

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Information Processing and Management

journal homepage: www.elsevier.com/locate/ipm

Large language models for scholarly ontology generation: An extensive analysis in the engineering field

Tanay Aggarwal ^a, Angelo Salatino ^a, Francesco Osborne ^{a,b}, Enrico Motta ^a^a Knowledge Media Institute, The Open University, Milton Keynes, UK^b Department of Business and Law, University of Milano Bicocca, Milan, Italy

ARTICLE INFO

Keywords:

Large language models
Zero-shot learning
Ontology generation
Research topics
Scholarly knowledge
Scientific knowledge graphs

ABSTRACT

Ontologies of research topics are crucial for structuring scientific knowledge, enabling scientists to navigate vast amounts of research, and forming the backbone of intelligent systems such as search engines and recommendation systems. However, manual creation of these ontologies is expensive, slow, and often results in outdated and overly general representations. As a solution, researchers have been investigating ways to automate or semi-automate the process of generating these ontologies. One of the key challenges in this domain is accurately assessing the semantic relationships between pairs of research topics. This paper presents an analysis of the capabilities of large language models (LLMs) in identifying such relationships, with a specific focus on the field of engineering. To this end, we introduce a novel benchmark based on the IEEE Thesaurus for evaluating the task of identifying three types of semantic relations between pairs of topics: *broader*, *narrower*, and *same-as*. Our study evaluates the performance of seventeen LLMs, which differ in scale, accessibility (open vs. proprietary), and model type (full vs. quantised), while also assessing four zero-shot reasoning strategies. Several models with varying architectures and sizes have achieved excellent results on this task, including Mixtral-8x 7B, Dolphin-Mistral-7B, and Claude 3 Sonnet, with F1-scores of 0.847, 0.920, and 0.967, respectively. Furthermore, our findings demonstrate that smaller, quantised models, when optimised through prompt engineering, can achieve strong performance while requiring very limited computational resources.

1. Introduction

Ontologies of research topics, along with other knowledge organisation systems (KOS), such as taxonomies and thesauri, play a vital role in organising and retrieving information from digital libraries (Salatino, Aggarwal, Mannocci, Osborne, & Motta, 2024). Major publishers and organisations in this domain, such as Elsevier, Springer Nature, ACM, IEEE, and PubMed, employ a variety of classification systems, including the All Science Journal Classification (ASJC), SpringerNature Taxonomy, ACM Computing Classification System,¹ the IEEE Thesaurus, and Medical Subject Headings. These systems are crucial for enabling efficient information discovery and supporting effective scientific dissemination (Dunne & Hulek, 2020; Lipscomb, 2000; Rous, 2012). Furthermore, ontologies of research topics play a vital role in supporting a wide range of intelligent systems to navigate and interpret academic literature (Beel, Gipp, Langer, & Breiting, 2016; Gusenbauer & Haddaway, 2020; Osborne, Motta, &

* Corresponding author.

E-mail addresses: tanay.aggarwal@open.ac.uk (T. Aggarwal), angelo.salatino@open.ac.uk (A. Salatino), francesco.osborne@open.ac.uk (F. Osborne), enrico.motta@open.ac.uk (E. Motta).

¹ ACM Computing Classification System - <https://dl.acm.org/ccs>.

<https://doi.org/10.1016/j.ipm.2025.104262>

Received 30 August 2024; Received in revised form 6 June 2025; Accepted 6 June 2025

Available online 21 July 2025

0306-4573/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Mulholland, 2013). These include search engines (Gusenbauer & Haddaway, 2020), conversational agents (Meloni et al., 2023), analytics dashboards (Angioni, Salatino, Osborne, Recupero, & Motta, 2021), and recommender systems (Beel et al., 2016).

Nevertheless, generating and maintaining ontologies of research topics is both expensive and time-consuming, typically requiring the involvement of several domain experts over an extended period of time (Osborne & Motta, 2015). Moreover, the ever-increasing volume of literature, with approximately 2.5 million new scientific papers published annually (Bornmann & Mutz, 2015; Wang & Barabási, 2021), demands continuous updates to ensure that these ontologies accurately reflect the latest research developments. Over the past 25 years, Artificial Intelligence (AI) researchers have developed several methods for the automated or semi-automated generation of such ontologies, with varying degrees of success (Han, Yang, Mishra, & Diesner, 2020; Osborne & Motta, 2012, 2015; Sanderson & Croft, 1999). Some notable achievements include the semi-automatic pipeline for generating OpenAlex Topics (OpenAlex, 2024) and the automatic development of the Computer Science Ontology (Salatino et al., 2020), along with a few other initiatives. However, most of the ontologies in this domain are still manually created because current solutions still struggle to generate large-scale, detailed representations of research topics. Consequently, this challenge remains far from being fully addressed.

Over the past two years, Large Language Models (LLMs) have transformed the field of Natural Language Processing (NLP) by significantly improving machine comprehension of human language. In particular, some studies have demonstrated that LLMs can achieve remarkable performance even in zero-shot settings, where the model makes predictions without prior exposure to specific tasks or examples (Kojima, Gu, Reid, Matsuo, & Iwasawa, 2023).

In this paper, we explore the potential of LLMs to facilitate ontology generation by identifying semantic relationships between research topics, with a specific focus on the field of engineering. This task is a fundamental step in the construction of scientific KOSs, and has traditionally been the most labour-intensive phase, often involving extensive deliberation among domain experts (Osborne & Motta, 2015). Specifically, we analyse the capability of several state-of-the-art LLMs to detect three semantic relations frequently adopted in such ontologies: *broader*, *narrower*, and *same-as*. In particular, our experimental analysis focuses on identifying such relations between research topics defined in the IEEE Thesaurus, the most widely adopted ontology in electrical and electronic engineering, which also offers broad coverage of computer science. To support this investigation, we introduce IEEE-Rel-1K, a gold-standard dataset derived from the IEEE Thesaurus, which describes the semantic relationships of 1000 topic pairs.

We evaluated 17 off-the-shelf LLMs on this dataset, varying in scale, accessibility (open or proprietary), and model type (full or quantised). The models included 3 full open models, 10 open quantised models, and 4 proprietary models. For each model, we employed four distinct zero-shot prompting strategies. The first strategy involved a straightforward prompt asking the model to determine the relationship between two topics, t_a and t_b . The second strategy extended this by applying the same method to both the original pair (t_a, t_b) and the reversed pair (t_b, t_a), with the results then being combined. The third strategy utilised a Chain-of-Thought (CoT) prompt, encouraging the model to reason through the relationship step by step. The fourth strategy was a two-way version of this CoT approach, where the process is repeated for both the original and reversed pairs before integrating the outcomes.

Our findings show that several models exhibit strong zero-shot reasoning capabilities for this task. Proprietary models performed best, with Claude 3 Sonnet achieving an impressive F1 score of 0.967. Nonetheless, several open models also delivered highly competitive results, particularly when enhanced with advanced prompting techniques based on Chain-of-Thought (CoT) reasoning. Among the open models, the top performer was Dolphin-Mistral-7B (0.920 F1), a Mistral variant fine-tuned on the Dolphin dataset. The Mixtral model also demonstrated solid performance, reaching an F1 score of 0.847.

Our study produced several notable insights. First, many models exhibit strong zero-shot reasoning capabilities for this task. Second, the performance of LLMs is highly dependent on the prompting strategy employed. In particular, the two-way variant of CoT prompting consistently outperforms alternative approaches. In fact, our results show that prompting strategies can lead to substantial performance gains, with improvements exceeding 0.2 F1 points in several instances. Finally, when the optimal prompting strategy described in this paper is adopted, smaller models, which are more scalable and cost-efficient, can achieve performance comparable to that of the best-performing architectures. These findings point to promising directions for the development of optimised and efficient models for ontology generation.

Ultimately, our objective is to develop a detailed, comprehensive ontology of research topics across all domains. To achieve this, we plan to develop and publish field-specific ontologies using the SKOS standard,² subsequently mapping them together. This paper represents the first step in this ambitious direction.

In summary, the contributions of this paper are as follows:

- We conduct a comprehensive analysis of 17 recent LLMs in identifying semantic relationships between research topics in a zero-shot setting.
- We introduce IEEE-Rel-1K, a gold standard for research topic relation prediction, designed to facilitate the evaluation of models in this domain.
- We explore this task under various conditions, examining different types of LLMs and evaluating the effectiveness of multiple prompting strategies.
- We provide the complete codebase for our methodologies, including the gold standard.³

² SKOS - <https://www.w3.org/TR/2009/REC-skos-reference-20090818/>.

³ Gold Standard and code for experiments can be found at <https://github.com/ImTanay/LLM-Automatic-Ontology-Generation>.

The remainder of this paper is structured as follows. Section 2 provides an overview of the literature. Section 3 defines the task, describes the gold standard, and presents the seventeen LLMs employed in this study. Section 4 describes the experiments and provides details on the implementation setup. Section 5 presents and discusses our results. Section 6 provides an analysis of the types of errors made by the models. Section 7 discusses the key findings and the limitations. Finally, in Section 8 we conclude the paper and provide future directions.

2. Related work

In this section, we review the relevant literature. We begin with a discussion on knowledge organisation systems (Section 2.1), then discuss ontologies of research areas (Section 2.2), and finally explore state-of-the-art methodologies for generating these ontologies (Section 2.3).

2.1. Knowledge organisation systems

Knowledge Organisation Systems (KOSs) are frameworks for structuring information to facilitate efficient knowledge management and retrieval (Zeng, 2008). These systems, varying in structure and functionality (e.g., hierarchical relationships, synonym control), are commonly categorised into four primary types: term lists, taxonomies, thesauri, and ontologies.

A *term list* constitutes a linear, non-hierarchical compilation of subject headings or descriptors, facilitating the organisation of document collections (Hedden, 2010; Zaharee, 2013). In literature, they are alternatively addressed as *pick list* (Guidelines for the construction, 2010) or as *terminology* (Information and documentation, 2013). These KOSs are distinguished by the absence of explicit relational definitions between subjects. A *taxonomy* arranges classes within a hierarchical framework, with parent–child relationships (Rasch, 1987). They are organised in tree structures, featuring a root node, which subsequently unfolds into sub-branches. A *thesaurus* extends the taxonomy structure in which subjects are still organised hierarchically and are further characterised by descriptive attributes, including definitions, related terms, and synonyms (Guidelines for the construction, 2010). Finally, an *ontology* is a formal and explicit representation of domain knowledge, classifying entities based on their relevant attributes (Gruber, 1993). In practice, ontologies include concepts, objects, and their relations, and represent the most functionally comprehensive type of KOSs (Genesereth & Nilsson, 2012). They offer advanced capabilities, including synonym control, establishment of diverse relationship types, and definition of object properties (Zeng, 2008).

In the context of this work, an ontology of research areas applies this concept to the scholarly domain, systematically organising the various fields of study, their sub-disciplines, and their connections.

2.2. Ontologies of research areas

In the literature we can find several ontologies of research areas, each differing in scope, size, depth, and curation. Notable examples include the ACM Computing Classification System, Medical Subject Headings, IEEE Thesaurus, and UNESCO Thesaurus (Lipscomb, 2000; OpenAlex, 2024; Rous, 2012; Salatino, Thanapalasingam, Mannocci, Osborne, & Motta, 2018).

The Medical Subject Headings (MeSH) is a single-field ontology with over 30K concepts in the field of Medicine (Lipscomb, 2000). It is maintained by the National Library of Medicine (US), and mostly employed to organised content in the Medline digital library. It receives yearly updates with minor amendments to reflect the latest developments in the field.

The field of Computer Science is represented by the ACM Computing Classification System⁴ (Rous, 2012) and the Computer Science Ontology (CSO) (Salatino et al., 2018). The former contains about 2K research topics and it is maintained by the Association for Computing Machinery (ACM) for its digital library. CSO, on the other hand, offers a more comprehensive view with over 14,000 research concepts. CSO was automatically generated using Klink-2 (Osborne & Motta, 2015), starting from a corpus of 16 million scientific publications; and is being employed by Springer Nature to enhance the metadata quality of their computer science proceedings, including the Lecture Notes in Computer Science and Lecture Notes in Artificial Intelligence (Salatino, Osborne, Birukou, & Motta, 2019).

In the field of Engineering, there is the IEEE Thesaurus, which contains 12K terms. Curated by the Institute of Electrical and Electronics Engineers (IEEE), it is updated annually and serves as a basis for classifying academic articles and research within the IEEE digital library.

With regard to ontologies that cover multiple fields we can find the UNESCO Thesaurus. Curated by UNESCO, it covers 4.4K subjects and it is mainly employed for indexing and searching resources within the UNESCO's document repository. It undergoes frequent minor revisions, typically every 2–3 months. Similarly, the ANZSRC Fields of Research (FoR) developed by the Australia and New Zealand research councils, covers around 4.4K topics across a wide range of subjects, but it receives less frequent updates (~10 years). It is employed by Digital Science to organise research content within their main datasets: Dimensions.ai and Figshare.

OpenAlex Topics is another multi-field taxonomy with 4800 topics, primarily used within OpenAlex, a free and open catalogue of scholarly articles; and currently at its first release (OpenAlex, 2024).

