

Do Large Language Models Understand Word Senses?

Domenico Meconi¹ Simone Stirpe¹ Federico Martelli²
Leonardo Lavalle² Roberto Navigli^{1,2}

¹Babelscape

²Sapienza NLP Group, Sapienza University of Rome
{surname}@babelscape.com, {surname}@diag.uniroma1.it

Abstract

Understanding the meaning of words in context is a fundamental capability for Large Language Models (LLMs). Despite extensive evaluation efforts, the extent to which LLMs show evidence that they truly grasp word senses remains underexplored. In this paper, we address this gap by evaluating both i) the Word Sense Disambiguation (WSD) capabilities of instruction-tuned LLMs, comparing their performance to state-of-the-art systems specifically designed for the task, and ii) the ability of two top-performing open- and closed-source LLMs to understand word senses in three generative settings: definition generation, free-form explanation, and example generation. Notably, we find that, in the WSD task, leading models such as GPT-4o and DeepSeek-V3 achieve performance on par with specialized WSD systems, while also demonstrating greater robustness across domains and levels of difficulty. In the generation tasks, results reveal that LLMs can explain the meaning of words in context up to 98% accuracy, with the highest performance observed in the free-form explanation task, which best aligns with their generative capabilities. We release our code and data at: <https://github.com/Babelscape/LLM-WSD>.

1 Introduction

Thanks to the remarkable representational power of recent neural approaches based on the Transformer architecture (Vaswani et al., 2017) and its subsequent developments, language modeling has become the cornerstone of nearly every NLP application (Chang et al., 2024): fueled by unprecedented scaling both in model size and training data, Large Language Models (LLMs) – such as Llama (Touvron et al., 2023), GPT (Brown et al., 2020; Achiam et al., 2023), and DeepSeek (Liu et al., 2024) among others – have set new performance standards in a wide range of tasks.

While extensive evaluation has been conducted

across numerous generative tasks, a crucial yet underexplored capability of LLMs is their handling of lexical ambiguity, a pervasive feature of natural language (Bevilacqua et al., 2021). Effectively resolving ambiguity is vital for generating contextually appropriate responses in NLP tasks such as translation and question answering, as misinterpretation can lead to hallucinations and reduced reliability (Guerreiro et al., 2023; Campolungo et al., 2022). Although many researchers believe that this issue is indirectly addressed by the current state-of-the-art performance across tasks, only a few studies have explicitly evaluated the models’ ability to disambiguate words in context. Importantly, recent research has revealed that LLMs struggle with ambiguous words, particularly when these are used in infrequent senses, either in translation tasks (Campolungo et al., 2022; Martelli et al., 2025) or in other disambiguation-related evaluation scenarios (Basile et al., 2024; Capone et al., 2024). However, no systematic study has yet addressed a fundamental research question: to what extent do LLMs understand the meaning of words in context? With our work, we aim to bridge this gap with the following contributions:

- We conduct an extensive experimental study on the explicit Word Sense Disambiguation (WSD) capabilities of open- and closed-weight LLMs, ranging in size from one billion to several hundred billion parameters. Our evaluation spans four distinct WSD benchmarks, including a newly curated dataset comprising over 5,500 items. We also compare the performance of the top LLM with that of an expert human annotator.
- We investigate the lexical understanding abilities of the best-performing open- and closed-source models, i.e. Llama-3.3-70B-Instruct and GPT-4o, beyond the constraints of predefined sense inventories. Specifically, we

evaluate them in three generation tasks: i) definition generation, ii) free-form explanation, iii) example generation. Our human-evaluated experiments reveal that both models demonstrate significantly enhanced understanding when allowed to define target words freely, achieving up to 98% accuracy in the most unconstrained setting.

2 Related Work

Over the course of the last few years, several benchmarks for LLMs have been proposed covering a wide range of language understanding tasks, such as the Massive Multitask Language Understanding benchmark (Hendrycks et al., 2020, MMLU), Language Model Evaluation Harness (Gao et al., 2023), BIG-bench (Srivastava et al., 2023), InstructEval (Chia et al., 2023) and the Holistic Evaluation of Language Models (Liang et al., 2023, HELM). However, despite their widespread adoption, such benchmarks do not specifically target disambiguation capabilities. Other studies have explored how LLMs handle various types of ambiguity, such as lexical, syntactic and pragmatic, via Natural Language Inference (Liu et al., 2023; Kamath et al., 2024, NLI). Nevertheless, these works do not directly examine LLMs’ ability to identify the correct sense of target words in context.

