

Inferring Adjective Hypernyms with Language Models to Increase the Connectivity of Open English Wordnet

Lorenzo Augello*

Università Cattolica del Sacro Cuore,
Milan, Italy.
lorenzo.augello01@icatt.it

John P. McCrae

Insight Research Ireland
Centre for Data Analytics,
University of Galway,
Galway, Ireland.
john@mccr.ae

Abstract

Open English Wordnet is a key resource published in OntoLex-lemon as part of the linguistic linked open data cloud. There are, however, many links missing in the resource, and in this paper, we look at how we can establish hypernymy between adjectives. We present a theoretical discussion of the hypernymy relation and how it differs for adjectives in contrast to nouns and verbs. We develop a new resource for adjective hypernymy and fine-tune large language models to predict adjective hypernymy, showing that the methodology of TaxoLLaMa can be adapted to this task.

1 Introduction

Open English Wordnet (McCrae et al., 2019a, OEWN) is an open-source fork of the Princeton WordNet (Fellbaum, 1998), which is modelled as and released as OntoLex data (McCrae et al., 2017). This work aims to continue the work of maintaining the resource, as well as providing a central resource for linked data resources to connect to. However, OEWN is not itself a completely connected graph, as recent work on the verb hierarchy has shown (McCrae, 2025). For adjectives and adverbs, there are a very large number of synsets¹ that are not linked and this reduces the effectiveness of this resource as a central resource in Linguistic Linked Open Data (LLOD) cloud (McCrae et al., 2016).

Recent advances in large language models have significantly reshaped the field of lexical semantics (Moskvoretskii et al., 2024b). These models, trained on huge corpora, have demonstrated an emerging ability to infer nuanced lexical relations, and in tasks of taxonomy extraction and hypernymy discovery (Bordea et al., 2016). Our goal

is to use this to fully connect the graph of OEWN 2024, however, the performance of these models has been predominantly evaluated on nouns and verbs, leaving open the question of whether they can meaningfully capture a complex relation such as hypernymy among adjectives, which is underrepresented in existing datasets and lexical resources.

Hypernymy has been extracted from large text corpora and modelled in widely used benchmarks for lexical relations classification such as BLESS (Baroni and Lenci, 2011), EvaLution (Santus et al., 2015) and HyperLex (Vulić et al., 2017). Yet, its application to adjectives remains considerably underexplored, despite its potential theoretical and practical relevance. Unlike noun and verb hypernymy, which often relies on well-defined hierarchical taxonomies, hypernymy between adjectives is inherently more fluid and context-dependent, influenced by issues such as polysemy, gradability, and contextual ambiguity (Kennedy, 2007), which are often investigated by existing resources only from other relations' points of view (Murphy, 2003).

Understanding hypernymy among adjectives is then critical and in order to explore its issues and despite their acknowledged complexity, our aim is to establish a clearer theory of adjective hypernymy that can offer insights into the broader organization of lexical meaning and support the development of more semantically aware language models.

In this work, we try to address the theoretical and practical gaps surrounding hypernymy between adjectives, with two main objectives. First, we propose a theoretical framework for interpreting hypernymy among adjectives and construct a gold-standard English dataset of hyponym-hypernym pairs reflecting this relation, which is released in RDF using the OntoLex model.² Starting from existing lexical resources for other languages such as Polish WordNet 3.0 (Maziarz et al., 2016) and

Work completed at Insight Research Ireland Centre for Data Analytics, University of Galway.

¹synonym set' or 'synset' for short are equivalent to lexical concepts in the OntoLex model

²<https://github.com/lorenzoaugello/adjective-hypernymy>

Open Dutch WordNet (Postma et al., 2016), we perform careful human annotation and validation to address issues of semantic ambiguity and ensure the reliability of the dataset for English by leveraging the OEWN. Second, we explore the ability of language models to recognise and interpret adjective hypernymy. We evaluate their performance both before and after training them on our gold standard, analysing how well they can capture this complex lexical relation³.

Examining models’ performance on fine-grained and subjective semantic relations can explore and unveil their underlying linguistic competence and biases, contributing to the ongoing effort to interpret and refine their behaviour. In order for the models to be better in understanding language dynamics, we think that the development of specialized evaluation datasets is a foundational step towards benchmarking and improving their ability to disentangle complex semantic phenomena (Putra et al., 2024).

In addition to its theoretical contribution that tries to address the lack of hierarchical structure of adjectives in English, our dataset offers several practical benefits: it fills a structural gap in lexical resources, enhancing their utility in knowledge representation and reasoning tasks; it supports improved performances in NLP tasks such as lexical entailment and word sense disambiguation; it serves as a novel benchmark for evaluating LLMs understanding of adjective semantics; it creates a foundation for possible multilingual extensions, towards better integration in the LLOD cloud.

The rest of this paper is structured as follows. In Section 2, we discuss related work on hypernymy, adjective semantics, and language model evaluation, and in Section 3, we introduce the Open English Wordnet. Section 4 presents our theoretical framework, while Section 5 describes the construction of our adjective hypernymy dataset. Section 6 focuses on the evaluation of language models on this task, including both zero-shot testing and fine-tuning experiments. Section 7 presents the results and a summary of the findings. Finally, Section 8 concludes the paper and suggests directions for future research, followed by the limitations of this study.