A comprehensive survey on research topic ontologies is available in Salatino et al. (2024).

⁴ The ACM Computing Classification System – <http://www.acm.org/publications/class-2012>.

Despite their differences, these systems consistently employ hierarchical structures, using broader/narrower relations, and manage synonyms through ‘same-as’ or related links. Our work strategically focuses on these common relations, allowing us to build a robust framework for creating or enhancing existing ontologies. Furthermore, most existing systems are manually curated, and their manual development is labour-intensive, time-consuming, and costly (Salatino et al., 2024). Although automatic and semi-automatic generation methods have recently emerged, resulting in the creation of KOS such as CSO and OpenAlex Topics, these approaches have yet to fully exploit the potential of LLMs.

Finally, it is worth noting that several of these ontologies are also integrated in large-scale knowledge graphs (Peng, Xia, Naseriparsa, & Osborne, 2023), that describe the metadata of research publications (e.g., SemOpenAlex (Färber, Lamprecht, Krause, Aung, & Haase, 2023), AIDA-KG (Angioni et al., 2021)) or specific research entities and concepts (e.g., ORKG (Jaradeh et al., 2019), CS-KG (Dessí, Osborne, Reforgiato Recupero, Buscaldi, & Motta, 2022), Nano-publications (Kuhn et al., 2016)).

2.3. Generation of ontologies of research topics

Previous decades have seen the emergence of several automated solutions for ontology creation, yet these have primarily been applied outside the scholarly domain, explaining why scholarly ontologies often remain manually curated. This section is structured into three parts. First, we examine established methodologies for automatic ontology generation, with a focus on traditional natural language processing (NLP) techniques. Second, we broaden the scope to include the application of large language models (LLMs). Finally, we conclude by reviewing the currently limited applications of these techniques in the scholarly domain.

Early work focused on semi-automatic methods, combining human expertise with computational techniques. Maedche and Staab (2001) introduced a method using linguistic and statistical techniques to extract concepts and relationships from text, but heavily relied on human input. Cimiano and Völker (2005) advanced this with Text2Onto, integrating techniques from machine learning and natural language processing to automate extraction, though it still required substantial human validation. Fully automatic methods emerged with OntoLearn (Navigli, Velardi, Cucchiarelli, & Neri, 2004), which used lexical resources and statistical methods to reduce human oversight. The same team (Velardi, Faralli, & Navigli, 2013) later improved it with a graph-based algorithm in OntoLearn Reloaded, enhancing accuracy and robustness in complex domains. The integration of deep learning marked a significant advancement for the field of ontology generation. BERT (Devlin, Chang, Lee, & Toutanova, 2018) revolutionised language understanding, enhancing concept extraction and relationship identification. BERTopic (Grootendorst, 2022) leveraged embeddings for coherent topic generation, influencing ontology and taxonomy extraction. Le, Roller, Papaxanthos, Kiela, and Nickel (2019) use hyperbolic embeddings to capture the hierarchical structure of concepts. As recent advancements, Chen, Lin, and Klein (2020) employed pre-trained models for developing hierarchical taxonomies, obtaining improved results compared to state-of-the-art solutions. Some efforts have also focused on the automatic or semi-automatic evolution of ontologies (Canito, Corchado, & Marreiros, 2022; Osborne & Motta, 2018).

This foundation has been significantly advanced by recent research leveraging LLMs for ontology learning. Babaei Giglou, D’Souza, and Auer (2023) introduced LLMs4OL, demonstrating how LLMs can extract concepts, hierarchical relations, and non-taxonomic relations from text, addressing the limitations of previous statistical approaches. However, their methodology suffered from critical limitations, particularly in extracting non-taxonomic relations consistently and maintaining semantic coherence across different domains, which introduces potential bias and hallucination risks. Building upon this foundational work, Saeedizade and Blomqvist (2024) conducted a comprehensive examination of LLMs’ capabilities across the entire ontology development lifecycle. Their analysis revealed that while LLMs excel at ontology requirements gathering and conceptualisation, they still require substantial human oversight for formalisation tasks. Critically, they show that models struggle with maintaining precise ontological consistency and their inability to autonomously validate ontological structures without human intervention. Kommineni, König-Ries, and Samuel (2024) proposed an innovative framework combining human expertise with LLM capabilities for knowledge graph construction. Their approach demonstrated significant efficiency improvements, reducing expert time requirements while maintaining quality. Nevertheless, the study highlighted consistent challenges, including performance inconsistencies across different knowledge domain contexts and potential semantic drift during automated graph generation. Zhang et al. (2024) explored the integration of conversational interfaces into ontology engineering and introduced OntoChat, a framework that enables domain experts to construct ontologies through natural language dialogue. While innovative, the approach revealed significant limitations, particularly in scalability for complex ontologies and the introduction of potential communication ambiguities that could compromise ontological precision. Sun et al. (2024) systematically evaluated the potential of LLMs to replace traditional taxonomies. Their comprehensive analysis revealed that LLMs demonstrate outstanding capabilities, but they cannot consistently replace curated taxonomies. Moreover, they found significant variations in performance across domain specialisations and challenges in maintaining long-term consistency.

While LLMs offer clear advantages over traditional ontology generation methods, particularly in concept extraction and the identification of hierarchical relationships, producing comprehensive and detailed ontologies remains an open challenge. However, for this work we hypothesise that carefully designed prompting strategies, combined with LLMs’ advanced language processing capabilities, can yield consistent and domain-specific modelisations.

With regard to generating ontologies in the scholarly domain, research endeavours have explored a variety of (semi-)automated methods. For instance, Shen, Ma, and Wang (2018) developed a method to discover research concepts and arrange them hierarchically. Their resulting extensive taxonomy was then implemented in Microsoft Academic Graph. Another notable method is Klink-2 (Osborne & Motta, 2015), which extracts large-scale ontologies of research topics and is currently used to generate CSO. The algorithm starts with a set of keywords and analyses their relationship based on the extent of their co-occurrence. It infers

semantic relationships between keyword pairs by utilising a modified subsumption method combined with temporal information to assess hierarchical relationships, and topic similarity strategies to identify keywords referring to the same subject. Finally, the OpenAlex team expanded the ASJC structure found in Scopus by incorporating over 4.5K new topics identified through citation clustering (OpenAlex, 2024). This involved clustering citation networks to identify coherent and thematically related groups of papers, using LLMs to label these groups, and then connecting the newly created concepts to Scopus' ASJC.

In conclusion, LLMs are increasingly being used for ontology generation, as well as a range of research-related tasks such as identifying relevant papers (Cadeddu et al., 2024; Khanal, Prasad, Alsadoon, & Maag, 2020), assisting with or even automating the creation of literature reviews (Bolanos, Salatino, Osborne, & Motta, 2024), and enhancing academic writing and referencing (Brody, 2021; Buscaldi et al., 2024). However, their application to research topic ontologies remains significantly underexplored. In this paper, we address this issue by conducting a thorough evaluation of a broad range of recent LLMs on the task of identifying relationships between scientific topics.

3. Background

This section defines the task under investigation (Section 3.1), describes the gold standard developed for this study (Section 3.2), and presents an overview of the LLMs used in the experiments (Section 3.3).

3.1. Task definition

This paper addresses the challenge of identifying semantic relationships between pairs of research topics, denoted in the following as t_A and t_B . We formalise this task as a single-label, multi-class classification problem, where each topic pair is assigned one of the following categories:

- *broader*: t_A is a parent topic of t_B . E.g., *machine learning* is a broader area than *deep neural networks*;
- *narrower*: t_A is a child topic of t_B . E.g., *distributed databases* is a specific area within *databases*. This is the inverse relationship of *broader*;
- *same-as*: t_A and t_B can be used interchangeably to refer to same concept. E.g., *haptic interface* and *haptic device*. This relationship is symmetric;
- *other*: t_A and t_B do not relate according to the above categories. E.g., *software algorithms* and *cyclones*. This category does not represent a semantic relation; rather, it serves to provide the classifier with a mechanism for labelling negative examples.

We focused our extraction on the broader, narrower, and same-as relations, as these constitute a fundamental set for constructing a Knowledge Organisation System (KOS) that effectively represents scientific disciplines (Smith, 2020). As outlined in Sections 2.1 and 2.2, these relations play a critical role in defining hierarchical structures and managing synonymy.

3.2. Gold standard

To evaluate the performance of LLMs on this task, we created a gold standard by selecting a sample of 1000 semantic relationships from the IEEE Thesaurus. We have named this dataset IEEE-Rel-1K (IEEE Relations). The IEEE thesaurus contains around 11,730 engineering, technical, and scientific terms, including IEEE-specific society terms. It is compiled from IEEE transactions, journal articles, conference papers, and standards, and reflects the vocabulary used within the engineering and scientific fields of IEEE. Specifically, we utilised the IEEE Thesaurus v1.02,⁵ which is dated back to July 2023, and is available as a PDF following the ANSI/NISO Z39.4-2021 standard (Ansi/niso z39.4-2021, 2021).

To create IEEE-Rel-1K, we developed a Python script⁶ for extracting the hierarchical structure and relationships between terms from the IEEE Thesaurus PDF file. Next, we randomly selected 250 relationships for *broader* and *narrower*. For *same-as*, we used a combination of 'use preferred term' and 'used for' relationships as defined in the ANSI/NISO Z39.19-2005 standard (Guidelines for the construction, 2010). While the IEEE Thesaurus offers numerous related topics, these do not consistently meet the strict definition of synonymy (Guidelines for the construction, 2010). For instance, '4G mobile communication' is related to '5G mobile communication' but while they belong to the same category of mobile networks, they represent entirely different technologies. Therefore, a team of three experts manually validated and selected only synonyms, lexical variants, and near-synonyms. This manual validation was time-consuming, limiting our dataset to 250 samples per category. To complete the dataset, we randomly generated 250 pairs of topics, ensuring they were not semantically related within the IEEE Thesaurus, and labelled these as *other*.

We released IEEE-Rel-1K within the GitHub repository: <https://github.com/ImTanay/LLM-Automatic-Ontology-Generation>.

⁵ A copy of version 1.02 of the IEEE Thesaurus is available at https://github.com/angelosalatino/ieee-taxonomy-thesaurus-rdf/blob/main/source/ieee-thesaurus_2023.pdf.

⁶ Script for extracting semantic relationships in IEEE Thesaurus - <https://github.com/angelosalatino/ieee-taxonomy-thesaurus-rdf>.

3.3. Large language models

Our experiments included a diverse set of seventeen LLMs, all transformer-based and utilising a decoder-only architecture. However, these models exhibit variations in training parameters, quantisation levels, openness, and fine-tuning methodologies. We categorised these models into two primary groups: open (13) and proprietary (4). This distinction was necessary because we lack comprehensive information about proprietary models in terms of their internal structures and any supplementary systems that may support them. As a result, a direct comparison between open and proprietary models would not be entirely fair.

We further divided the 13 open models into two groups: 3 full models and 10 quantised models, the latter being quantised to 8 bits. The quantised models were employed to assess the performance of more computationally efficient solutions. Since some quantised models have been fine-tuned on specific datasets, they may outperform the original models from which they were derived in certain tasks.

Detailed descriptions of each model are provided below.

3.3.1. Open models

Mistral-7B (shortened as *mistral*) is trained on Open-Web data, has 7 billion parameters, and utilises grouped-query attention and sliding window attention with a context window of 32,000 tokens (Jiang et al., 2023).

Mixtral-8 × 7B (shortened as *mixtral*) is a sparse mixture-of-experts network incorporating 46.7 billion parameters. Only 12.9 billion parameters are used per token, whereas the remaining parameters are part of the “expert” layers within its sparse mixture-of-experts architecture. It has a context window of 32,000 tokens (Jiang et al., 2024).

Llama-2-70B (shortened as *llama-2*) has 70 billion parameters and a context window of 4096 tokens. It is trained on a corpus that includes 2 trillion tokens of data from publicly available sources (Touvron et al., 2023).

3.3.2. Quantised models

Dolphin-2.1-Mistral-7B (shortened as *dolphin-mistral*⁷) is a model with 7 billion parameters and a token context capacity of 4096. It is based on Mistral-7B and fine-tuned with the Dolphin⁸ dataset, which is an open-source implementation of Microsoft’s Orca (Mukherjee et al., 2023), with an addition of Airoboros⁹ dataset.

Dolphin-2.6-Mistral-7B-dpo-laser (shortened as *dolphin-mistral-dpo*¹⁰) is based on Mistral-7B with 7 billion parameters and fine-tuned on top of Dolphin DPO using Layer Selective Rank Reduction (LASER) (Sharma, Ash, & Misra, 2023). It offers a context window of 4096 tokens.

Dolphin2.1-OpenOrca-7B (shortened as *dolphin-openorca*¹¹) is a model that blends Dolphin-2.1-Mistral-7B and Mistral-7B-OpenOrca models. These models were merged using the “ties merge” (Yadav, Tam, Choshen, Raffel, & Bansal, 2024) technique, keeping the same number of training parameters and token context window size, respectively 7 billion and 4096.

OpenChat-3.5-1210 (shortened as *openchat*¹²) is a model fine-tuned on top of *Mistral-7B-v0.1*. It possesses similar properties and consists of 7 billion parameters with a context window size of 8192 tokens (Wang et al., 2023).