Among the research efforts which focus on this goal, Kocóń et al. (2023) analyze the performance of ChatGPT 3.5¹ across various tasks, including Word Sense Disambiguation (WSD) and Word-in-Context (Pilehvar and Camacho-Collados, 2019). However, their study is limited to a single model and does not offer a comprehensive view of LLMs’ overall disambiguation capabilities. More recently, relatively low performance in terms of WSD capabilities has been reported by both Yae et al. (2024) and Kibria et al. (2024). However, while the former relies on only approximately 800 test instances, the latter does not adopt a robust evaluation strategy capable of handling unexpected LLM behaviors, such as hallucinations, inconsistencies, and deviations from the required response format. Sumanathilaka et al. (2025) explore a prompt augmentation strategy that incorporates additional linguistic cues such as part-of-speech tags, synonyms, and usage examples; their evaluation, nonetheless, is limited to a selected subset of the FEWS test set (Blevins et al., 2021).

¹<https://openai.com/index/chatgpt/>

Most closely related to our work, Basile et al. (2024) propose a new benchmark in Italian for evaluating the disambiguation capabilities of LLMs, later extended to English by Basile et al. (2025). However, both studies have limited model coverage and rely on automatically generated data such as XL-WSD (Pasini et al., 2021) and translated glosses derived from BabelNet (Navigli and Ponzetto, 2012).

While insightful, the aforementioned studies lack robust evaluation settings and do not offer comprehensive comparisons across LLMs spanning a wide range of model sizes – from a few billion to hundreds of billions of parameters – neither do they provide human-evaluated analyses of model performance. In this paper, we bridge this gap by first testing several instruction-tuned LLMs, i) on the WSD task, ii) in comparison with state-of-the-art WSD systems, iii) in zero-, one- and few-shot scenarios. We then evaluate the best open- and closed-source LLMs in three different generation scenarios in which the models are free to express their ability to understand word senses.

3 Lexical Understanding through WSD

We now investigate our first research question (**RQ1**): what is the capability of LLMs to carry out the definition selection task, i.e., Word Sense Disambiguation, where the model is provided with a word in context along with a list of dictionary definitions for the target word?

3.1 Experimental setup

Answer Extraction. Given a prompt containing a word in context and the list of candidate dictionary definitions for that word (see Table 1 for an example), we request the model to select the most appropriate definition among those provided and compute the highest lexical overlap score between the output definition and the candidate ones. Formally, let Δ_w be the set of candidate definitions associated with a given target word w to be disambiguated, δ_σ the definition reproduced by the model and $\hat{\Delta}_w$ the set of gold definitions ($|\hat{\Delta}_w| \geq 1$). Subsequently, we tokenize² both δ_σ , resulting in δ_σ^T and all definitions contained in Δ_w , thus obtaining Δ_w^T . Then, we compute a lexical overlap score between δ_σ and all candidate definitions:

²We use the NLTK library, <https://www.nltk.org/>

Setting	Example Prompt
WSD	<p>Read the sentence: <i>I heard parts of the building above my head cracking</i>. Choose the correct dictionary definition of the word “<i>building</i>” from the options:</p> <ol style="list-style-type: none"> 1) A structure that has a roof and walls and stands more or less permanently in one place. 2) The act of constructing something. 3) The commercial activity involved in repairing old structures or constructing new ones. 4) The occupants of a building. <p>Provide as output only the correct dictionary definition.</p>

Table 1: Prompt example for WSD in a zero-shot scenario.

Model	Prompt Templates																			
	instruction-based					question-based					conversational			synonyms-based			QA-based		neg.-based	
	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15	P16	P17	P18	P19	P20
Llama-3.2-1B-it	36.7	50.3	35.4	45.7	33.9	52.5	40.9	39.1	44.4	51.2	46.8	45.5	33.6	38.2	33.2	45.1	51.6	46.6	36.7	31.2
Phi-3.5-mini-4B-it	58.5	58.2	58.9	60.0	58.0	55.8	57.1	58.0	59.6	54.3	56.9	56.3	55.2	53.8	53.2	56.5	55.9	56.3	57.6	55.8
Phi-3-small-128k-7B-it	65.7	64.6	68.6	65.1	66.6	64.2	67.3	61.1	64.8	60.0	62.9	63.3	65.7	66.2	66.6	62.4	68.1	60.4	67.3	62.6
gemma-2-9b-it	67.5	67.5	66.6	66.2	65.5	65.7	65.9	64.4	66.8	61.8	63.5	64.0	63.5	64.8	63.3	65.5	67.3	62.0	68.6	67.3
DeepSeek-V3	69.9	71.0	69.7	68.6	67.0	68.1	68.8	69.2	69.4	68.1	70.1	67.7	65.3	64.6	65.5	68.6	68.8	68.1	68.4	68.1
GPT-4o	71.6	72.1	69.9	70.1	71.9	69.7	70.1	71.6	70.8	71.2	70.1	72.1	68.6	71.2	70.8	71.4	71.2	70.6	71.6	66.8
Mean	61.6	63.9	61.5	62.6	60.5	62.7	61.7	60.6	62.6	61.1	61.7	61.5	58.6	59.8	58.8	61.6	63.8	60.7	61.7	58.6