³<https://huggingface.co/collections/loraug/hyper-discovery-68123c85fef0889c6559e674>

2 Background and Related Work

Theorizing the nature of adjectives, their function, taxonomy, and interrelations has traditionally begun with categorizing how they modify nouns (Raskin and Nirenburg, 1995), reflecting the general agreement around the definition for which adjectives serve as “modifiers of the nouns with which they are combined” (Lyons, 1977). However, the semantic categorization of adjectives proves considerably more complex than that of nouns and verbs, as their behaviour often eludes straightforward ontological modelling (McCrae et al., 2014). While several classification attempts have been made, adjectives remain understudied in lexical semantics and lack a universal agreement on a theoretical framework. Existing classifications typically remain at a high level of general categorization, grouping adjective synsets into a few broad semantic classes known as “supersenses” to create taxonomies, as the first examples of Dixon (1982) and Hundsnurscher and Splett (1982) show.

Moving the focus to the task of lexical entailment and the specific relation of hypernymy, we face more ambiguity. Traditional definitions of hypernymy, such as the proposal that one term A is a hypernym of another term B if A’s meaning covers the meaning of B or much broader (Tjong Kim Sang, 2007), are more straightforward if applied to nouns and verbs, whose hierarchical relations and structures are clearer. However, adjectives introduce challenges related to gradability, scalar structure, and context dependency (Liu et al., 2023). Unlike nouns, which can be organized in chains where each hyponym naturally inherits properties from its hypernyms (Heyvaert, 2010), adjectives resist simple hierarchical arrangement and prefer more compact and less vertical scales. Notably, hypernymy relations between adjectives are absent from foundational lexical resources such as Princeton WordNet (Miller, 1995; Fellbaum, 1998). Instead, adjectives are organized via antonymy (e.g., wet–dry), semantic similarity (e.g., dry–arid), and pertainymy (e.g., crime–criminal), with an absence of explicit modelling of hierarchical relations that could lead to a conflation between synonymy and hypernymy in adjectival semantics (Scheible and Schulte im Walde, 2014).

The question is now around the existence of hypernymy for adjectives. Recent works suggest that lexical relations should be conceptualized not in binary terms, but rather as graded semantic phe-

nomena, with terms represented along a continuous scale according to principles of category membership (Vulić et al., 2017). However, given the fact that most of the existing lexical resources and datasets created for lexical semantics tasks derive from the structure of the Princeton WordNet - which does not incorporate adjectival hypernymy - the representation and formalization of this relation remains underdeveloped.

Some lexical resources for other languages have taken steps towards modelling adjectival hypernymy. For example, the GermaNet (Hamp and Feldweg, 1997) abandons the cluster-based approach of Princeton WordNet, adopting instead a hierarchical structuring for adjectives similar to that of nouns and verbs, where, for example, “gut” (“good”) is the hypernym of “toll” (“great”). Similar hierarchical mappings are observed in the Polish WordNet, where the original “dumbbell model” of Princeton WordNet has been transformed into a vertical hyponymy structure (Rudnicka et al., 2016). In the Open Dutch WordNet, although antonymy remains the dominant relation for adjectives, instances of hierarchical hypernymy are also introduced, such as “knotsgek, stapelgek, krankjorum, knettergek” (“very mad”) as hyponyms of “gek, dwaas” (“mad”) (Maks et al., 2008).

In the context of the NLP community, the hypernymy relation has been explored in many works, from its automatic extraction from text corpora (Hearst, 1992) to investigations around language models’ knowledge and classification capabilities (Ushio et al., 2021; Wang et al., 2021). Various evaluation studies have been conducted on language models’ performances, but they overwhelmingly focus on nouns (Camacho-Collados et al., 2018), verbs (Greco et al., 2024), or general natural language inference tasks (Madaan et al., 2024). Specific techniques such as prompting-based evaluation of hypernymy knowledge (Hanna and Mareček, 2021), dataset augmentation and benchmarking strategies (Kober et al., 2021) have been proposed, but still none of them are dedicated specifically to adjective hypernymy. Thus, while significant advances have been made in understanding lexical entailment for other parts of speech, an investigation specifically tailored to adjectives remains an unexplored field.

3 Open English Wordnet

The *Open English Wordnet* (McCrae et al., 2019b) is a comprehensive, open-source lexical resource for the English language, derived from the original Princeton WordNet (Fellbaum, 1998), and developed using the W3C OntoLex-Lemon model (McCrae et al., 2017). As an OntoLex resource, OEWN publishes lexical data on the Web in accordance with linked data principles, thereby facilitating interoperability across linguistic and semantic resources.

OEWN plays a central role in the Linguistic Linked Open Data (LLOD) cloud (McCrae et al., 2016), acting as a foundational resource for representing lexical semantics, enabling cross-linguistic comparison, and linking to other language resources. It uses the OntoLex-Lemon model as a standardized vocabulary for representing lexical information as RDF. The Global WordNet Association has developed formats for the representation of wordnets (McCrae et al., 2021), which extend the core principles of OntoLex. These formats support multiple serializations in XML, JSON(-LD), and RDF/Turtle—providing interoperability with other OntoLex resources. These formats have been adopted by OEWN and other wordnets in the Open Multilingual Wordnet (OMW) to ensure compatibility and extensibility.