OpenChat-3.5-0106-Gemma (shortened as *openchat-gemma*¹³) is a model trained on *openchat-3.5-0106*¹⁴ data using Conditioned Reinforcement Learning Fine-Tuning (C-RLFT) framework (de Bruin, Kober, Tuyls, & Babuška, 2020; Wang et al., 2023). This model shares the same properties of the *openchat-3.5-0106* model, which was fine-tuned¹⁵ on top of Mistral-7B. It consists of 7 billion parameters with a context window of 8192 tokens.

SOLAR-10.7B-Instruct-v1.0 (shortened as *solar*¹⁶) a 10.7 billion parameters model with a Depth Up-Scaling (DUS) architecture that includes architectural modifications and continued pretraining (Kim et al., 2023). It is trained on the OpenOrca¹⁷ dataset.

Mistral-7B-OpenOrca (shortened as *mistral-openorca*¹⁸) is based on *Mistral-7B-Instruct-v0.1* and has been fine-tuned using the OpenOrca dataset (Mukherjee et al., 2023). This dataset is a comprehensive collection of augmented Fine-tuned LLanguage Net (FLAN) data aligned with the distributions outlined in the Orca paper (Lian et al., 2023; Longpre et al., 2023). The model has 7 billion parameters and a context window of 4096.

Eurus-7b-sft (shortened as *eurus*¹⁹) is fine-tuned on top of *Mistral-7B-Instruct-v0.2* using all correct actions in UltraInteract, mixing a small proportion of UltraChat, ShareGPT, and OpenOrca examples, and possess 7 billion parameters with a context window size of 4096 (Yuan et al., 2024).

⁷ Dolphin-2.1-Mistral-7B - <https://huggingface.co/TheBloke/dolphin-2.1-mistral-7B-GGUF>.

⁸ Dolphin Dataset - <https://huggingface.co/datasets/cognitivecomputations/dolphin>.

⁹ Airoboros - <https://huggingface.co/datasets/jondurbin/airoboros-2.1>.

¹⁰ Dolphin-2.6-Mistral-7B-dpo-laser - <https://huggingface.co/TheBloke/dolphin-2.6-mistral-7B-dpo-laser-GGUF>.

¹¹ Dolphin2.1-OpenOrca-7B - <https://huggingface.co/TheBloke/Dolphin2.1-OpenOrca-7B-GGUF>.

¹² OpenChat-3.5-1210 - <https://huggingface.co/TheBloke/openchat-3.5-1210-GGUF>.

¹³ OpenChat-3.5-0106-Gemma - <https://huggingface.co/gguf/openchat-3.5-0106-gemma-GGUF>.

¹⁴ openchat-3.5-0106 - <https://huggingface.co/openchat/openchat-3.5-0106>.

¹⁵ <https://huggingface.co/openchat/openchat-3.5-1210/discussions/4#658288f1168803bdee13d6b3>.

¹⁶ SOLAR-10.7B-Instruct-v1.0 - <https://huggingface.co/TheBloke/SOLAR-10.7B-Instruct-v1.0-GGUF>.

¹⁷ OpenOrca dataset - <https://huggingface.co/datasets/Open-Orca/OpenOrca>.

¹⁸ Mistral-7B-OpenOrca - <https://huggingface.co/Open-Orca/Mistral-7B-OpenOrca>.

¹⁹ Eurus-7b-sft - <https://huggingface.co/mradermacher/Eurus-7b-sft-GGUF>.

Llama-3-8B-Instruct (shortened as `llama-320`) has 8 billion parameters and a context length of almost 8192 tokens (AI@Meta, 2024). It utilises supervised fine-tuning (SFT) and reinforcement learning with human feedback (RLHF).

Orca-2-13B (shortened as `orca-221`) is a model with 13 billion parameters and a 2048-token context window. It was trained on a mix of publicly available and curated datasets to enhance its reasoning and comprehension abilities. The fine-tuning process was based on datasets like Super-NaturalInstructions and FLAN to ensure strong performance across NLP tasks (Mitra et al., 2023).

3.3.3. Proprietary models

We considered four proprietary models. Due to undisclosed details, we can only provide limited information. For example, their number of parameters is not officially disclosed.

GPT-3.5-Turbo-Instruct (shortened as `gpt-3.5`) is trained using RLHF, following similar methods as *InstructGPT*, but with some differences in the data collection setup. It has a context window size of 4096 tokens and was trained with data up to September 2021 (Ye et al., 2023).

GPT-4-Turbo (shortened as `gpt-4`) has been pre-trained to predict the next token in a document. It was trained using publicly available data and data licensed from third-party providers up to December 2023. The model was then fine-tuned using RLHF. It has a context window of 128,000 tokens (OpenAI, 2023).

Claude 3 Haiku (shortened as `haiku`) is a model with a context size of 200,000 tokens, offering a maximum output of 4096 tokens with training data up to August 2023 (Anthropic, 2024).

Claude 3 Sonnet (shortened as `sonnet`) offers a context size of 200,000 tokens. It has the capability to generate a maximum output of 4096 tokens and has been trained on data up to August 2023. While `haiku` is optimised for speed and cost-effectiveness in applications requiring immediate responses, `sonnet` features a larger model architecture giving it advantage in handling complex tasks that require in-depth understanding of language (Anthropic, 2024).

4. Experiments

We evaluated the 17 models described in Section 3.3 against IEEE-Rel-1K, the gold standard defined in Section 3.2, employing four different experiments that varied based on the prompting strategy utilised.

We investigated two primary prompting paradigms: a standard prompting approach based on a simple prompt that directly asks the model to perform the classification and Chain-of-Thought (CoT) prompting (Wei et al., 2023), which encourages the model to reason through a predefined sequence of steps before providing an answer. For each of these two paradigms, we employed two additional strategies, resulting in four distinct experimental settings. The first, labelled *one-way*, involves the model inferring the relationship between t_A and t_B directly. The second, referred to as *two-way*, builds upon the first by having the model initially infer the relationship between t_A and t_B , then infer the relationship between t_B and t_A , and finally combine the results using heuristic rules.

Fig. 1 illustrates the flowcharts of these experiments.

Based on these strategies, we can characterise our four experiments as follows:

1. standard prompting, one-way strategy (left branch of Fig. 1(a));
2. standard prompting, two-way strategy (both branches of Fig. 1(a));
3. chain-of-thought prompting, one-way strategy (left branch of Fig. 1(b));
4. chain-of-thought prompting, two-way strategy (both branches of Fig. 1(b)).

In all experiments, the performance of the model was assessed using standard precision, recall and F1.

The standard prompting, one-way strategy serves as a baseline for classification tasks. We introduced CoT prompting due to its demonstrated effectiveness in complex reasoning tasks (Kojima, Gu, Reid, Matsuo, & Iwasawa, 2022; Wei et al., 2023). Indeed, as discussed in Section 5, CoT prompting led to improved performance for most models. We investigated the two-way validation strategy to identify and rectify potential inconsistencies. It draws inspiration from prior research showing that approaches to improve logical consistency can enhance the overall performance of LLMs (Liu, Guo et al., 2024). This approach leverages the inverse relationship between *broader* and *narrower*, as well as the symmetric nature of the *same-as* and *other* relationships (Guidelines for the construction, 2010). Conflicting answers suggest inconsistencies that may be addressed using heuristic methods.

In the following subsections, we provide a detailed description of each strategy, followed by a discussion of the experimental setup and implementation details.

4.1. Standard prompting

In the experiments using standard prompting (Fig. 1(a)), we generated a prompt for each pair of research topics from the gold standard using a predefined template. The prompt template (detailed in Appendix A) describes the task and clearly defines the four categories, providing one generic example for each. Additionally, the template specifies that the answer should be provided in a numeric code format to facilitate parsing and ensure unambiguous responses. The resulting prompts were then submitted to the LLM to obtain an answer for each pair of topics. To ensure a fair comparison, all models received the same prompt.

²⁰ Llama-3 - <https://huggingface.co/lmstudio-community/Meta-Llama-3-8B-Instruct-GGUF>.

²¹ Orca-2-13B - <https://huggingface.co/TheBloke/Orca-2-13B-GGUF>.

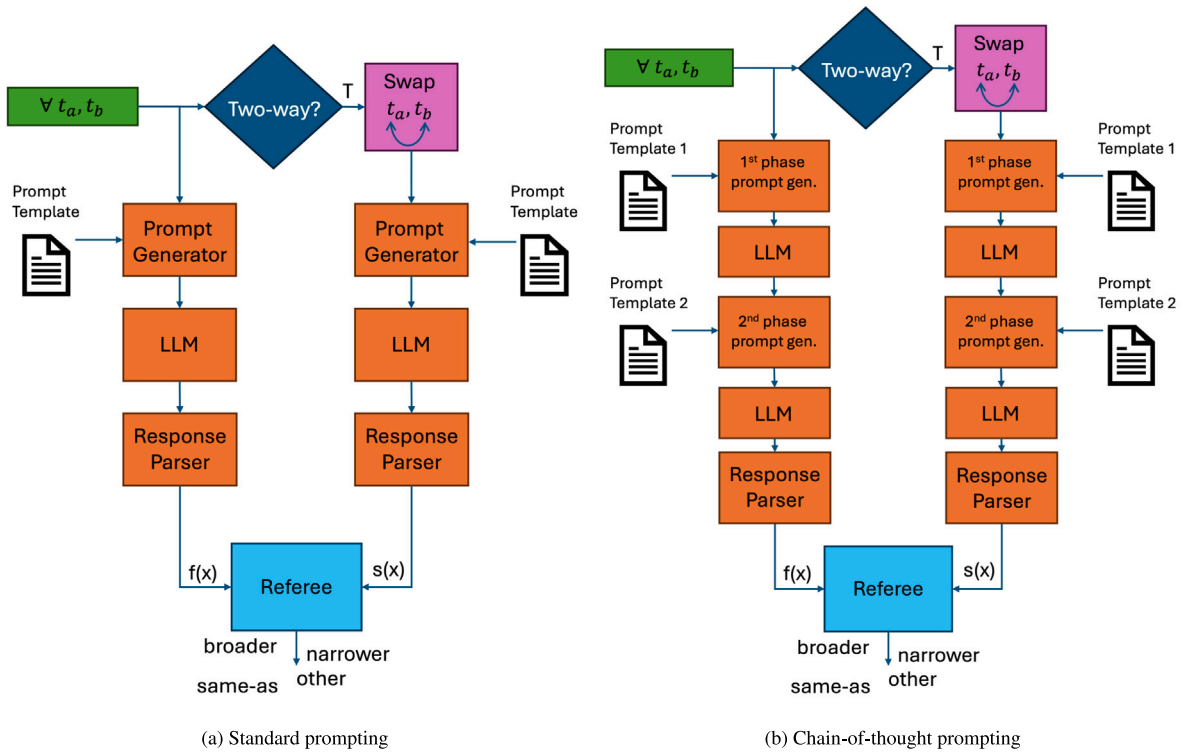


Fig. 1. Flowcharts of our experiments. On the left there is the flowchart of our standard prompting experiments. On the right is the flowchart of our chain-of-thought prompting experiments. All boxes except the orange are similar across experiments.

4.2. Chain-of-thought prompting

CoT prompting experiments operated similarly to the previous solution but involved a two-phase interaction with the LLM, as illustrated in Fig. 1(b). In the first phase, the model is asked to provide a definition of both topics, formulate a sentence that incorporates both of them, and finally discuss their potential semantic relationships. The resulting response is then fed back into the model in the second phase, with the sole task of determining the type of semantic relationship. The two templates used to produce the prompts for the CoT experiments are available in Appendix B.

4.3. One-way vs. two-way strategies

The experiments for both standard prompting and CoT prompting are conducted in two different configurations. The one-way strategy simply asks for the relationship between t_A and t_B . The two-way strategy involves asking the LLM to identify the relationship between topics t_A and t_B , as well as the relationship between topics t_B and t_A , where the order of the topics is reversed. This approach can naturally produce inconsistent responses, which may indicate potential issues.

In order to detect and mitigate logical inconsistencies, we designed a set of empirical rules (cyan boxes as referees in both Figs. 1(a) and 1(b)). The rules were designed to prioritise the development of a hierarchical taxonomy.

Considering $f(X)$ and $s(X)$ as the relationship types respectively obtained via the first and second branches of the two-way strategy; and $len(t_A)$ as the length of the topic's surface form, the rules are established as follows:

1. broader :- $f(\text{broader}) \wedge s(\text{narrower})$
2. narrower :- $f(\text{narrower}) \wedge s(\text{broader})$
3. broader :- $((f(\text{narrower}) \wedge s(\text{narrower})) \vee (f(\text{broader}) \wedge s(\text{broader}))) \wedge len(t_A) \leq len(t_B)$
4. narrower :- $((f(\text{narrower}) \wedge s(\text{narrower})) \vee (f(\text{broader}) \wedge s(\text{broader}))) \wedge len(t_A) > len(t_B)$
5. same-as :- $f(\text{same-as}) \wedge s(\text{same-as})$
6. broader :- $(f(\text{broader}) \wedge s(\text{other})) \vee (f(\text{other}) \wedge s(\text{narrower}))$
7. narrower :- $(f(\text{narrower}) \wedge s(\text{other})) \vee (f(\text{other}) \wedge s(\text{broader}))$
8. :- $f(X)$

4.4. Experimental setup

To interact with the various LLMs we employed two services, e.g., Amazon Bedrock,²² and OpenAI API²³; as well as the KoboldAI²⁴ library for running LLMs locally.