Table 2: F1 scores obtained by representative models when evaluated in the WSD task on the SemEval-2007 dataset using 20 different prompts (see Appendix A for the list of prompts). Best result for each model in bold.

$$\delta_w^{max} = \operatorname{argmax}_{\delta_w \in \Delta_w} \frac{|\delta_\sigma^\tau \cap \delta_w^\tau|}{|\delta_w^\tau|}$$

Finally, if $\delta_w^{max} \in \hat{\Delta}_w$, we classify the answer as correct. Manual inspection of a substantial subset of model outputs confirms that the extracted answer consistently matches the model’s intended selection in over 99.9% of the cases on average.

We do not rely on perplexity to extract the model’s answer, as we also evaluate closed-source models which do not provide access to logits. Nonetheless, we conduct some experiments using perplexity-based extraction on a subset of models (see appendix B). Also, we do not utilize next-token probability, given our request to provide a full definition, and the well-known issues in the literature (Wang et al., 2024).

LLMs. We test instruction-tuned LLMs from the Gemma (Team et al., 2024), GPT (Hurst et al., 2024), DeepSeek (Liu et al., 2024), Llama (Touvron et al., 2023), Mistral (Jiang et al., 2023), Qwen (Bai et al., 2023) and Phi (Li et al., 2023) families. As **open-weight models**, we test **1-4B parameter models** (Llama-3.2-1B-Instruct, Llama-3.2-3B-Instruct, gemma-2-2b-it, gemma-3-4b-it, Phi-3.5-mini-instruct, Phi-3-small-128k-instruct); **7-12B parameter models** (Mistral-7B-Instruct-v0.3, Ministral-8B-Instruct-2410, Llama-3.1-8B-Instruct, gemma-2-9b-it, gemma-3-12b-it); **27-70B parameter models** (gemma-3-27b-it, Qwen2.5-

32B-Instruct, Llama-3.3-70B-Instruct); as **closed-source models**, we test two very large LLMs, namely, DeepSeek-V3 and GPT-4o. We test all the LLMs in zero-shot, one-shot and 3-shot settings (the latter providing examples with the same part of speech as the test item under consideration). More information about inference and hyperparameters can be found in Appendix C.

Comparison models and baseline. We test the LLMs against the highest-performance WSD systems, namely ConSec (Barba et al., 2021b), ESCHER (Barba et al., 2021a) and BEM (Blevins and Zettlemoyer, 2020), all based on Transformer models. As our baseline, we compute the Most Frequent Sense (MFS), that always selects as the answer the definition corresponding to the most frequent sense based on frequency in SemCor (Miller et al., 1993), as customary in the WSD literature.

Test sets. To evaluate the WSD capabilities of the models, we compute F1 performance on four different benchmarks, all annotated according to the WordNet sense inventory. Our primary test set is the English dataset proposed by Maru et al. (2022), a manually refined version of the popular WSD evaluation dataset by Raganato et al. (2017), featuring 4917 instances from the Senseval and SemEval competitions. We also include two supplementary datasets to test robustness across domains and difficulty levels, 42D and hardEN, also proposed by

Maru et al. (2022). 42D contains 370 instances sampled from 42 diverse domains, while hardEN is a subset of the primary test set, consisting of 476 challenging cases on which all supervised WSD models failed. Finally, to address potential concerns about contamination from publicly available benchmarks – which could compromise the validity of LLM evaluations – we include a new WSD dataset explicitly curated for this study, hereafter referred to as WikiPortal WSD. It comprises 1,000 manually annotated sentences from news sources, totaling 5549 instances, with a part-of-speech distribution that mirrors that of the Maru et al. 2022 dataset. See Appendix D for further details.

Prompt design. We design 20 distinct prompt templates that task the models to select the most appropriate definition for a given word in context. The prompts vary not only in phrasing but also in the type of instruction they convey, enabling us to assess the sensitivity of model performance to different prompting styles. Specifically, we categorize the templates into six types: instruction-based, question-based, conversational, synonyms-based, question-answer-based and negation-based (see Appendix A for the full list). Each template includes a context sentence, the target word to be disambiguated, and the list of WordNet definitions for that word. We evaluate all prompts on the SemEval-2007 dataset (Agirre and Soroa, 2007), which is typically used for development purposes. To balance coverage and feasibility, we consider two representative LLMs from each parameter cluster introduced above (1-4B, 7-12B, 27-70B, and closed-source). As shown in Table 2, there is considerable variation in model performance across different prompt templates. The gap between the best and worst-performing prompts is about 5 points on average, with Prompt 2 achieving the highest overall mean score (63.9) and Prompts 13 and 20 the lowest (58.6). This variation highlights the sensitivity of LLMs (especially smaller ones) to prompt phrasing and structure in the context of lexical disambiguation. Given its top performance across models, we select Prompt 2 for further evaluation (an instance of which is shown in Table 1).