In RDF, the GWA format directly encodes `LexicalEntry`, `LexicalSense`, and `LexicalConcept` using OntoLex classes, while WordNet-specific metadata and relations (e.g., sense keys, synset ordering) are modelled via a dedicated ontology at <https://globalwordnet.github.io/schemas/wn>.

4 Theoretical proposal

Even though it is true that hypernymy for adjectives is difficult to define and has unclear boundaries, we think that it is a necessary relation that could help in the semantic organization of a language by introducing a hierarchy of meaning and usage that is not limited to a simple categorization in broad classes or to fuzzy and often confused horizontal and opposition relations. The distinction with synonymy is in fact very narrow, but we think that in the scope of the same semantic field, there are adjectives with a broader meaning than others, and those need to be distinguished (e.g., cognitive - mental - immaterial; displeasing - unpleasant - negative).

The Princeton WordNet has 117,659 synsets and

84,428 hypernym-hyponym relations, but those are only for nouns and verbs, and not for adjectives and adverbs. Furthermore, there is currently no reliable dataset for the English language that represents this relation for adjectives. Facing the limited availability of high-quality training data, one of the main objectives of this study was thus the creation of a reliable initial dataset containing adjective pairs and their definitions, drawn from the OEWN, in order to address the problem of word-sense disambiguation.

Throughout the annotation and reliability evaluation of hyponym-hypernym pairs extracted from existing lexical resources, a principle of substitution was followed (see [McCarthy and Navigli \(2007\)](#) for other applications of lexical substitution): the hypernym should be substitutable in place of the hyponym in context, such that the resulting sentence preserves the original meaning at an acceptable general level, without contradiction or need for forced interpretation. While a semantic broadening is both expected and required, the sense of the hyponym needs to be included within the possible interpretations of the hypernym without ambiguity. Thus, hypernymy is characterized by inclusion of meaning, but not equivalence, distinguishing itself from synonymy. In the annotation process, if the connection between the two adjectives of the candidate pair was not perceived as direct and intuitive by the annotators, the pair was discarded, as we do not want dubious inference or multiple plausible readings. Given that hypernymy is inherently difficult to define - being vague and often subject to flexible and debatable boundaries - only those pairs deemed most reliable were retained (even though, undoubtedly, some may still be debated, highlighting once again the subjective nature of such a relation).

One of the main challenges in identifying a hypernymy relation between two adjectives lies in the polysemous nature of many of them. A first example is “cold”, which can refer either to the semantic field of temperature or to aspects of human behaviour and personality. Another case is “hard”, which in one sense refers to a physically solid material, and in another to the more abstract concept of difficulty. Adjectives’ meaning and their relations are therefore defined not only in isolation but, more importantly, depending on the associated noun and the semantic domain they evoke: “cold” can describe both an environment and a person; “hard” can qualify both a material and a problem.

In this regard, context plays a crucial role. While this is already a challenge for human annotators, it becomes even more complex for language models that have to deal with semantic disambiguation (see [Section 5](#)). Moreover, many English words can assume more than one part of speech depending on context and usage, without undergoing inflectional changes, as English has a very limited morphological richness. A first example is the word “clean”, which can function as both a verb and an adjective. A broader and more frequent case involves present and past participles, which are at times interpreted as verbs and at other times as adjectives. This ambiguity can be extended to nouns as well (the above mentioned word “cold” can have multiple adjectival meanings, but it can also be a noun: “a mild viral infection involving the nose and respiratory passages”), and given the pervasiveness of this phenomenon in English, it was deemed necessary to incorporate OEWN definitions for both hyponyms and hypernyms into the construction of the gold standard dataset (see also the experiments of [Moskvoretskii et al. \(2024b\)](#) for why definitions are crucial).

In contrast to the more straightforward case of nominal hypernymy, we think that adjectival hypernymy needs to be grounded in sense disambiguation, operationalized through substitution-based inclusion tests, and evaluated in a context. Our gold-standard dataset leverages those principles and aims to provide an initial but reliable resource for further studies of adjectives’ semantics and its modelling in computational systems.

5 Dataset creation

The construction of the gold-standard dataset was based on pre-existing lexical resources, leveraging wordnets from languages that explicitly encode adjectival hypernymy relations: the Open Dutch WordNet (ODWN) and the Polish WordNet (plWN). These two resources were chosen because, unlike the Princeton WordNet, they organize adjectives hierarchically, making them suitable starting points for our purposes.

Hyponym-hypernym pairs of adjectives were automatically extracted from those two resources by using the `wn` Python library ([Goodman and Bond, 2021](#)), and an initial pool of 450 pairs was randomly picked from the total (166 from ODWN and 284 from plWN). Only adjectives were considered, and both child hyponym and father hy-

pernym nodes were included. Each extracted pair was then manually translated into English using bilingual online dictionaries (Wiktionary and the Cambridge Dictionary), consistently selecting the first suggested translation to minimize subjective bias and arbitrary choice.