Specifically, through **Amazon Bedrock**, we interacted with MistralAI family LLMs (e.g., Mistral-7B and Mixtral-8 × 7B), Anthropic models (e.g., Claude 3 Haiku, and Claude 3 Sonnet), and Llama-2-70B by Meta. This facility provides a pre-structured generic wrapper, which encapsulates the interaction with the different LLMs offered by Amazon Bedrock.

With the **OpenAI** API we interacted with GPT models, such as GPT3.5-Turbo-Instruct and GPT4-Turbo.

Finally, we employed **KoboldAI** to interface with Dolphin-2.1-Mistral-7B, Mistral-7B-OpenOrca, Dolphin-2.6-Mistral-7B-dpo-laser, Dolphin-2.1-OpenOrca-7B, SOLAR-10.7B-Instruct-v1.0, OpenChat-3.5-1210, OpenChat-3.5-0106-Gemma, EURUS-7B-SFT. KoboldAI is an open-source tool built on llama.cpp and provides model access through an API endpoint. We run it on a Google Colaboratory instance equipped with Nvidia's V100 and L4 GPUs.

Within the GitHub repository, we report the various parameters we set when interacting with the models across the three services.

5. Results and discussion

This section presents the results of the complete set of experiments described in Section 4. First we will discuss the standard prompting experiments (see Fig. 1(a) for flowchart) and then we will report the results of the CoT prompting experiments (see Fig. 1(b)). Further, we will compare the results across the various strategies.

5.1. Standard prompting, one-way strategy

Table 1 reports the results of the experiments with standard prompting when applying one-way strategy. Among the full open models (first three rows), llama-2 and mixtral demonstrate good performance in terms of average F1 with 0.669 and 0.779, respectively. In particular, mixtral shows a high value of precision with the broader (0.819) and same-as (1), whereas llama-2 obtains high values of precision with narrower (1) and other (0.877). Complementarily, mixtral reaches high values of recall for narrower (0.888) and other (0.940), whereas llama-2 for broader (0.872) and same-as (0.992). In this case, mistral consistently performs low, except for the recall of the relationship other (0.992). The observed performance differences may be partially attributed to variations in the pre-training datasets. Indeed, mixtral (Jiang et al., 2023) was trained on a more diverse corpus that included scientific and technical documents (Jiang et al., 2024). This broader exposure may have enhanced its ability to model the relationships between scientific domains, resulting in superior performance in identifying hierarchical structures. In contrast, llama-2 (Touvron et al., 2023) was pre-trained primarily on web data.

Among quantised models, dolphin-openorca and openchat stand out with the highest average F1-score of 0.724 and 0.705, respectively. While both consistently excel in precision, they face some challenges in recall for narrower and same-as relations. Other quantised models, although falling below 0.7 in average F1-score, still demonstrate remarkable performance in specific relationships. For instance, mistral-openorca excels in precision for the same-as (1) and other (0.992) relations and recall for narrower (0.996), despite lower performance in all other metrics. The varied performance among models with the same base architecture but different fine-tuning approaches highlights the critical role of instruction-tuning datasets (Zhang et al., 2023). dolphin-openorca and openchat were both fine-tuned on datasets that emphasise following detailed instructions (Mukherjee et al., 2023; Wang et al., 2023), which may make them more suitable for our specific task.

When evaluating the proprietary models (as shown in the last four rows of **Table 1**), sonnet stands out, achieving an outstanding average F1-score of 0.967, with a well-balanced precision and recall across all relationship types. gpt-4 also demonstrates strong performance (0.948 F1). These two models exhibit high precision for same-as and other relations, and high recall for broader relations. In contrast, the two other models do not perform as well. Haiku, with an F1-score of 0.712, is surpassed by both mixtral and dolphin-openorca. Surprisingly, gpt-3.5 also shows lower-than-expected results. The exceptional performance of sonnet and gpt-4 is most likely attributable to both their size and their pre-training on datasets that include extensive scientific and technical corpora (Anthropic, 2024; OpenAI, 2023). While the specific composition of their training data remains proprietary, these models have demonstrated superior capabilities on several tasks that require a nuanced understanding of scientific domains (Bubeck et al., 2023). Naturally, both model size and recency play a significant role, as evidenced by the superior performance of the most recent and largest proprietary models (gpt-4, sonnet) compared to their earlier and smaller counterparts (gpt-3.5, haiku).

Overall, we observe a clear performance gap between full open models and proprietary models. For example, sonnet achieves an average F1-Score of 0.967, substantially outperforming mixtral which has an average F1-Score of 0.779.

While none of the models explicitly claim to have been trained on the IEEE digital library specifically, the superior performance of proprietary models suggests they may have encountered scientific publications during pre-training (Zhao et al., 2023). Kandpal, Wallace, and Raffel (2022) demonstrated that LLMs can effectively memorise portions of their training data, particularly for well-structured technical content.

²² Amazon Bedrock - bedrock<https://aws.amazon.com/bedrock>.

²³ OpenAI API - <https://platform.openai.com/docs/overview>.

²⁴ Kobold API - <https://github.com/KoboldAI/KoboldAI-Client>.

Table 1

Precision, Recall, and F1-Score for standard prompting, one-way strategy. BR are the performance associated to *broader* relationships, NA to *narrower*, SA to *same-as*, OT to *others*, and AVG are average performance considering the four types of relationships. In **bold** are the best performing Models for a given class.

MODEL	PRECISION					RECALL					F1-SCORE				
	AVG	BR	NR	SA	OT	AVG	BR	NR	SA	OT	AVG	BR	NR	SA	OT
mistral	0.401	0.667	0.000	0.511	0.428	0.462	0.008	0.000	0.848	0.992	0.313	0.016	0.000	0.638	0.598
mixtral	0.817	0.819	0.634	1.000	0.816	0.781	0.688	0.888	0.608	0.940	0.779	0.748	0.740	0.756	0.874
llama-2	0.809	0.796	1.000	0.565	0.877	0.722	0.872	0.140	0.992	0.884	0.669	0.832	0.246	0.720	0.880
dolphin-mistral	0.676	0.613	0.736	0.901	0.455	0.599	0.652	0.412	0.508	0.824	0.599	0.632	0.528	0.650	0.586
dolphin-mistral-dpo	0.734	0.613	0.868	0.960	0.497	0.603	0.988	0.236	0.284	0.904	0.552	0.757	0.371	0.438	0.641
dolphin-openorca	0.780	0.612	0.810	0.960	0.737	0.733	0.952	0.496	0.576	0.908	0.724	0.745	0.615	0.720	0.814
openchat	0.826	0.930	0.991	0.902	0.481	0.701	0.740	0.416	0.664	0.984	0.705	0.824	0.586	0.765	0.646
openchat-gemma	0.679	0.396	0.471	1.000	0.849	0.523	0.968	0.296	0.268	0.560	0.506	0.562	0.364	0.423	0.675
solar	0.723	0.642	0.877	0.850	0.523	0.657	0.916	0.544	0.784	0.384	0.646	0.755	0.672	0.529	0.627
mistral-openorca	0.815	0.866	0.401	1.000	0.992	0.605	0.568	0.996	0.384	0.472	0.613	0.686	0.572	0.555	0.640
eurus	0.596	0.410	0.638	0.424	0.911	0.493	0.984	0.204	0.332	0.452	0.466	0.579	0.309	0.372	0.604
llama-3	0.715	0.560	0.446	1.000	0.853	0.582	0.732	0.808	0.256	0.532	0.568	0.634	0.575	0.408	0.655
orca-2	0.500	0.342	0.382	0.486	0.788	0.456	0.944	0.052	0.068	0.760	0.372	0.502	0.091	0.119	0.774
gpt-3.5	0.668	0.267	0.407	1.000	1.000	0.295	0.976	0.096	0.032	0.076	0.195	0.419	0.155	0.062	0.141
gpt-4	0.952	0.919	0.934	1.000	0.954	0.949	1.000	0.968	0.832	0.996	0.948	0.958	0.951	0.908	0.975
haiku	0.820	0.496	0.870	1.000	0.912	0.710	0.980	0.564	0.468	0.828	0.712	0.659	0.684	0.638	0.868
sonnet	0.969	0.929	0.950	0.996	1.000	0.967	0.996	0.996	0.932	0.944	0.967	0.961	0.973	0.963	0.971

5.2. Standard prompting, two-way strategy

Table 2 presents the performance of the models when using standard prompting with the two-way strategy.

Among the full open models, mixtral achieves an average F1-score of 0.847, significantly outperforming both llama-2 and mistral. It is able to identify the broader, narrower, and same-as with F1-score around 0.82, but shows exceptional capabilities (0.925) in determining the relationship other. The superior performance of mixtral in the two-way classification strategy likely arises from its mixture-of-experts architecture, which can dynamically activate different subsets of experts depending on the input (Jiang et al., 2024). Therefore, mixtral may be able to selectively engage experts specialised in reasoning along specific hierarchical directions (Zadouri et al., 2023).

Regarding the quantised models, there are three models (e.g., dolphin-mistral-dpo, dolphin-openorca, and openchat) that obtain an average F1-score above 0.78. The strong performance of these quantised models when employing the two-way strategy likely reflects recent advancements in their fine-tuning methodologies. Specifically, dolphin-mistral-dpo benefits from Direct Preference Optimisation (DPO), a technique that improves the model’s capacity to follow complex instructions (Rafailov et al., 2023), while dolphin-openorca leverages demonstration-based learning from more capable teacher models (Mukherjee et al., 2023).

Among the proprietary models, sonnet obtained comparable results to previous experiments (0.965 F1), where the performance of gpt-3.5, gpt-4, and haiku obtained a slight improvement.

In this case, the full open models and their quantised counterparts exhibit similar performance levels (e.g., 0.847 for mixtral and 0.853 for dolphin-openorca). Moreover, the two-way strategy has effectively reduced the performance gap between these models and their larger, proprietary counterparts, narrowing the difference to just 0.113 F1 points. This enhancement suggests that the two-way strategy is effective in mitigating the performance gap between small and large models. This finding aligns with prior research demonstrating that similar self-verification mechanisms can improve model accuracy in information extraction tasks (Gero et al., 2023). Indeed, the significant performance gains observed in open models indicate that they may be capable of accessing relevant knowledge when prompted within an appropriate verification framework.

5.3. Chain-of-thought prompting, one-way strategy

Table 3 presents the results of the CoT, one-way strategy experiments. All models demonstrate significantly better performance compared to the analogous experiment using standard prompting. This confirms that CoT is a highly effective technique for enhancing model efficacy. A notable phenomenon observed here is that almost all models exhibit high precision on the same-as relationship. However, this often comes at the cost of recall, with the highest recall achieved by openchat, reaching 0.780.

Among full open models, mixtral again outperformed the other model with an average F1-score of 0.808. As per Section 5.1, mixtral’s strong performance is likely attributed to its mixture-of-experts architecture and extensive pre-training on diverse datasets, including scientific and technical literature (Jiang et al., 2024; Zhao et al., 2023).

Regarding the quantised models, dolphin-mistral achieves the highest average F1-score of 0.869. It is closely followed by dolphin-mistral-dpo with a score of 0.826, solar at 0.825, and openchat at 0.823. This suggests that models fine-tuned on the OpenOrca corpus are particularly effective for this task. The OpenOrca dataset’s inclusion of scientific and technical content likely contributes to this efficacy (Mukherjee et al., 2023). In this context, the use of the CoT strategy enables the quantised model dolphin-mistral to surpass the much larger mixtral model. This performance differential may stem from

Table 2

Precision, Recall, and F1-Score for standard prompting, two-way strategy. BR are the performance associated to *broader* relationships, NA to *narrower*, SA to *same-as*, OT to *others*, and AVG are average performance considering the four types of relationships. In **bold** are the best performing Models for a given class.

