ayasofya>

6. Conclusion and Future Work

This work shows that LLMs can roughly identify and assert disjointness axioms in ontologies, with a different degree of reliability depending on the model. By harnessing their inherent background knowledge and employing strategic prompt engineering, we showed that these models can classify ontological disjointness with minimal human intervention. This capability simplifies ontology management and supports more robust reasoning in knowledge graphs. Our findings underscore the potential of LLMs as valuable tools for the automated enrichment of ontologies, which encourages future exploration and innovation in this domain.

Future works include testing the approach proposed in Section 4 on other ontologies, to assess its effectiveness on different types of ontologies, including domain-specific ontologies. Additionally, comprehensive validation by human domain experts would be required to obtain conclusive insights into the degree of reliability of the axioms asserted by the LLM.

Moreover, using different LLMs with different numbers of parameters and improving and expanding our strategies for testing disjointness constitutes interesting future work.

It could be worthwhile to look into heuristics for – given a large list of disjointness candidate pairs – picking those entries that are particularly “promising”. One option would be to follow the strong disjointness assumption [6] and pick “sibling classes”, that is, classes A and B that have a common direct superclass C . Furthermore, it could be interesting to test class pairs with just one or two examples of non-disjointness, as these instances may be errors to remove from the KG. On another note, one could develop strategies for gauging the reliability of an LLM response by rephrasing the question asked. This involves adding a description of classes in prompts to see if it improves the answers, relying on proper ontology serialization techniques [29]. Finally, using advanced prompting techniques, such as chain-of-thought, may improve the results alongside RAG techniques to pick the few-shot examples. Similarly, a richer prompt, including more qualifying phrases such as “at the same time” to check the temporality of the disjointness or “theoretically” to force abstraction might instruct the model toward a more effective framing of the problem.

Acknowledgments

Elias Crum acknowledges funding provided by VITO NV (UG_PhD_2303_contract). Antonio De Santis’s doctoral scholarship is funded by the Italian Ministry of University and Research (MUR) under the National Recovery and Resilience Plan (NRRP), by Thales Alenia Space, and by the European Union (EU) under the NextGenerationEU project. Alessia Pisu acknowledges MUR and EU-FSE for financial support of the PON Research and Innovation 2014-2020 (D.M. 1061/2021). Nicolas Lazzari has received funding from the FAIR – Future Artificial Intelligence Research Foundation as part of the grant agreement MUR n. 341. Sebastian Rudolph is funded by the Bundesministerium für Bildung und Forschung (BMBF, Federal Ministry of Education and Research) and DAAD (German Academic Exchange Service) in project 57616814 (SECAI, School of Embedded and Composite AI).