After this, each pair was reviewed individually by two annotators. If both adjectives of the pair appeared in the same synset within the OEWN - indicating synonymy rather than a hierarchical relation - the pair was discarded and the original relation was kept: for example, “difficult” and “hard”, which are hypernym-hyponym in pIWN, were not included because they are synonyms in OEWN (“not easy; requiring great physical or mental effort to accomplish or comprehend or endure”). When multiple hypernyms were available for a hyponym, if they belonged to the same OEWN synset they were kept; otherwise, only the most reliable one was chosen and agreed upon by the annotators, while the others were discarded.

So, starting from the initial 450 pairs, 148 were discarded, and the final gold standard for English comprises 302 adjective pairs: 92 sourced from ODWN, 170 from pIWN, and 40 derived from either of the two (18 ODWN, 22 pIWN) but for which a more reliable alternative hypernym was proposed and approved by the annotators (e.g., the pIWN had “effective” as hypernym for “deft”, but “skillful” was suggested and approved as an alternative and kept in the dataset). Acceptance rates were similar across resources, with a retention ratio of 0.55 for ODWN (92/166) and 0.60 for pIWN (170/284). Furthermore, as already mentioned in Section 4, for each adjective included in the gold standard, the corresponding English definition was included after retrieving it from the OEWN. This was essential to ensure that the annotated relations reflected the intended senses of the adjectives rather than possible polysemous variants.

The annotation process began with an initial small dataset of 50 adjective pairs, annotated by two annotators using a three-label classification system (yes, no, maybe). All cases of negative agreement (no-no), strong disagreement (yes-no), weaker disagreement (no-maybe) and shared doubt (maybe-maybe) were discarded, while complete agreement (yes-yes) cases were retained and partial agreement (yes-maybe) ones were further discussed (examples are shown in Table 1). From this first sample, 62% (31 adjective pairs) were selected

Annotation	hypo	hyper
yes-no	intelligent	rational
yes-maybe	limitless	vast
yes-yes	lucid	aware
no-maybe	multiple	plural
maybe-maybe	unneeded	useless
no-no	productive	rich

Table 1: Examples of agreements and disagreements in the annotation of the pairs of adjectives performed by the two annotators.

(Cohen’s kappa $\kappa = 0.65$). From a second sample of 100 pairs, 69% were retained ($\kappa = 0.64$), and from a third sample of 300 pairs, 67% (201 pairs) were selected ($\kappa = 0.61$). A third annotator was then included for the annotation of 100 pairs randomly selected from the third larger sample (300 pairs), providing additional discussion and validation (Fleiss’ kappa $\kappa = 0.48$).

Given the difficulty of semantic relations tasks which introduce a high level of subjectivity, inter-annotator agreement levels were moderate, but still acceptable, consistent and comparable across the different samples, with most disagreements involving “maybe” labels rather than direct contradictions (15 occurrences of the “maybe” label in the first sample, 24 in the second and 93 in the third).

In addition to the main gold standard where each hyponym is associated to one single exact hypernym, a second version of the dataset was later developed for model fine-tuning and evaluation (see Section 6.1). Here, synonyms of each hypernym were added based on synset membership in OEWN, in order to account for cases where multiple semantically correct hypernyms exist (see Table 2). This was motivated both by the natural multiplicity of hypernyms for a given adjective and for reducing the possible influence of definition-based prompts during the evaluation of the models.

6 Methodology

From the definition of hypernymy for adjectives to the issues in creating a benchmark dataset, it is clear that the task of lexical entailment already presents many issues for human understanding. As for the capabilities of language models, there have been attempts to assess their reliability following either a binary classification approach (choosing the correct hypernym) or a generation one (given a hyponym, predicting its most probable hypernym).

In this study, we use two models: TaxoLLaMa and SmolLM-360M-Instruct. We first test their

Dataset	hyponym-lemma	hypo_definition	hypernym-lemma	hyper_definition
single	relaxed	without strain or anxiety	calm	not agitated; without losing self-possession
multiple	relaxed	without strain or anxiety	calm, serene, tranquil, unagitated	not agitated; without losing self-possession

Table 2: One example of a pair taken from the two versions of the gold standard, first showing one single exact hypernym (“calm”) for the input hyponym (“relaxed”), and then showing its synonyms found in OEWN too.

capabilities in the hypernymy discovery task and then fine-tune them on both our benchmark datasets in order to explore their capabilities of capturing semantic knowledge from them.

TaxoLLaMa (Moskvoretskii et al., 2024a) is the finetuned version of the LLaMA-2-7b model (Touvron et al., 2023), which was trained on the WordNet dataset for 16 taxonomy-related tasks and reached SoTA results on hypernymy discovery in different domains and languages. As it was neither trained nor tested on adjective examples, however, its performance in predicting them was expected to be lower.

SmolLM-360M-Instruct is part of SmolLM, a family of state-of-the-art small models⁴. It is optimized for instruction-following tasks through supervised fine-tuning on multiple instruction datasets. Even though it is much smaller than TaxoLLaMa, its size makes it ideal for rapid fine-tuning and evaluation on specialized tasks such as adjective hypernymy discovery, enabling us to explore if smaller models can acquire semantic relations when provided with limited data.

While recent models offer significantly improved performances across many tasks, our focus was on small to mid-sized models that could be fine-tuned efficiently, guided by high-quality and focused supervision. This emphasizes practical accessibility and cost-efficiency, but future work could investigate whether larger models perform significantly better in zero-shot or few-shot settings, or whether fine-tuning continues to provide a clear advantage.