MODEL	PRECISION					RECALL					F1-SCORE				
	AVG	BR	NR	SA	OT	AVG	BR	NR	SA	OT	AVG	BR	NR	SA	OT
mistral	0.239	0.000	0.000	0.451	0.503	0.475	0.000	0.000	0.908	0.992	0.318	0.000	0.000	0.603	0.668
mixtral	0.865	0.751	0.773	1.000	0.935	0.846	0.904	0.856	0.708	0.916	0.847	0.820	0.812	0.829	0.925
llama-2	0.759	0.832	0.819	0.426	0.959	0.621	0.336	0.308	0.996	0.844	0.605	0.479	0.448	0.596	0.898
dolphin-mistral	0.718	0.686	0.673	0.884	0.629	0.704	0.708	0.732	0.704	0.672	0.706	0.697	0.701	0.764	0.664
dolphin-mistral-dpo	0.834	0.759	0.764	0.968	0.844	0.808	0.956	0.932	0.476	0.868	0.795	0.846	0.840	0.638	0.856
dolphin-openorca	0.862	0.799	0.774	0.928	0.947	0.852	0.888	0.892	0.776	0.852	0.853	0.841	0.829	0.845	0.897
openchat	0.838	0.979	0.964	0.823	0.586	0.777	0.736	0.648	0.744	0.980	0.783	0.840	0.775	0.782	0.734
openchat-gemma	0.734	0.484	0.512	1.000	0.939	0.590	0.804	0.792	0.332	0.432	0.579	0.605	0.622	0.499	0.592
solar	0.723	0.731	0.753	0.784	0.623	0.715	0.836	0.816	0.508	0.700	0.710	0.780	0.783	0.617	0.659
mistral-openorca	0.781	0.542	0.584	1.000	1.000	0.650	0.888	0.908	0.492	0.312	0.630	0.673	0.711	0.660	0.476
eurus	0.587	0.480	0.495	0.374	1.000	0.489	0.584	0.576	0.492	0.304	0.488	0.527	0.532	0.425	0.466
llama-3	0.746	0.476	0.544	1.000	0.964	0.586	0.856	0.836	0.328	0.324	0.562	0.611	0.659	0.494	0.485
orca-2	0.561	0.429	0.440	0.434	0.940	0.519	0.668	0.688	0.092	0.628	0.491	0.523	0.537	0.152	0.753
gpt-3.5	0.690	0.360	0.401	0.400	1.000	0.389	0.748	0.744	0.032	0.032	0.283	0.486	0.521	0.062	0.062
gpt-4	0.965	0.936	0.936	1.000	0.988	0.963	1.000	0.996	0.868	0.988	0.962	0.967	0.965	0.929	0.988
haiku	0.842	0.704	0.668	0.994	1.000	0.792	0.896	0.908	0.640	0.724	0.794	0.789	0.769	0.779	0.840
sonnet	0.967	0.929	0.940	1.000	1.000	0.965	0.992	0.996	0.948	0.924	0.965	0.959	0.967	0.973	0.961

Table 3

Precision, Recall, and F1-Score for chain-of-thought prompting, one-way strategy. BR are the performance associated to *broader* relationships, NA to *narrower*, SA to *same-as*, OT to *others*, and AVG are average performance considering the four types of relationships. In **bold** are the best performing models for a given class.

MODEL	PRECISION					RECALL					F1-SCORE				
	AVG	BR	NR	SA	OT	AVG	BR	NR	SA	OT	AVG	BR	NR	SA	OT
mistral	0.766	0.667	0.749	0.956	0.695	0.720	0.936	0.632	0.348	0.964	0.696	0.779	0.685	0.510	0.807
mixtral	0.849	0.732	0.807	1.000	0.858	0.820	0.948	0.904	0.488	0.940	0.808	0.826	0.853	0.656	0.897
llama-2	0.775	0.466	0.674	1.000	0.960	0.612	0.976	0.860	0.224	0.388	0.576	0.630	0.756	0.366	0.553
dolphin-mistral	0.892	0.749	0.912	1.000	0.907	0.871	1.000	0.872	0.680	0.932	0.869	0.856	0.892	0.810	0.919
dolphin-mistral-dpo	0.863	0.685	0.947	0.994	0.827	0.827	0.992	0.708	0.688	0.920	0.826	0.811	0.810	0.813	0.871
dolphin-openorca	0.840	0.524	0.865	1.000	0.971	0.736	1.000	0.848	0.560	0.536	0.738	0.688	0.857	0.718	0.691
openchat	0.844	0.692	0.874	0.947	0.864	0.822	0.988	0.780	0.780	0.740	0.823	0.814	0.825	0.855	0.797
openchat-gemma	0.822	0.429	0.864	1.000	0.993	0.650	1.000	0.408	0.588	0.604	0.662	0.600	0.554	0.741	0.751
solar	0.858	0.739	0.771	1.000	0.924	0.830	0.940	0.928	0.580	0.872	0.825	0.828	0.842	0.734	0.897
mistral-openorca	0.810	0.551	0.865	0.987	0.836	0.736	1.000	0.744	0.608	0.592	0.739	0.710	0.800	0.753	0.693
eurus	0.796	0.661	0.635	1.000	0.887	0.723	1.000	0.912	0.416	0.564	0.706	0.796	0.749	0.588	0.690
llama-3	0.723	0.547	0.491	1.000	0.852	0.626	0.696	0.748	0.232	0.828	0.605	0.613	0.593	0.377	0.840
orca-2	0.787	0.622	0.624	1.000	0.904	0.711	0.868	0.928	0.404	0.644	0.700	0.725	0.746	0.576	0.752
gpt-3.5	0.727	0.534	0.501	0.896	0.976	0.573	0.904	0.880	0.344	0.164	0.522	0.672	0.639	0.497	0.281
gpt-4	0.924	0.828	0.921	1.000	0.945	0.914	0.984	0.980	0.724	0.968	0.911	0.899	0.950	0.840	0.957
haiku	0.861	0.664	0.809	1.000	0.969	0.812	0.996	0.952	0.540	0.760	0.806	0.797	0.875	0.701	0.852
sonnet	0.909	0.821	0.865	1.000	0.950	0.896	0.992	0.976	0.696	0.920	0.893	0.899	0.917	0.821	0.935

dolphin-mistral’s specialised fine-tuning on instruction-following datasets that emphasise reasoning processes (Wang et al., 2022).

Among the proprietary models, gpt-4 achieved the highest performance, reaching an average F1-score of 0.911. Notably, both gpt-4 and sonnet exhibit a slight decline in performance when using the CoT prompting strategy, whereas haiku and gpt-3.5 show improved results under the same conditions. The counterintuitive behaviour of the larger models may arise from redundancy and cognitive overhead. Since these models have already internalised complex semantic relationships during pre-training, explicitly prompting them to articulate intermediate reasoning steps could introduce unnecessary noise (Liu, Geng et al., 2024; Wei et al., 2023). This additional complexity may disrupt their internal inference mechanisms, ultimately impairing performance and reducing accuracy.

5.4. Chain-of-thought prompting, two-way strategy

Table 4 reports the performance of the models when using CoT prompting, two-way strategy. The results are generally consistent with those obtained from the CoT one-way strategy. Specifically, mixtral, dolphin-mistral, and gtp-4 emerge as the top-performing models among the full open models, the quantised models, and the proprietary models, respectively. The average F1 score increases by a few percentage points in 15 out of 17 models, confirming that the two-way strategy remains effective when combined with CoT. The consistent improvement across models when employing the two-way strategy aligns with previous findings demonstrating that techniques aimed at enhancing logical consistency can lead to more robust judgements (Liu, Guo et al., 2024).

Table 4

Precision, Recall, and F1-Score for chain-of-thought prompting, two-way strategy. BR are the performance associated to *broader* relationships, NA to *narrower*, SA to *same-as*, OT to *others*, and AVG are average performance considering the four types of relationships. In **bold** are the best performing Models for a given class.

MODEL	PRECISION					RECALL					F1-SCORE				
	AVG	BR	NR	SA	OT	AVG	BR	NR	SA	OT	AVG	BR	NR	SA	OT
mistral	0.830	0.775	0.734	0.932	0.880	0.810	0.936	0.904	0.492	0.908	0.799	0.848	0.810	0.644	0.894
mixtral	0.876	0.768	0.795	1.000	0.940	0.850	0.992	0.948	0.576	0.884	0.843	0.866	0.865	0.731	0.911
llama-2	0.766	0.502	0.563	1.000	1.000	0.588	0.936	0.928	0.280	0.208	0.534	0.654	0.701	0.438	0.344
dolphin-mistral	0.930	0.874	0.852	1.000	0.995	0.920	0.996	0.992	0.844	0.848	0.920	0.931	0.917	0.915	0.916
dolphin-mistral-dpo	0.915	0.859	0.831	0.986	0.982	0.905	0.952	0.964	0.840	0.864	0.906	0.903	0.893	0.907	0.919
dolphin-openorca	0.828	0.640	0.670	1.000	1.000	0.745	0.968	0.968	0.704	0.340	0.724	0.771	0.792	0.826	0.508
openchat	0.857	0.763	0.777	0.903	0.986	0.834	0.940	0.932	0.892	0.572	0.828	0.842	0.847	0.897	0.724
openchat-gemma	0.798	0.601	0.597	0.994	1.000	0.711	0.868	0.852	0.672	0.452	0.709	0.710	0.702	0.802	0.623
solar	0.886	0.759	0.815	1.000	0.971	0.861	0.980	0.968	0.704	0.792	0.860	0.855	0.885	0.826	0.872
mistral-openorca	0.846	0.686	0.698	1.000	1.000	0.784	0.968	0.972	0.756	0.440	0.772	0.803	0.813	0.861	0.611
eurus	0.833	0.672	0.670	1.000	0.990	0.756	0.992	0.992	0.628	0.412	0.739	0.801	0.800	0.772	0.582
llama-3	0.766	0.556	0.541	1.000	0.968	0.661	0.808	0.812	0.292	0.732	0.649	0.659	0.650	0.452	0.834
orca-2	0.809	0.598	0.651	1.000	0.985	0.724	0.904	0.912	0.544	0.536	0.720	0.720	0.760	0.705	0.694
gpt-3.5	0.741	0.547	0.553	0.863	1.000	0.606	0.904	0.920	0.528	0.072	0.540	0.682	0.691	0.655	0.134
gpt-4	0.941	0.899	0.889	1.000	0.975	0.935	1.000	0.996	0.792	0.952	0.933	0.947	0.940	0.884	0.964
haiku	0.877	0.755	0.765	1.000	0.989	0.842	0.984	0.988	0.680	0.716	0.839	0.854	0.862	0.809	0.831
sonnet	0.916	0.822	0.855	1.000	0.986	0.901	1.000	0.992	0.768	0.844	0.900	0.902	0.918	0.869	0.909

For quantised models, the good performance of *dolphin-mistral* further suggests that pretraining on OpenOrca (Mukherjee et al., 2023) yields a consistent advantage on this task. In this case, the quantisation process (Dettmers, Lewis, Belkada, & Zettlemoyer, 2022) appears to preserve the model’s core representational capacity for handling technical terminology. This preservation of performance post-quantisation is particularly noteworthy and indicates that the model’s parameters encode knowledge efficiently, even under reduced precision constraints (Frantar, Ashkboos, Hoefler, & Alistarh, 2023).

The variation in performance across different versions of the same base model (e.g., different *mistral* derivatives) suggests that fine-tuning procedures can either enhance or diminish domain-specific capabilities present in the base model (Chang et al., 2024). This highlights the trade-offs inherent in optimising for general capabilities versus domain-specific knowledge retention (Zhou et al., 2023). The consistent improvement achieved through two-way verification with CoT reasoning across nearly all models suggests that, regardless of pre-training, explicit reasoning helps models more effectively leverage their knowledge of technical terminology and scientific concepts (Miao, Teh, & Rainforth, 2023; Zheng et al., 2023).

5.5. Comparative analysis

Table 5 presents the average F1-scores for all 17 models across the four strategies, including the relative differences between one-way and two-way strategies for both standard and CoT prompting. The results indicate that shifting from a one-way to a two-way strategy improved performance in 14 models for standard prompting and in 15 models for CoT prompting.

Generally, we observe a gradual increase in the average F1-score from left to right in Table 5, demonstrating the beneficial effects of both CoT and the two-way strategy. Thirteen models showed overall improvement, while four models (*llama-2*, *dolphin-openorca*, *gpt-4*, and *sonnet*) experienced a decline in performance. When applying a 0.7 average F1-score as a quality benchmark, only *dolphin-openorca* and *openchat* surpassed this threshold under a standard prompting approach with a one-way strategy. However, with a two-way strategy in standard prompting, five models exceeded the threshold. This number increased to eight models when a CoT, one-way strategy was used. Ultimately, nine models achieved an average F1-score above the 0.7 threshold when employing a CoT, two-way strategy. These results underscore the critical importance of a good prompting strategy when working in a zero-shot setting.

Finally, we can analyse the best model for each of the three categories: fully open models, quantised models, and proprietary models. The *mixtral* model stands out among fully open models, particularly when used with the standard prompting and two-way strategy (0.847 F1). In contrast, the *dolphin-mistral* model surpasses other quantised models when applied with the CoT prompting and two-way strategy (0.920 F1). Notably, *dolphin-mistral* outperforms *mixtral*, likely due to the additional fine-tuning on datasets like OpenOrca, which seem particularly beneficial for this task. Among proprietary models, *sonnet* delivers the best performance when using standard prompts and a straightforward, one-way approach (0.967 F1). Interestingly, unlike most open models, its performance declines with more complex prompts. Since *sonnet* is a black-box system that may incorporate additional components, pinpointing the exact cause of this behaviour is challenging.

6. Error analysis

In this section, we examine common error types displayed by the top-performing LLMs across the three categories: (a) *mixtral* (standard prompting, two-way), (b) *dolphin-mistral* (CoT, two-way), and (c) *sonnet* (standard prompting, one-way). Fig. 2 reports the relevant confusion matrices.

Table 5

Average F1-Scores obtained from the 17 Models on all experiments. In **bold** are the best performing Models for a given strategy.