References

- [1] T. D. Wang, Gauging Ontologies and Schemas by Numbers, in: D. Vrandečić, M. C. Suárez-Figueroa, A. Gangemi, Y. Sure (Eds.), Proceedings of 4th International EON Workshop 2006 Evaluation of Ontologies for the Web Co-located with the WWW2006 Edinburgh, UK, May 22, 2006, volume 179 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2006. URL: <https://ceur-ws.org/Vol-179/eon2006wang.pdf>.
- [2] J. Völker, D. Vrandečić, Y. Sure, A. Hotho, Learning Disjointness, in: E. Franconi, M. Kifer, W. May (Eds.), The Semantic Web: Research and Applications, 4th European Semantic Web Conference, ESWC 2007, Innsbruck, Austria, June 3-7, 2007, Proceedings, volume 4519 of *Lecture Notes in Computer Science*, Springer, 2007, pp. 175–189. URL: https://doi.org/10.1007/978-3-540-72667-8_14. doi:10.1007/978-3-540-72667-8_14.
- [3] J. Völker, D. Fleischhacker, H. Stuckenschmidt, Automatic acquisition of class disjointness, *J. Web Semant.* 35 (2015) 124–139. URL: <https://doi.org/10.1016/j.websem.2015.07.001>. doi:10.1016/J.WEBSEM.2015.07.001.
- [4] G. Rizzo, C. d’Amato, N. Fanizzi, An unsupervised approach to disjointness learning based on terminological cluster trees, *Semantic Web* 12 (2021) 423–447. URL: <https://doi.org/10.3233/SW-200391>. doi:10.3233/SW-200391.
- [5] S. Schlobach, Debugging and Semantic Clarification by Pinpointing, in: A. Gómez-Pérez, J. Euzenat (Eds.), The Semantic Web: Research and Applications, Second European Semantic Web Conference, ESWC 2005, Heraklion, Crete, Greece, May 29 - June 1, 2005, Proceedings, volume 3532 of *Lecture Notes in Computer Science*, Springer, 2005, pp. 226–240. URL: https://doi.org/10.1007/11431053_16. doi:10.1007/11431053_16.
- [6] R. Cornet, A. Abu-Hanna, Usability of expressive description logics—a case study in UMLS, in: AMIA 2002, American Medical Informatics Association Annual Symposium, San Antonio, TX, USA, November 9-13, 2002, AMIA, 2002. URL: <https://knowledge.amia.org/amia-55142-a2002a-1.610020/t-001-1.612667/f-001-1.612668/a-036-1.613143/a-037-1.613140>.
- [7] J. Lehmann, DL-Learner: Learning Concepts in Description Logics, *J. Mach. Learn. Res.* 10 (2009) 2639–2642. URL: <https://dl.acm.org/doi/10.5555/1577069.1755874>. doi:10.5555/1577069.1755874.
- [8] B. P. Allen, L. Stork, P. Groth, Knowledge Engineering Using Large Language Models, *Transactions on Graph Data and Knowledge* 1 (2023) 3:1–3:19. URL: <https://drops.dagstuhl.de/entities/document/10.4230/TGDK.1.1.3>. doi:10.4230/TGDK.1.1.3.
- [9] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, X. Wu, Unifying large language models and knowledge graphs: A roadmap, *IEEE Transactions on Knowledge and Data Engineering* 36 (2024) 3580–3599. doi:10.1109/TKDE.2024.3352100.
- [10] A. D. Santis, M. Balduini, F. D. Santis, A. Proia, A. Leo, M. Brambilla, E. D. Valle, Integrating large language models and knowledge graphs for extraction and validation of textual test data, 2024. URL: <https://arxiv.org/abs/2408.01700>. arXiv:2408.01700.
- [11] F. Petroni, T. Rocktäschel, S. Riedel, P. Lewis, A. Bakhtin, Y. Wu, A. Miller, Language models as knowledge bases?, in: K. Inui, J. Jiang, V. Ng, X. Wan (Eds.), Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP),

Association for Computational Linguistics, Hong Kong, China, 2019, pp. 2463–2473. URL: <https://aclanthology.org/D19-1250>. doi:10.18653/v1/D19-1250.

- [12] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is All you Need, in: I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, December 4-9, 2017, Long Beach, CA, USA, 2017, pp. 5998–6008. URL: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- [13] J. D. M.-W. C. Kenton, L. K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [14] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language Models are Few-shot Learners, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020*, NeurIPS 2020, December 6-12, 2020, virtual, 2020. URL: <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.
- [15] R. OpenAI, Gpt-4 technical report. arxiv 2303.08774, View in Article 2 (2023).
- [16] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, G. Lample, LLaMA: Open and Efficient Foundation Language Models, CoRR abs/2302.13971 (2023). URL: <https://doi.org/10.48550/arXiv.2302.13971>. doi:10.48550/ARXIV.2302.13971. arXiv:2302.13971.
- [17] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. Canton-Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, T. Scialom, Llama 2: Open Foundation and Fine-tuned Chat Models, CoRR abs/2307.09288 (2023). URL: <https://doi.org/10.48550/arXiv.2307.09288>. doi:10.48550/ARXIV.2307.09288. arXiv:2307.09288.
- [18] M. Reid, N. Savinov, D. Teplyashin, D. Lepikhin, T. P. Lillicrap, J. Alayrac, R. Soricut, A. Lazaridou, O. Firat, J. Schrittwieser, I. Antonoglou, R. Anil, S. Borgeaud, A. M. Dai, K. Millican, E. Dyer, M. Glaese, T. Sottiaux, B. Lee, F. Viola, M. Reynolds, Y. Xu, J. Molloy, J. Chen, M. Isard, P. Barham, T. Hennigan, R. McIlroy, M. Johnson, J. Schalkwyk, E. Collins, E. Rutherford, E. Moreira, K. Ayoub, M. Goel, C. Meyer, G. Thornton, Z. Yang, H. Michalewski, Z. Abbas, N. Schucher, A. Anand, R. Ives, J. Keeling, K. Lenc, S. Haykal, S. Shakeri, P. Shyam, A. Chowdhery, R. Ring, S. Spencer, E. Sezener, et al., Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context, CoRR abs/2403.05530 (2024). URL: <https://doi.org/10.48550/arXiv.2403.05530>.

doi:10.48550/ARXIV.2403.05530. arXiv:2403.05530.