6.1 Training Details

We performed the fine-tuning using the Unsloth method⁵, which allowed us to quantize the models to 4 bits and train them using LoRA (Hu et al., 2022), reducing memory and computational requirements without affecting performance. This

approach was particularly suited for using our small datasets and enabled the fine-tuning of both smaller and larger models like TaxoLLaMa with limited hardware resources and easy accessibility⁶. We fine-tuned the models using the SFTTrainer class, conducting the training for 60 optimization steps with a learning rate of $2e-4$. The per-device batch size was set to 2, and gradient accumulation was used with 4 steps, simulating a batch size of 8. Leveraging our gold standard dataset, we fine-tuned the models on a training set composed of 211 items (70% of the dataset). The one below is a training sample, with an input question by the human user and an expected output answer by the GPT assistant, following a chat template:

(from: human) “What are the hypernyms of the hyponym: “complicated” (definition: “difficult to analyze or understand”)?”,

(from: gpt) “The hypernyms are: “difficult, hard” (definition: “not easy; requiring great physical or mental effort to accomplish or comprehend or endure”).”

The training was performed in two separate settings: first the models were fine-tuned using the original gold standard dataset (hereafter referred to as “single”) with a one-to-one correspondence between each hyponym and its relative hypernym; then a different gold standard was developed (“multiple”), implementing the single dataset with the synonyms of the hypernyms in order to have a one-to-many correspondence between each hyponym and its relative hypernyms (see Table 2). In order to have a consistent criterion and not fall into ambiguities, we relied on the OEWN and added all the adjectives pertaining to the same synsets of the original single hypernyms. This was done for two main reasons. First, given that a hyponym can have more than one hypernym, including only one correct exact hypernym in the gold standard

⁴<https://huggingface.co/blog/smollm>

⁵<https://unsloth.ai/>

⁶<https://colab.research.google.com/>

could leave out other possible candidates (e.g., “opportune” has “suitable” as its exact hypernym in the single dataset, but “appropriate” and “suited” were added in the multiple dataset, pertaining to the same synset: “meant or adapted for an occasion or use”). Secondly, the models trained on the single dataset would output only one hypernym for each hyponym when tested, while if trained with multiple possible hypernyms they would include more predictions and improve their semantic knowledge.

6.2 Prompting

Both during the zero-shot evaluation and after the fine-tuning, we used two different prompt settings: first, we just provided the models with an input hyponym adjective, and then we also gave them the hyponym definition. For the zero-shot prompting, we followed the below format originally used for TaxoLLaMa:

```
<s>[INST] «SYS» You are a helpful assistant. List all the possible words divided with a comma. Your answer should not include anything except the words divided by a comma«/SYS»
```

```
hyponym: humorous (full of or characterized by humor) | hypernyms: [/INST]
```

While after the fine-tuning we followed the below format:

```
messages = [{"from": "human", "value": "What are the hypernyms of the hyponym: "in-vigorating" (definition: "imparting strength and vitality")?"},]
```

7 Results

The original base TaxoLLaMa model reached SoTA results and scored an MRR (Mean Reciprocal Rank) of 54.39 for English, but this was trained only on verbs and nouns sampled from the WordNet-3.0 graph. So, its performance on hypernymy discovery for adjectives was expected to be lower. In order to assess this, we test it in a zero-shot setting on our test set (91 pairs, 30% of the total gold standard), and we record an MRR of 9.4 when the model is not prompted with the definition of the input hyponym, and an increase to 25.8 when the definition is given. After the fine-tuning on the multiple dataset with the synonyms, those scores improve, respectively, to 23.6 and 33.3.

The difficulty in treating adjectives and distinguishing them from other parts of speech was a major one with the base TaxoLLaMa: when evaluated on the test set in a zero-shot setting without the definition, 58% of times it gives as output only nouns, 21% only adjectives, and 21% both (by providing it with the definition, the numbers change respectively to 14, 44 and 42%). After the fine-tuning on the single dataset, the amount of predicted adjectives significantly increases: 95% without the definition and 100% with the definition (both 100% when fine-tuned on the multiple dataset). This improvement is achieved by both TaxoLLaMa and SmoLLM-360M-Instruct, as Table 3 shows. Given the difficulty with POS recognition and disambiguation, we consider this result to be almost as relevant and significant as the correct hypernym prediction.

Model	Setting	No def	With def
TaxoLLaMa	Zero-shot	0.39	0.78
TaxoLLaMa	ft-single	0.96	1.00
TaxoLLaMa	ft-multi	1.00	1.00
SmoLLM-360M-Instr	Zero-shot	0.69	0.77
SmoLLM-360M-Instr	ft-single	0.79	0.96
SmoLLM-360M-Instr	ft-multi	1.00	1.00

Table 3: F1-score performances on predicting the correct POS (ADJ), before the fine-tuning (Zero-shot), when trained on the single dataset (ft-single) and when trained on the multiple dataset (ft-multi), without and with the definition.

After the fine-tuning, the models were tested in two settings against the two different datasets, in order to first evaluate their capability of inferring the exact correct hypernym of a given hyponym (against the single dataset), and then introducing also the possibility of giving synonyms in output (against the multiple dataset).