MODEL	Standard Prompting			CoT Prompting			Overall Improvement
	One-way	Two-way	Difference	One-way	Two-way	Difference	
mistral	0.313	0.318	0.005	0.696	0.799	0.103	0.486
mixtral	0.779	0.847	0.068	0.808	0.843	0.035	0.068
llama-2	0.669	0.605	-0.064	0.576	0.534	-0.042	0.135
dolphin-mistral	0.599	0.706	0.107	0.869	0.920	0.051	0.321
dolphin-mistral-dpo	0.552	0.795	0.243	0.826	0.906	0.080	0.354
dolphin-openorca	0.724	0.853	0.129	0.738	0.724	-0.014	0.129
openchat	0.705	0.783	0.078	0.823	0.828	0.005	0.123
openchat-gemma	0.506	0.579	0.073	0.662	0.709	0.047	0.203
solar	0.646	0.710	0.064	0.825	0.860	0.035	0.214
mistral-openorca	0.613	0.630	0.017	0.739	0.772	0.033	0.159
eurus	0.466	0.488	0.022	0.706	0.739	0.033	0.273
llama-3	0.568	0.562	-0.006	0.605	0.649	0.044	0.087
orca-2	0.372	0.491	0.119	0.700	0.720	0.020	0.348
gpt-3.5	0.195	0.283	0.088	0.522	0.540	0.018	0.345
gpt-4	0.948	0.962	0.014	0.911	0.933	0.022	0.051
haiku	0.712	0.794	0.082	0.806	0.839	0.033	0.127
sonnet	0.967	0.965	-0.002	0.893	0.900	0.007	0.074

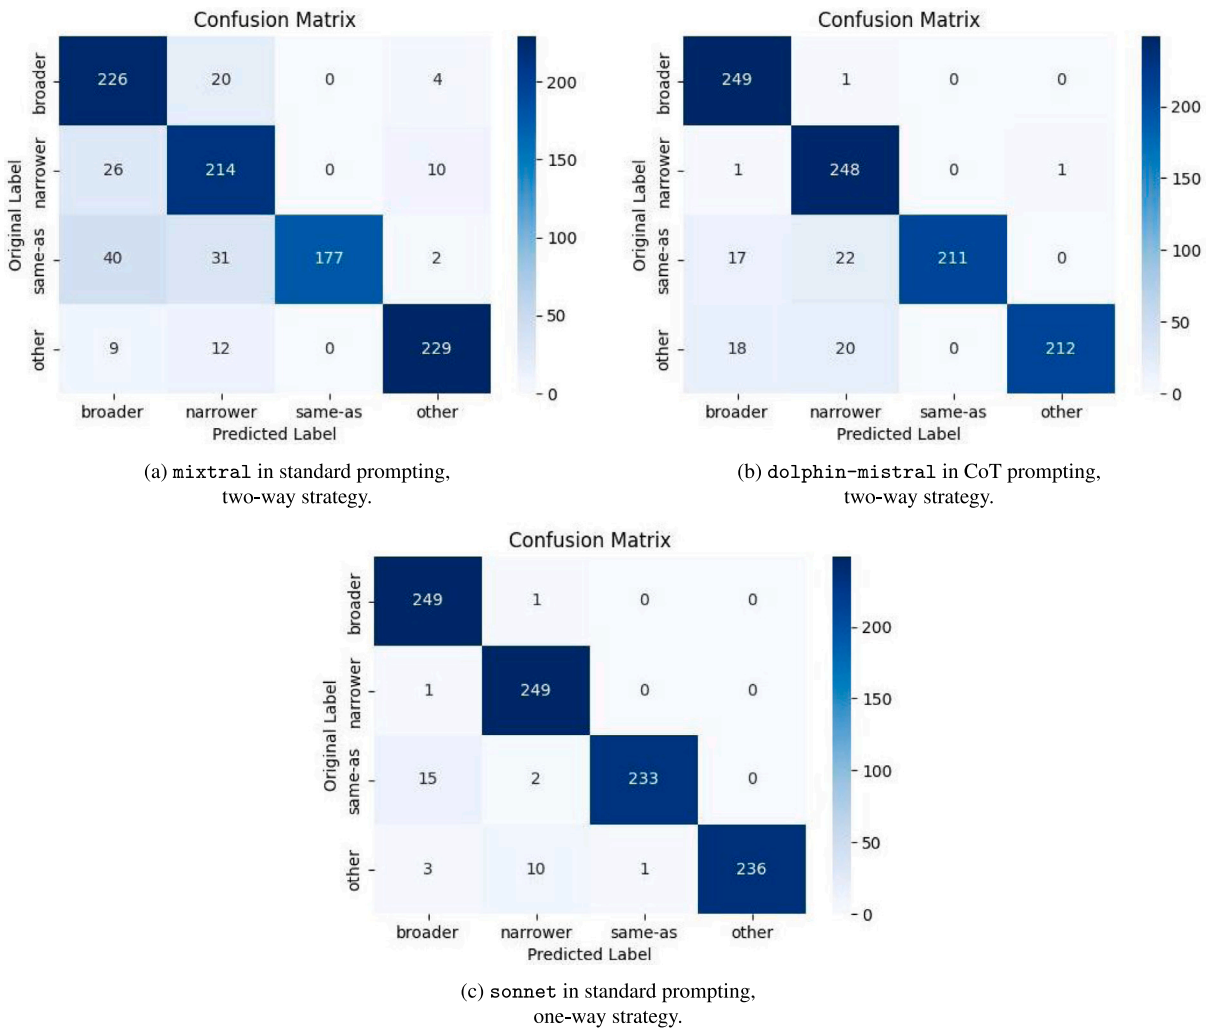


Fig. 2. Confusion matrices: (a) for the best open model, (b) for the best quantised model, and (c) for the best proprietary model.

The *mixtral* model misclassified 154 out of 1000 topic pairs. Specifically, 75 pairs were incorrectly identified as broader, 63 as narrower, 16 as other, and none as same-as. *mixtral* demonstrates a highly conservative approach in predicting the same-as category, achieving a precision of 1, as indicated in Table 2. This comes at a cost, as the model exhibits a tendency to misclassify pairs that are labelled as same-as in the gold standard, often categorising them incorrectly as either broader or narrower. This misclassification yields a low recall for the same-as category (0.708), which represents a significant limitation of the model. Specifically, 29.2% of same-as pairs were misclassified: 16% were labelled as broader, 12.4% as narrower, and 0.8% as other. Additionally, the *mixtral* model occasionally confuses broader pairs with narrower ones and vice versa. Notably, 9.6% of broader pairs were incorrectly predicted, with 8% labelled as narrower and 1.6% as other. Similarly, 14.4% of narrower pairs were misclassified, with 10.4% labelled as broader and 4% as other.

The *dolphin-mistral* model exhibits a different pattern compared to *mixtral*, particularly excelling in pairs labelled as broader and narrower in the gold standard. It makes 80 misclassifications, predominantly in the same-as (39) and other (38) categories. Although the model still faces challenges with same-as relationships, it performs better than *mixtral*. Specifically, for the same-as relationships, a small percentage (6.8%) are misclassified as broader, while 8.8% are incorrectly identified as narrower. Notably, the other relationships are misclassified equally as broader or narrower.

The *sonnet* model exhibits a similar performance pattern to *dolphin-mistral*, achieving near-perfect accuracy across both broader and narrow categories, while showing notable improvements in the other two categories. The model made only 33 classification errors, with the majority occurring in same-as (17) and other (14). Specifically, within the same-as misclassifications, 6% were incorrectly identified as broader, and 0.8% as narrower. In contrast, for the other category, 1.6% of errors were misclassified as broader or same-as, while 4% were erroneously categorised as other.

The error analysis provides several valuable insights. First, the best-performing model excels in identifying hierarchical relationships, such as broader and narrower. This is a promising indication of the potential for LLMs to assist in ontology generation, where hierarchical relationships are fundamental to structuring the ontology. Second, while the very best model shows reasonable success in identifying “same-as” relationships, this task proves challenging. This difficulty may arise because same-as relationships between research topics are not purely syntactic synonyms and need deep expertise in the field to be interpreted correctly. Fine-tuning the model on task-specific data could help address this issue, and we plan to explore this approach in future work.

7. Implications and limitations

The experiments presented in this study led to several notable insights. *First*, many models demonstrate strong zero-shot reasoning capabilities for this task, at least within the domain of engineering and computer science as represented by the IEEE Thesaurus. This performance may be partly due to the nature of relevant concepts in electrical engineering and computer science, which are likely well represented in the training data of LLMs, including code repositories, technical documentation, manuals, and patents. We intend to extend our analysis to other disciplines to determine whether this strong performance is consistent across fields and to evaluate the potential benefits of fine-tuning models for this specific task. *Second*, the performance of LLMs is strongly influenced by the choice of prompting strategy. In particular, CoT prompting clearly outperforms standard prompting, aligning with findings from prior research (Wei et al., 2023). Additionally, running the classification process twice, switching the topic position, and applying simple rules to integrate the two outputs further reduces the number of errors. Indeed, the two-way CoT variant outperforms all other approaches, especially for smaller, more scalable, and cost-efficient models. The analysis indicates that selecting appropriate prompting strategies can lead to notable performance improvements, with gains exceeding 0.2 F1 points in several cases (see Table 5, last column). *Third*, smaller models achieve performance comparable to that of the best-performing architectures when adopting the best prompting strategy. Such models, particularly when fine-tuned on relevant datasets, can be highly effective for this task and may even outperform their base versions. For example, *dolphin-mistral* and *dolphin-mistral-dpo* surpass the standard *mistral* and are only outperformed by the largest proprietary models. These results suggest promising directions for the development of optimised and efficient models for ontology generation.

While we believe that these insights represent a substantial advancement in the state of the art, our current research settings present certain limitations that warrant discussion. We plan to address and explore these further in future work.

The first limitation lies in the domain-specific nature of our findings, which, as previously discussed, are currently confined to the IEEE taxonomy and the engineering field. A broader evaluation is required to assess the performance of LLMs in other domains. To this end, we are actively working on expanding our dataset to include classifications from additional disciplines, such as the Medical Subject Headings and the Physics Subject Headings, which represent the domains of biomedicine and physics, respectively. A second limitation is that our evaluation only focuses on LLM-based strategies. We do not compare these approaches with traditional AI methods, such as rule-based, statistical, or hybrid techniques. Nevertheless, we recognise the importance of such comparisons and intend to include them in future experiments. A third limitation arises from the dynamic nature of scientific research, where disciplines are continuously redefined by emerging paradigms and discoveries. Although our current results provide valuable insights, LLMs trained on present-day data may struggle to accurately classify relationships involving topics that have not yet emerged. Therefore, while our methodological approach offers long-term relevance, the specific LLMs employed will need to be updated periodically to reflect the most recent developments. We posit that future generations of LLMs will likely continue the observed trend of improvement, offering greater efficiency and enhanced performance at a reduced computational cost.

In addition, while this paper focused exclusively on identifying relationships among a subset of topics, future work will need to address the construction of a complete ontology. This entails resolving several challenges, including the evaluation of relationships across a much larger set of topics. To improve scalability and efficiency, it is crucial to identify strategies that reduce the search

space. One promising approach proposed in the literature (Pisu, Pompianu, Salatino, Osborne, Riboni, Motta, & Reforgiato Recupero, 2024a; Pisu et al., 2024b) involves leveraging topic co-occurrence patterns in research paper abstracts and titles. Topic pairs that rarely or never co-occur can be automatically classified under the “other” category, thereby significantly reducing the computational cost of model inference.

Finally, while the results produced by LLMs demonstrate high accuracy, in realistic settings they must be further verified and refined by human experts. We envision a collaborative workflow in which LLMs generate initial ontology drafts that are subsequently reviewed and refined by domain experts. This process ensures both accuracy and the preservation of critical human judgement, which remains essential for developing robust and reliable KOS. Recent studies have proposed various strategies for integrating human-in-the-loop methodologies into the construction of knowledge graphs and ontologies (Tsaneva, Dessi, Osborne, & Sabou, 2025).

8. Conclusions

This paper conducted an in-depth analysis of the capability of LLMs to identify semantic relationships between research topics for the purpose of ontology generation. We evaluated seventeen LLMs, categorised into three groups: fully open models (3), open quantised models (10), and proprietary models (4). We evaluated their performance against a gold standard of 1000 relationships from the IEEE Thesaurus. We also investigated four distinct zero-shot prompting strategies, which yielded significantly different outcomes.

Our results demonstrated the impressive zero-shot reasoning capabilities of state-of-the-art models. Three models emerged as top performers within their respective categories: Mixtral-8 × 7B (mixtral: 0.847 average F1-score) for the full open models; Dolphin-2.1-Mistral-7B (dolphin-mistral: 0.920) for open quantised models; and Claude 3 Sonnet (sonnet: 0.967) for proprietary models.