- [19] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. de Las Casas, E. B. Hanna, F. Bressand, G. Lengyel, G. Bour, G. Lample, L. R. Lavaud, L. Saulnier, M. Lachaux, P. Stock, S. Subramanian, S. Yang, S. Antoniak, T. L. Scao, T. Gervet, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, *Mixtral of Experts*, CoRR abs/2401.04088 (2024). URL: <https://doi.org/10.48550/arXiv.2401.04088>. doi:10.48550/ARXIV.2401.04088. arXiv:2401.04088.
- [20] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, J. Dean, S. Ghemawat, *Language Models are Unsupervised Multitask Learners*, in: OSDI'04: Sixth Symposium on Operating System Design and Implementation, 2018, pp. 137–150.
- [21] B. Chen, Z. Zhang, N. Langrené, S. Zhu, *Unleashing the potential of prompt engineering: a comprehensive review*, 2024. arXiv:2310.14735.
- [22] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, D. Zhou, *Chain-of-thought Prompting Elicits Reasoning in Large Language Models*, in: S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022, 2022*. URL: http://papers.nips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html.
- [23] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, Y. Iwasawa, *Large Language Models are Zero-shot Reasoners*, in: S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022, 2022*. URL: http://papers.nips.cc/paper_files/paper/2022/hash/8bb0d291acd4acf06ef112099c16f326-Abstract-Conference.html.
- [24] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, C. Bizer, *DBpedia - A large-scale, multilingual knowledge base extracted from Wikipedia*, *Semantic Web 6 (2015)* 167–195. URL: <https://doi.org/10.3233/SW-140134>. doi:10.3233/SW-140134.
- [25] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de Las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, *Mistral 7B*, CoRR abs/2310.06825 (2023). URL: <https://doi.org/10.48550/arXiv.2310.06825>. doi:10.48550/ARXIV.2310.06825. arXiv:2310.06825.
- [26] T. Mesnard, C. Hardin, R. Dadashi, S. Bhupatiraju, S. Pathak, L. Sifre, M. Rivière, M. S. Kale, J. Love, P. Tafti, L. Hussenot, A. Chowdhery, A. Roberts, A. Barua, A. Botev, A. Castro-Ros, A. Slone, A. Héliou, A. Tacchetti, A. Bulanova, A. Paterson, B. Tsai, B. Shahriari, C. L. Lan, C. A. Choquette-Choo, C. Crepy, D. Cer, D. Ippolito, D. Reid, E. Buchatskaya, E. Ni, E. Noland, G. Yan, G. Tucker, G. Muraru, G. Rozhdestvenskiy, H. Michalewski, I. Tenney, I. Grishchenko, J. Austin, J. Keeling, J. Labanowski, J. Lespiau, J. Stanway, J. Brennan, J. Chen, J. Ferret, J. Chiu, et al., *Gemma: Open Models Based on Gemini Research and Technology*, CoRR abs/2403.08295 (2024). URL: <https://doi.org/10.48550/arXiv.2403.08295>. doi:10.48550/ARXIV.2403.08295. arXiv:2403.08295.
- [27] A. Yang, B. Yang, B. Hui, B. Zheng, B. Yu, C. Zhou, C. Li, C. Li, D. Liu, F. Huang, et al.,

Qwen2 Technical Report, arXiv preprint arXiv:2407.10671 (2024).

- [28] L. Reynolds, K. McDonell, Prompt Programming for Large Language Models: Beyond the Few-shot Paradigm, in: Y. Kitamura, A. Quigley, K. Isbister, T. Igarashi (Eds.), CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama Japan, May 8-13, 2021, Extended Abstracts, ACM, 2021, pp. 314:1–314:7. URL: <https://doi.org/10.1145/3411763.3451760>. doi:10.1145/3411763.3451760.
- [29] C. Ringwald, F. Gandon, C. Faron, F. Michel, H. Abi Akl, 12 shades of RDF: Impact of Syntaxes on Data Extraction with Language Models, in: ESWC 2024 Extended Semantic Web Conference, May 2024, Hersonissos, Greece., 2024.