Table 4 shows the results in the first setting, where TaxoLLaMa-ft-multi reaches the best scores both without and with the definition in the prompt. Interestingly, providing the definition does not improve the performance of TaxoLLaMa-ft-single. At first glance, this may seem surprising. However, we need to consider that this model produces only one adjective as output, and this choice is often strongly influenced by lexical overlap with words in the input definition: e.g., “extant” (defined as “still in existence; not extinct or destroyed or lost”) has “real” as its correct hypernym, but TaxoLLaMa-ft-single predicts “existent”, being biased by the vocabulary of the definition.

This variability and influence due to definitions was also one of the motivations to introduce syn-

Model	Setting	No def	With def
TaxoLLaMa	Zero-shot	0.15	0.39
TaxoLLaMa	ft-single	0.32	0.31
TaxoLLaMa	ft-multi	0.35	0.44
SmolLM-360M-Instr	Zero-shot	0.13	0.16
SmolLM-360M-Instr	ft-single	0.14	0.15
SmolLM-360M-Instr	ft-multi	0.21	0.25

Table 4: Performances evaluated against the single dataset on predicting the exact correct hypernym before the fine-tuning, when trained on the single dataset and when trained on the multiple dataset, without and with the definition. Given that apart from TaxoLLaMa-Zero-shot all the other models output a single hypernym, precision and recall are equal and only one value is reported.

onyms into the multiple gold standard (“real” and “existent” are in the same synset in OEWN). By allowing synonyms, we can consider acceptable also cases where the model predicts a semantically correct hypernym which is not exactly the one we expected. Additionally, it is important to note that the base TaxoLLaMa is trained to output a list of hypernyms, whereas TaxoLLaMa-ft-single produces only a single prediction. This structural difference increases the chances of including a correct hypernym in the output for the base model (resulting in a higher F1-score of 0.39 against 0.31). However, when evaluating only the first-ranked hypernym from the base TaxoLLaMa’s output list, its performance drops, with only 7 out of 14 correct hypernyms appearing in the first position.

The results shown in Table 5, obtained in the second evaluation setting against the multiple dataset, reveal that TaxoLLaMa-ft-multi performs better than the other models, and we observe a consistent improvement in the performance of the models when definitions are included in the prompt, unveiling the importance of incorporating them for disambiguation. Additionally, when compared to their counterparts trained on the single dataset (TaxoLLaMa-ft-single and SmolLM-ft-single), both TaxoLLaMa-ft-multi and SmolLM-ft-multi demonstrate clear gains, highlighting the benefit of allowing multiple valid hypernyms during training for capturing the nuanced nature of the hypernymy relation.

8 Conclusion

In this paper, we explore the understudied relation of adjective hypernymy, both from a theoretical and a computational perspective. We propose a definition for it, grounded in semantic inclusion

and contextual substitutability, distinguishing it from synonymy and other relations. Making use of this framework, we construct a two-version gold-standard English dataset for adjective hypernymy by adapting the lexical information stored in the Polish, Dutch and Open English wordnets. Our dataset was validated through human annotation and synset-based disambiguation, offering a reliable, small and initial benchmark for future research.

We then evaluate the capabilities of language models - the large-scale TaxoLLaMa and the smaller SmolLM-360M-Instruct - on the task of hypernymy discovery, both in zero-shot and fine-tuned settings. Our results show that the models initially struggle with adjective hypernymy, particularly with issues of POS ambiguity and semantic polysemy. However, after fine-tuning on our dataset, especially the synonym-augmented variant, their performances improve, highlighting the value of task-specific training data. We also found that providing the models with explicit definitions of the input adjectives improves their ability to identify correct hypernyms. This underscores the central role of word sense disambiguation in first identifying and then modelling adjectival meaning.

As future work, we would like to a) expand the gold-standard dataset to reach a higher coverage and more generalizability, and allow for a better theoretical description and more reliability in training and evaluating language models, b) model adjective hypernymy in OEWN, identifying and representing hypernymy relations for all adjectives in the resource, c) support the interoperability between the OntoLex and NLP communities, promoting the use of adjective hypernymy in downstream applications, d) extend language models evaluation to other semantic relations between adjectives to explore how they differ theoretically and how well they are captured computationally.

Limitations

1. Dataset size: The gold standard is limited in size, which may restrict both its theoretical completeness (in modelling the full spectrum of adjectival hypernymy) and its practical utility (for model training and evaluation, especially on unseen or ambiguous adjective pairs), constraining generalizability and lexical coverage.
2. Models evaluated: Only two language mod-

Model	Setting	No def			With def		
		P	R	F-M	P	R	F-M
TaxoLLaMa	Zero-shot	0.04	0.13	0.06	0.07	0.23	0.11
TaxoLLaMa	ft-single	0.14	0.14	0.14	0.16	0.16	0.16
TaxoLLaMa	ft-multi	0.15	0.20	0.17	0.28	0.25	0.26
SmoLLM-360M-Instr	Zero-shot	0.07	0.07	0.07	0.09	0.09	0.09
SmoLLM-360M-Instr	ft-single	0.09	0.09	0.09	0.10	0.10	0.10
SmoLLM-360M-Instr	ft-multi	0.16	0.10	0.12	0.20	0.14	0.16

Table 5: Performances evaluated against the multiple dataset on predicting a list of possible hypernyms before the fine-tuning, when trained on the single dataset and when trained on the multiple dataset, without and with the definition. For the models that output only one hypernym (TaxoLLaMa-ft-single, SmoLLM-Zero-shot and SmoLLM-ft-single, precision, recall and F-measure values are the same.

els were used, so this limits the conclusions regarding models capabilities by necessarily leaving out other existing architectures, sizes, and training data diversities.