We are currently working on several fronts. We are fine-tuning off-the-shelf models like BERT, Mistral, LLaMa 3, Flan T5, Gemma, Phi 3, and DeepSeek with the aim of enhancing the performance of lightweight models to reduce computational costs and environmental impact when generating ontologies of research areas. Furthermore, we plan to incorporate a semantic reasoner, such as RDFox (Nenov et al., 2015), into our pipeline to identify and resolve inconsistencies. We will also make all the resulting ontologies publicly available, formatted according to the SKOS standard. Other directions include exploring explainability techniques, including SHAP values (Marcilio & Eler, 2020), attention visualisation (Vig, 2019), and the identification of monosemantic features (Templeton et al., 2024) to understand how LLMs process ontological relationships. We also plan to investigate potential biases in LLM-generated ontologies, such as examining whether inference error rates vary across different sub-domains of Engineering. Additionally, we will benchmark our findings against additional state-of-the-art ontology generation methods, including rule-based systems and statistical approaches. Ultimately, our goal is to develop a comprehensive, cross-disciplinary ontology of research topics. To this end, we plan to expand our analysis to cover a variety of scientific domains, including materials science, medicine, and physics.

CRedit authorship contribution statement

Tanay Aggarwal: Writing – original draft, Software, Formal analysis, Data curation. **Angelo Salatino:** Writing – review & editing, Supervision, Software, Methodology, Data curation. **Francesco Osborne:** Writing – review & editing, Supervision, Project administration, Data curation. **Enrico Motta:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) used Grammarly in order to improve readability. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This publication is based upon work from COST Action CA23147 GOBLIN - Global Network on Large-Scale, Cross-domain and Multilingual Open Knowledge Graphs, supported by COST (European Cooperation in Science and Technology, <https://www.cost.eu>).

Appendix A. Prompt for standard prompting

For consistency we applied the same prompt across all the models. We engineered this prompt through various refinements, to ensure optimal comprehension of the task and accurate responses. Below is the template of our prompt, customised for each topic pair by substituting [TOPIC-A] for the first topic and [TOPIC-B] for the second.

Classify the relationship between '[TOPIC-A]' and '[TOPIC-B]' by applying the following relationship definitions:

1. '[TOPIC-A]' is-broader-than '[TOPIC-B]' if '[TOPIC-A]' is a super-category of '[TOPIC-B]', that is '[TOPIC-B]' is a type, a branch, or a specialised aspect of '[TOPIC-A]' or that '[TOPIC-B]' is a tool or a methodology mostly used in the context of '[TOPIC-A]' (e.g., car is-broader-than wheel).
2. '[TOPIC-A]' is-narrower-than '[TOPIC-B]' if '[TOPIC-A]' is a sub-category of '[TOPIC-B]', that is '[TOPIC-A]' is a type, a branch, or a specialised aspect of '[TOPIC-B]' or that '[TOPIC-A]' is a tool or a methodology mostly used in the context of '[TOPIC-B]' (e.g., wheel is-narrower-than car).
3. '[TOPIC-A]' is-same-as-than '[TOPIC-B]' if '[TOPIC-A]' and '[TOPIC-B]' are synonymous terms denoting an identical concept (e.g., beautiful is-same-as-than attractive), including when one is the plural form of the other (e.g., cat is-same-as-than cats).
4. '[TOPIC-A]' is-other-than '[TOPIC-B]' if '[TOPIC-A]' and '[TOPIC-B]' either have no direct relationship or share a different kind of relationship that does not fit into the other defined relationships.

Given the previous definitions, determine which one of the following statements is correct:

1. '[TOPIC-A]' is-broader-than '[TOPIC-B]'
2. '[TOPIC-B]' is-narrower-than '[TOPIC-A]'
3. '[TOPIC-A]' is-narrower-than '[TOPIC-B]'
4. '[TOPIC-B]' is-broader-than '[TOPIC-A]'
5. '[TOPIC-A]' is-same-as-than '[TOPIC-B]'
6. '[TOPIC-A]' is-other-than '[TOPIC-B]'

Answer by only stating the correct statement and its number.

Appendix B. Prompt for chain-of-thought prompting

In the CoT prompting we have a two-phase interaction with the model. In the first interaction the model is asked to provide a definition of both topics, formulate a sentence that incorporates both topics, and finally discuss their potential semantic relationships. The resulting response is then fed back into the model in the second phase, with the sole task of determining the type of semantic relationship. Below are our prompt templates, customised for each topic pair by substituting [TOPIC-A] for the first topic and [TOPIC-B] for the second.

First interaction

Classify the relationship between '[TOPIC-A]' and '[TOPIC-B]' by applying the following relationship definitions:

1. '[TOPIC-A]' is-broader-than '[TOPIC-B]' if '[TOPIC-A]' is a super-category of '[TOPIC-B]', that is '[TOPIC-B]' is a type, a branch, or a specialised aspect of '[TOPIC-A]' or that '[TOPIC-B]' is a tool or a methodology mostly used in the context of '[TOPIC-A]' (e.g., car is-broader-than wheel).
2. '[TOPIC-A]' is-narrower-than '[TOPIC-B]' if '[TOPIC-A]' is a sub-category of '[TOPIC-B]', that is '[TOPIC-A]' is a type, a branch, or a specialised aspect of '[TOPIC-B]' or that '[TOPIC-A]' is a tool or a methodology mostly used in the context of '[TOPIC-B]' (e.g., wheel is-narrower-than car).
3. '[TOPIC-A]' is-same-as-than '[TOPIC-B]' if '[TOPIC-A]' and '[TOPIC-B]' are synonymous terms denoting a very similar concept (e.g., 'beautiful' is-same-as-than 'attractive'), including when one is the plural form of the other (e.g., cat is-same-as-than cats).
4. '[TOPIC-A]' is-other-than '[TOPIC-B]' if '[TOPIC-A]' and '[TOPIC-B]' either have no direct relationship or share a different kind of relationship that does not fit into the other defined relationships.

Think step by step by following these sequential instructions:

- 1) Provide a precise definition for '[TOPIC-A]'.
- 2) Provide a precise definition for '[TOPIC-B]'.
- 3) Formulate a sentence that includes both '[TOPIC-A]' and '[TOPIC-B]'.
- 4) Discuss '[TOPIC-A]' and '[TOPIC-B]' usage and relationship (is-narrower-than, is-broader-than, is-same-as-than, or is-other-than).

Second interaction

[PREVIOUS-RESPONSE]

Given the previous discussion, determine which one of the following statements is correct:

1. '[TOPIC-A]' is broader than '[TOPIC-B]'
2. '[TOPIC-B]' is narrower than '[TOPIC-A]'
3. '[TOPIC-A]' is narrower than '[TOPIC-B]'
4. '[TOPIC-B]' is broader than '[TOPIC-A]'
5. '[TOPIC-A]' is same as than '[TOPIC-B]'
6. '[TOPIC-A]' is other than '[TOPIC-B]'

Answer by only stating the number of the correct statement.

Data availability

We have created a github repo, linked in the paper.

References

- AI@Meta (2024). Llama 3 model card. URL https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md.
- Angioni, S., Salatino, A., Osborne, F., Recupero, D. R., & Motta, E. (2021). AIDA: A knowledge graph about research dynamics in academia and industry. *Quantitative Science Studies*, 2(4), 1356–1398.
- ANSI/NISO Z39.4-2021, criteria for indexes. (2021). <http://dx.doi.org/10.3789/ansi.niso.z39.4-2021>, URL <http://dx.doi.org/10.3789/ansi.niso.z39.4-2021>.
- Anthropic (2024). The claude 3 model family: Opus, sonnet, haiku. *Online*.
- Babaei Giglou, H., D'Souza, J., & Auer, S. (2023). LLMs4OL: Large language models for ontology learning. In *International semantic web conference* (pp. 408–427). Springer.
- Beel, J., Gipp, B., Langer, S., & Breiting, C. (2016). Paper recommender systems: a literature survey. *International Journal on Digital Libraries*, 17, 305–338.
- Bolanos, F., Salatino, A., Osborne, F., & Motta, E. (2024). Artificial intelligence for literature reviews: Opportunities and challenges. *Artificial Intelligence Review*, 57(10), 259.
- Bornmann, L., & Mutz, R. (2015). Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references. *Journal of the Association for Information Science and Technology*, 66(11), 2215–2222. <http://dx.doi.org/10.1002/asi.23329>.
- Brody, S. (2021). Scite. *Journal of the Medical Library Association: JMLA*, 109(4), 707.
- Bubeck, S., Chadrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., et al. (2023). Sparks of artificial general intelligence: Early experiments with gpt-4.
- Buscaldi, D., Dessì, D., Motta, E., Murgia, M., Osborne, F., & Recupero, D. R. (2024). Citation prediction by leveraging transformers and natural language processing heuristics. *Information Processing & Management*, 61(1), Article 103583.
- Cadeddu, A., Chessa, A., De Leo, V., Fenu, G., Motta, E., Osborne, F., et al. (2024). A comparative analysis of knowledge injection strategies for large language models in the scholarly domain. *Engineering Applications of Artificial Intelligence*, 133, Article 108166.
- Canito, A., Corchado, J., & Marreiros, G. (2022). A systematic review on time-constrained ontology evolution in predictive maintenance. *Artificial Intelligence Review*, 55(4), 3183–3211.
- Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., et al. (2024). A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3), 1–45.
- Chen, C., Lin, K., & Klein, D. (2020). Constructing taxonomies from pretrained language models. arXiv preprint arXiv:2010.12813.
- Cimiano, P., & Völker, J. (2005). Text2onto: A framework for ontology learning and data-driven change discovery. In *International conference on application of natural language to information systems* (pp. 227–238). Springer.
- de Bruin, T., Kober, J., Tuyls, K., & Babuška, R. (2020). Fine-tuning deep RL with gradient-free optimization. *IFAC-PapersOnLine*, 53(2), 8049–8056.
- Dessì, D., Osborne, F., Reforgiato Recupero, D., Buscaldi, D., & Motta, E. (2022). CS-KG: A large-scale knowledge graph of research entities and claims in computer science. In *International semantic web conference* (pp. 678–696). Springer.
- Dettmers, T., Lewis, M., Belkada, Y., & Zettlemoyer, L. (2022). Gpt3.int8(): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35, 30318–30332.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
- Dunne, E., & Hulek, K. (2020). Mathematics subject classification 2020. *EMS Newsletter*, [ISSN: 1027-488X] 2020–3(115), 5–6. <http://dx.doi.org/10.4171/news/115/2>.
- Färber, M., Lamprecht, D., Krause, J., Aung, L., & Haase, P. (2023). SemOpenAlex: The scientific landscape in 26 Billion RDF triples. In *International semantic web conference* (pp. 94–112). Springer.
- Frantar, E., Ashkboos, S., Hoefler, T., & Alistarh, D.-A. (2023). OPTQ: Accurate post-training quantization for generative pre-trained transformers. In *11th international conference on learning representations*.
- Genesereth, M. R., & Nilsson, N. J. (2012). *Logical foundations of artificial intelligence*. Morgan Kaufmann, <http://dx.doi.org/10.1016/C2009-0-27551-9>.
- Gero, Z., Singh, C., Cheng, H., Naumann, T., Galley, M., Gao, J., et al. (2023). Self-verification improves few-shot clinical information extraction. arXiv preprint arXiv:2306.00024.
- Grootendorst, M. (2022). BERTopic: Neural topic modeling with a class-based TF-IDF procedure. arXiv preprint arXiv:2203.05794.
- Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, [ISSN: 1042-8143] 5(2), 199–220. <http://dx.doi.org/10.1006/knac.1993.1008>, URL <https://www.sciencedirect.com/science/article/pii/S1042814383710083>.
- (2010). *Guidelines for the construction, format, and management of monolingual controlled vocabularies: Standard*, Baltimore, Maryland: National Information Standards Organization, <http://dx.doi.org/10.3789/ansi.niso.z39.19-2005R2010>.
- Gusenbauer, M., & Haddaway, N. R. (2020). Which academic search systems are suitable for systematic reviews or meta-analyses? Evaluating retrieval qualities of Google Scholar, PubMed, and 26 other resources. *Research Synthesis Methods*, 11(2), 181–217.

- Han, K., Yang, P., Mishra, S., & Diesner, J. (2020). WikiCSSH: extracting computer science subject headings from Wikipedia. In *ADBIS, TPD and EDA 2020 common workshops and doctoral consortium: international workshops: DOING, MADEISD, SKG, BBIGAP, SIMPDA, aiMinScience 2020 and doctoral consortium, Lyon, France, August 25–27, 2020, proceedings 24* (pp. 207–218). Springer.
- Hedden, H. (2010). Taxonomies and controlled vocabularies best practices for metadata. *Journal of Digital Asset Management*, 6, 279–284. <http://dx.doi.org/10.1057/dam.2010.29>.
- (2013). *Information and documentation. Thesauri and interoperability with other vocabularies. Interoperability with other vocabularies: Standard*, International Organization for Standardization.
- Jaradeh, M. Y., Oelen, A., Farfar, K. E., Prinz, M., D'Souza, J., Kismihók, G., et al. (2019). Open research knowledge graph: next generation infrastructure for semantic scholarly knowledge. In *Proceedings of the 10th international conference on knowledge capture* (pp. 243–246).
- Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., Casas, D. d. l., et al. (2023). Mistral 7B. arXiv preprint arXiv:2310.06825.
- Jiang, A. Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., et al. (2024). Mixtral of experts. arXiv preprint arXiv:2401.04088.
- Kandpal, N., Wallace, E., & Raffel, C. (2022). Deduplicating training data mitigates privacy risks in language models. In *International conference on machine learning* (pp. 10697–10707). PMLR.
- Khanal, S. S., Prasad, P., Alsadoon, A., & Maag, A. (2020). A systematic review: machine learning based recommendation systems for e-learning. *Education and Information Technologies*, 25(4), 2635–2664.
- Kim, D., Park, C., Kim, S., Lee, W., Song, W., Kim, Y., et al. (2023). SOLAR 10.7B: Scaling large language models with simple yet effective depth up-scaling. arXiv:2312.15166.
- Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2022). Large language models are zero-shot reasoners. *Advances in Neural Information Processing Systems*, 35, 22199–22213.
- Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2023). Large language models are zero-shot reasoners. arXiv:2205.11916 URL <https://arxiv.org/abs/2205.11916>.
- Komminen, V. K., König-Ries, B., & Samuel, S. (2024). From human experts to machines: An LLM supported approach to ontology and knowledge graph construction. arXiv preprint arXiv:2403.08345.
- Kuhn, T., Chichester, C., Krauthammer, M., Queralt-Rosinach, N., Verborgh, R., Giannakopoulos, G., et al. (2016). Decentralized provenance-aware publishing with nanopublications. *PeerJ Computer Science*, 2, Article e78.
- Le, M., Roller, S., Papaxanthos, L., Kiela, D., & Nickel, M. (2019). Inferring concept hierarchies from text corpora via hyperbolic embeddings. arXiv preprint arXiv:1902.00913.
- Lian, W., Goodson, B., Wang, G., Pentland, E., Cook, A., Vong, C., et al. (2023). MistralOrca: Mistral-7B model instruct-tuned on filtered OpenOrcaV1 GPT-4 dataset. <https://huggingface.co/Open-Orca/Mistral-7B-OpenOrca>.
- Lipscomb, C. E. (2000). Medical subject headings (MeSH). *Bulletin of the Medical Library Association*, 88(3), 265.
- Liu, R., Geng, J., Wu, A. J., Sucholutsky, I., Lombrozo, T., & Griffiths, T. L. (2024). Mind your step (by step): Chain-of-thought can reduce performance on tasks where thinking makes humans worse. arXiv preprint arXiv:2410.21333.
- Liu, Y., Guo, Z., Liang, T., Shareghi, E., Vulić, I., & Collier, N. (2024). Aligning with logic: Measuring, evaluating and improving logical consistency in large language models. arXiv preprint arXiv:2410.02205.
- Longpre, S., Hou, L., Vu, T., Webson, A., Chung, H. W., Tay, Y., et al. (2023). The flan collection: Designing data and methods for effective instruction tuning. arXiv:2301.13688.
- Maedche, A., & Staab, S. (2001). Learning ontologies for the semantic web. In *SemWeb*.
- Marcilio, W. E., & Eler, D. M. (2020). From explanations to feature selection: assessing SHAP values as feature selection mechanism. In *2020 33rd SIBGRAPI conference on graphics, patterns and images SIBGRAPI*, (pp. 340–347). Ieee.
- Meloni, A., Angioni, S., Salatino, A., Osborne, F., Recupero, D. R., & Motta, E. (2023). Integrating conversational agents and knowledge graphs within the scholarly domain. *Ieee Access*, 11, 22468–22489.
- Miao, N., Teh, Y. W., & Rainforth, T. (2023). Selfcheck: Using llms to zero-shot check their own step-by-step reasoning. arXiv preprint arXiv:2308.00436.
- Mitra, A., Corro, L. D., Mahajan, S., Codas, A., Simoes, C., Agrawal, S., et al. (2023). Orca 2: Teaching small language models how to reason. arXiv:2311.11045.
- Mukherjee, S., Mitra, A., Jawahar, G., Agarwal, S., Palangi, H., & Awadallah, A. (2023). Orca: Progressive learning from complex explanation traces of GPT-4. arXiv:2306.02707.
- Navigli, R., Velardi, P., Cucchiarelli, A., & Neri, F. (2004). Quantitative and qualitative evaluation of the OntoLearn ontology learning system. In *COLING 2004: proceedings of the 20th international conference on computational linguistics* (pp. 1043–1050). Geneva, Switzerland: COLING, URL <https://aclanthology.org/C04-1150>.
- Neenov, Y., Piro, R., Motik, B., Horrocks, I., Wu, Z., & Banerjee, J. (2015). RDFox: A highly-scalable RDF store. In *The semantic web-ISWC 2015: 14th international semantic web conference, bethlehem, PA, USA, October 11-15, 2015, proceedings, part II 14* (pp. 3–20). Springer.
- OpenAI (2023). GPT-4 technical report. arXiv:2303.08774.
- OpenAlex (2024). OpenAlex: End-to-End process for topic classification.
- Osborne, F., & Motta, E. (2012). Mining semantic relations between research areas. In *The semantic web-ISWC 2012: 11th international semantic web conference, boston, MA, USA, November 11-15, 2012, proceedings, part I 11* (pp. 410–426). Springer.
- Osborne, F., & Motta, E. (2015). Klink-2: integrating multiple web sources to generate semantic topic networks. In *The semantic web-ISWC 2015: 14th international semantic web conference, bethlehem, PA, USA, October 11-15, 2015, proceedings, part I 14* (pp. 408–424). Springer.
- Osborne, F., & Motta, E. (2018). Pragmatic ontology evolution: reconciling user requirements and application performance. In *The semantic web-ISWC 2018: 17th international semantic web conference, Monterey, CA, USA, October 8–12, 2018, proceedings, part I 17* (pp. 495–512). Springer.
- Osborne, F., Motta, E., & Mulholland, P. (2013). Exploring scholarly data with rexplore. In H. Alani, L. Kagal, A. Fokoue, P. Groth, C. Biemann, J. X. Parreira, L. Aroyo, N. Noy, C. Welty, & K. Janowicz (Eds.), *The semantic web – ISWC 2013* (pp. 460–477). Berlin, Heidelberg: Springer Berlin Heidelberg, ISBN: 978-3-642-41335-3.
- Peng, C., Xia, F., Naseriparsa, M., & Osborne, F. (2023). Knowledge graphs: Opportunities and challenges. *Artificial Intelligence Review*, 56(11), 13071–13102.
- Pisu, A., Pompianu, L., Salatino, A., Osborne, F., Riboni, D., Motta, E., et al. (2024a). Classifying scientific topic relationships with SciBERT. vol. 3759, In *CEUR workshop proceedings*. CEUR.
- Pisu, A., Pompianu, L., Salatino, A., Osborne, F., Riboni, D., Motta, E., et al. (2024b). Leveraging language models for generating ontologies of research topics. vol. 3747, In *CEUR WORKSHOP PROCEEDINGS* (p. 11). CEUR-WS.
- Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., & Finn, C. (2023). Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 53728–53741.
- Rasch, R. F. (1987). The nature of taxonomy. *Image: Journal of Nursing Scholarship*, 19(3), 147–149.
- Rous, B. (2012). Major update to ACM's computing classification system. *Communications of the ACM*, 55(11), 12.
- Saeedizade, M. J., & Blomqvist, E. (2024). Navigating ontology development with large language models. In *European semantic web conference* (pp. 143–161). Springer.
- Salatino, A., Aggarwal, T., Mannocci, A., Osborne, F., & Motta, E. (2024). A survey on knowledge organization systems of research fields: Resources and challenges. <http://dx.doi.org/10.48550/arXiv.2409.04432>, URL <http://arxiv.org/abs/2409.04432> arXiv:2409.04432 [cs].

- Salatino, A. A., Osborne, F., Birukou, A., & Motta, E. (2019). Improving editorial workflow and metadata quality at springer nature. In *The semantic web–ISWC 2019: 18th international semantic web conference, auckland, New zealand, October 26–30, 2019, proceedings, part ii 18* (pp. 507–525). Springer.
- Salatino, A. A., Thanapalasingam, T., Mannocci, A., Birukou, A., Osborne, F., & Motta, E. (2020). The computer science ontology: A comprehensive automatically-generated taxonomy of research areas. *Data Intelligence*, [ISSN: 2641-435X] 2(3), 379–416. http://dx.doi.org/10.1162/dint_a_00055, arXiv:https://direct.mit.edu/dint/article-pdf/2/3/379/1857480/dint_a_00055.pdf.
- Salatino, A. A., Thanapalasingam, T., Mannocci, A., Osborne, F., & Motta, E. (2018). The computer science ontology: a large-scale taxonomy of research areas. In *The semantic web–ISWC 2018: 17th international semantic web conference, Monterey, CA, USA, October 8–12, 2018, proceedings, part II 17* (pp. 187–205). Springer.
- Sanderson, M., & Croft, B. (1999). Deriving concept hierarchies from text. In *Proceedings of the 22nd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 206–213).
- Sharma, P., Ash, J. T., & Misra, D. (2023). The truth is in there: Improving reasoning in language models with layer-selective rank reduction. arXiv preprint arXiv:2312.13558.
- Shen, Z., Ma, H., & Wang, K. (2018). A web-scale system for scientific knowledge exploration. arXiv preprint arXiv:1805.12216.
- Smith, A. (2020). Physics subject headings (physh). *KO Knowledge Organization*, 47(3), 257–266.
- Sun, Y., Xin, H., Sun, K., Xu, Y. E., Yang, X., Dong, X. L., et al. (2024). Are large language models a good replacement of taxonomies?. arXiv preprint arXiv:2406.11131.
- Templeton, A., Conerly, T., Marcus, J., Lindsey, J., Bricken, T., Chen, B., et al. (2024). Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. Transformer Circuits Thread.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., et al. (2023). Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288.
- Tsaneva, S., Dessì, D., Osborne, F., & Sabou, M. (2025). Knowledge graph validation by integrating LLMs and human-in-the-loop. *Information Processing & Management*.
- Velardi, P., Faralli, S., & Navigli, R. (2013). Ontolearn reloaded: A graph-based algorithm for taxonomy induction. *Computational Linguistics*, 39(3), 665–707.
- Vig, J. (2019). A multiscale visualization of attention in the transformer model. In *Proceedings of the 57th annual meeting of the association for computational linguistics: system demonstrations* (pp. 37–42). Florence, Italy: Association for Computational Linguistics, <http://dx.doi.org/10.18653/v1/P19-3007>, URL <https://www.aclweb.org/anthology/P19-3007>.
- Wang, D., & Barabási, A.-L. (2021). *The science of science*. Cambridge University Press.
- Wang, G., Cheng, S., Zhan, X., Li, X., Song, S., & Liu, Y. (2023). Openchat: Advancing open-source language models with mixed-quality data. arXiv preprint arXiv:2309.11235.
- Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., Narang, S., et al. (2022). Self-consistency improves chain of thought reasoning in language models. arXiv preprint arXiv:2203.11171.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., et al. (2023). Chain-of-thought prompting elicits reasoning in large language models. arXiv:2201.11903 URL <https://arxiv.org/abs/2201.11903>.
- Yadav, P., Tam, D., Choshen, L., Raffel, C. A., & Bansal, M. (2024). Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36.
- Ye, J., Chen, X., Xu, N., Zu, C., Shao, Z., Liu, S., et al. (2023). A comprehensive capability analysis of gpt-3 and gpt-3.5 series models. arXiv preprint arXiv:2303.10420.
- Yuan, L., Cui, G., Wang, H., Ding, N., Wang, X., Deng, J., et al. (2024). Advancing llm reasoning generalists with preference trees. arXiv preprint arXiv:2404.02078.
- Zadouri, T., Üstün, A., Ahmadian, A., Ermiş, B., Locatelli, A., & Hooker, S. (2023). Pushing mixture of experts to the limit: Extremely parameter efficient moe for instruction tuning. arXiv preprint arXiv:2309.05444.
- Zaharee, M. (2013). Building controlled vocabularies for metadata harmonization. *Bulletin of the American Society for Information Science and Technology*, 39(2), 39–42. <http://dx.doi.org/10.1002/bult.2013.1720390211>, arXiv:<https://asistdl.onlinelibrary.wiley.com/doi/pdf/10.1002/bult.2013.1720390211> URL <https://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/bult.2013.1720390211>.
- Zeng, M. L. (2008). Knowledge organization systems (KOS). *KO Knowledge Organization*, 35(2–3), 160–182.
- Zhang, B., Carriero, V. A., Schreiberhuber, K., Tsaneva, S., González, L. S., Kim, J., et al. (2024). OntoChat: a framework for conversational ontology engineering using language models. In *European semantic web conference* (pp. 102–121). Springer.
- Zhang, S., Dong, L., Li, X., Zhang, S., Sun, X., Wang, S., et al. (2023). Instruction tuning for large language models: A survey. arXiv preprint arXiv:2308.10792.
- Zhao, W. X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., et al. (2023). A survey of large language models. arXiv preprint arXiv:2303.18223, 1 (2).
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., et al. (2023). Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36, 46595–46623.
- Zhou, C., Liu, P., Xu, P., Iyer, S., Sun, J., Mao, Y., et al. (2023). Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36, 55006–55021.