3. Subjectivity of annotation: The annotation process, although based on carefully defined criteria and multi-annotator agreement, introduces a degree of subjectivity, which is even more accentuated by the graded, nuanced and polysemous nature of adjective hypernymy itself.
4. Language: The dataset, the theories and the evaluation are all limited to English. Consequently, some of the observed phenomena may not be generalized cross-linguistically, as the semantic behaviour of English adjectives may be different from those of other languages.

Acknowledgements

This research is supported by Taighde Éireann - Research Ireland under Grant Number SFI/12/RC/2289_P2 Insight_2, Insight SFI Centre for Data Analytics and Grant Number 13/RC/2106_P2, ADAPT SFI Research Centre.

References

- Marco Baroni and Alessandro Lenci. 2011. [How we BLESSEd distributional semantic evaluation](#). In *Proceedings of the GEMS 2011 Workshop on GEometrical Models of Natural Language Semantics*, pages 1–10, Edinburgh, UK. Association for Computational Linguistics.
- Georgeta Bordea, Els Lefever, and Paul Buitelaar. 2016. [SemEval-2016 task 13: Taxonomy extraction evaluation \(TExEval-2\)](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1081–1091, San Diego, California. Association for Computational Linguistics.
- Jose Camacho-Collados, Claudio Delli Bovi, Luis Espinosa-Anke, Sergio Oramas, Tommaso Pasini, Enrico Santus, Vered Shwartz, Roberto Navigli, and Horacio Saggion. 2018. [SemEval-2018 task 9: Hypernym discovery](#). In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 712–724, New Orleans, Louisiana. Association for Computational Linguistics.
- R. M. W. Dixon. 1982. *Where have All the Adjectives Gone?* De Gruyter Mouton, Berlin, New York.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. The MIT Press.
- Michael Wayne Goodman and Francis Bond. 2021. [Intrinsically interlingual: The wn python library for wordnets](#). In *Proceedings of the 11th Global Wordnet Conference*, pages 100–107, University of South Africa (UNISA). Global Wordnet Association.
- Candida M. Greco, Lucio La Cava, and Andrea Tagarelli. 2024. [Talking the talk does not entail walking the walk: On the limits of large language models in lexical entailment recognition](#).
- Birgit Hamp and Helmut Feldweg. 1997. [GermaNet - a lexical-semantic net for German](#). In *Automatic Information Extraction and Building of Lexical Semantic Resources for NLP Applications*.
- Michael Hanna and David Mareček. 2021. [Analyzing BERT’s knowledge of hypernymy via prompting](#). In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 275–282, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Marti A. Hearst. 1992. [Automatic acquisition of hyponyms from large text corpora](#). In *Proceedings of the 14th Conference on Computational Linguistics - Volume 2, COLING ’92*, page 539–545, USA. Association for Computational Linguistics.
- Frans Heyvaert. 2010. An outline for a semantic categorization of adjectives. In *Proceedings of the 14th EURALEX International Congress*, pages 1309–1318, Leeuwarden/Ljouwert, The Netherlands. Fryske Akademy.

- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Franz Hundsnurscher and Jochen Splett. 1982. *Grundlegung Einer Semantischen Beschreibung der Adjektive des Deutschen*, pages 16–47. VS Verlag für Sozialwissenschaften, Wiesbaden.
- Cristopher Kennedy. 2007. [Vagueness and grammar: the semantics of relative and absolute gradable adjectives](#). *Linguistics & Philosophy*, 30:1–45.
- Thomas Kober, Julie Weeds, Lorenzo Bertolini, and David Weir. 2021. [Data augmentation for hypernymy detection](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1034–1048, Online. Association for Computational Linguistics.
- Wei Liu, Ming Xiang, and Nai Ding. 2023. [Adjective scale probe: Can language models encode formal semantics information?](#) *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):13282–13290.
- John Lyons. 1977. *Semantics*. Cambridge University Press.
- Lovish Madaan, David Esiobu, Pontus Stenetorp, Barbara Plank, and Dieuwke Hupkes. 2024. [Lost in inference: Rediscovering the role of natural language inference for large language models](#).
- Isa Maks, Piek Vossen, Roxane Segers, and Hennie van der Vliet. 2008. [Adjectives in the Dutch semantic lexical database CORNETTO](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Marek Maziarz, Maciej Piasecki, Ewa Rudnicka, Stan Szpakowicz, and Paweł Kędzia. 2016. [plWordNet 3.0 – a comprehensive lexical-semantic resource](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2259–2268, Osaka, Japan. The COLING 2016 Organizing Committee.
- Diana McCarthy and Roberto Navigli. 2007. [SemEval-2007 task 10: English lexical substitution task](#). In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 48–53, Prague, Czech Republic. Association for Computational Linguistics.
- John P. McCrae. 2025. [Renovating the verb hierarchy of english wordnet](#). In *Global WordNet Conference 2025*.
- John P. McCrae, Julia Bosque-Gil, Jorge Gracia, Paul Buitelaar, and Philipp Cimiano. 2017. [The ontolex-lemon model: development and applications](#). In *Proceedings of eLex 2017*, pages 587–597.
- John P. McCrae, Christian Chiarcos, Francis Bond, Philipp Cimiano, Thierry Declerck, Gerard de Melo, Jorge Gracia, Sebastian Hellmann, Bettina Klimek, Steven Moran, Petya Osenova, Antonio Pareja-Lora, and Jonathan Pool. 2016. [The open linguistics working group: Developing the linguistic linked open data cloud](#). In *10th Language Resource and Evaluation Conference (LREC)*, pages 2435–2441.
- John P. McCrae, Michael Wayne Goodman, Francis Bond, Alexandre Rademaker, Ewa Rudnicka, and Luis Morgado Da Costa. 2021. [The globalwordnet formats: Updates for 2020](#). In *Proceedings of the 11th Global Wordnet Conference*, pages 91–99.
- John P. McCrae, Francesca Quattri, Christina Unger, and Philipp Cimiano. 2014. [Modelling the semantics of adjectives in the ontology-lexicon interface](#). In *Proceedings of the 4th Workshop on Cognitive Aspects of the Lexicon (CogALex)*, pages 198–209, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.
- John P. McCrae, Alexandre Rademaker, Francis Bond, Ewa Rudnicka, and Christiane Fellbaum. 2019a. [English WordNet 2019 – an open-source WordNet for English](#). In *Proceedings of the 10th Global Wordnet Conference*, pages 245–252, Wroclaw, Poland. Global Wordnet Association.
- John P. McCrae, Alexandre Rademaker, Francis Bond, Ewa Rudnicka, and Christiane Fellbaum. 2019b. [English wordnet 2019 – an open-source wordnet for english](#). In *Proceedings of the 10th Global WordNet Conference – GWC 2019*.
- George A. Miller. 1995. [WordNet: a lexical database for English](#). *Commun. ACM*, 38(11):39–41.
- Viktor Moskvoretskii, Ekaterina Neminova, Alina Lobanova, Alexander Panchenko, and Irina Nikishina. 2024a. [TaxoLLaMA: WordNet-based model for solving multiple lexical semantic tasks](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 2331–2350. Association for Computational Linguistics.
- Viktor Moskvoretskii, Alexander Panchenko, and Irina Nikishina. 2024b. [Are large language models good at lexical semantics? a case of taxonomy learning](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1498–1510, Torino, Italia. ELRA and ICCL.
- M. Lynne Murphy. 2003. [Semantic relations and the lexicon: antonymy, synonymy, and other paradigms](#). *Acta Linguistica Hafniensia*, 36(1):185–189.
- Marten Postma, Emiel van Miltenburg, Roxane Segers, Anneleen Schoen, and Piek Vossen. 2016. [Open Dutch WordNet](#). In *Proceedings of the 8th Global WordNet Conference (GWC)*, pages 302–310, Bucharest, Romania. Global Wordnet Association.

- I. Made Suwija Putra, Daniel Siahaan, and Ahmad Saikhu. 2024. [Recognizing textual entailment: A review of resources, approaches, applications, and challenges](#). *ICT Express*, 10(1):132–155. Publisher Copyright: © 2023 The Author(s).
- Victor Raskin and Sergei Nirenburg. 1995. [Lexical semantics of adjectives a microtheory of adjectival meaning](#).
- Ewa Rudnicka, Wojciech Witkowski, and Katarzyna Podlaska. 2016. [Challenges of adjective mapping between plWordNet and Princeton WordNet](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 2413–2418, Portorož, Slovenia. European Language Resources Association (ELRA).
- Enrico Santus, Frances Yung, Alessandro Lenci, and Chu-Ren Huang. 2015. [EVALution 1.0: an evolving semantic dataset for training and evaluation of distributional semantic models](#). In *Proceedings of the 4th Workshop on Linked Data in Linguistics: Resources and Applications*, pages 64–69, Beijing, China. Association for Computational Linguistics.
- Silke Scheible and Sabine Schulte im Walde. 2014. [A database of paradigmatic semantic relation pairs for German nouns, verbs, and adjectives](#). In *Proceedings of Workshop on Lexical and Grammatical Resources for Language Processing*, pages 111–119, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.
- Erik Tjong Kim Sang. 2007. [Extracting hypernym pairs from the web](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 165–168, Prague, Czech Republic. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Asahi Ushio, Jose Camacho-Collados, and Steven Schockaert. 2021. [Distilling relation embeddings from pretrained language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9044–9062, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ivan Vulić, Daniela Gerz, Douwe Kiela, Felix Hill, and Anna Korhonen. 2017. [HyperLex: A large-scale evaluation of graded lexical entailment](#). *Computational Linguistics*, 43(4):781–835.
- Chengyu Wang, Minghui Qiu, Jun Huang, and Xiaofeng He. 2021. [KEML: A knowledge-enriched meta-learning framework for lexical relation classification](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13924–13932.