3.2 Results

Maru et al. 2022. We first report in Table 3 the results on the Maru et al. 2022 WSD dataset (scores broken down by part of speech are available in Appendix E). Overall, and as expected, among the

Model	zero-shot	one-shot	few-shot
Baseline	65.2	65.2	65.2
Llama-3.2-1B-it	56.3	52.5	60.3
gemma-2-2b-it	72.4	72.4	73.0
Llama-3.2-3B-it	68.4	70.2	68.3
Phi-3.5-mini-4B-it	71.0	72.3	73.4
gemma-3-4b-it	<u>72.8</u>	<u>74.0</u>	<u>73.5</u>
Phi-3-small-128k-7B-it	77.8	<u>80.1</u>	79.3
Mistral-7B-it	74.3	75.9	75.3
Ministral-8B-it	76.4	75.8	75.4
Llama-3.1-8B-it	73.3	73.9	74.0
gemma-2-9b-it	<u>78.9</u>	78.4	<u>79.5</u>
gemma-3-12b-it	78.2	78.0	78.2
gemma-3-27b-it	78.8	78.6	78.8
Qwen2.5-32B-it	78.8	78.5	79.5
Llama-3.3-70B-it	<u>81.4</u>	<u>80.4</u>	<u>81.2</u>
DeepSeek-V3	81.9	82.0	82.0
GPT-4o	82.3	82.6	83.2
Mean	75.2	75.4	75.9
BEM	79.5	79.5	79.5
ESCHER	81.6	81.6	81.6
ConSeC	<u>83.0</u>	<u>83.0</u>	<u>83.0</u>

Table 3: F1 scores obtained by models when evaluated in the WSD task on the Maru et al. 2022 test set. Best LLM in bold, best results for each cluster underlined.

evaluated LLMs, GPT-4o consistently achieves the highest F1 scores, followed by DeepSeek-V3, both of which are the largest models in the pool. Among the open-weight models, gemma-3-4b-it, Phi-3-small-128k-instruct, gemma-2-9b-it and Llama-3.3-70B-Instruct stand out as the leading options. The only model that does not surpass the MFS baseline (65.2) is the smallest one, Llama-3.2-1B-Instruct.

With a few exceptions, the impact of one and few-shot is negligible across all models, with the overall mean value (fourth-to-last row in the table) increasing from 75.2 in the zero-shot setting to 75.9 in the few-shot setting. For the best-performing model, i.e., GPT-4o, we observe a statistically significant improvement from zero to few-shot by almost one percentage point ($\chi^2, p < 0.05$).

A key question is how LLMs compare to supervised systems explicitly trained for WSD. On this benchmark dataset, ConSeC and ESCHER attain 83.0 and 81.6, respectively, indicating that GPT-4o and DeepSeek-V3 are in the same ballpark, with no statistically significant difference compared to ConSeC ($\chi^2, p < 0.05$).

hardEN. We now turn to the hardEN dataset (Table 4, left column). Unsurprisingly, F1 scores drop significantly across all models with respect to the

Model	hardEN	42D
Llama-3.2-1B-it	13.9	24.3
gemma-2-2b-it	31.3	57.3
Llama-3.2-3B-it	25.6	53.5
Phi-3.5-mini-4B-it	<u>41.4</u>	<u>66.2</u>
gemma-3-4b-it	34.5	62.9
Phi-3-small-128k-7B-it	38.0	67.0
Mistral-7B-it	36.9	63.5
Ministral-8B-it	39.2	67.6
Llama-3.1-8B-it	31.5	59.5
gemma-2-9b-it	<u>43.7</u>	<u>70.0</u>
gemma-3-12b-it	37.4	68.9
gemma-3-27b-it	<u>44.1</u>	72.2
Qwen2.5-32B-it	36.8	<u>72.7</u>
Llama-3.3-70B-it	37.8	71.3
DeepSeek-V3	39.5	74.6
GPT-4o	45.6	76.8
BEM	0.0	53.2
ESCHER	0.0	<u>58.9</u>
ConSeC	<u>7.4</u>	56.6

Table 4: F1 scores obtained by models when evaluated in the WSD task on the hardEN and 42D datasets in zero-shot settings. Best LLM in bold, best results for each cluster underlined.

Maru et al. 2022 benchmark. Even the strongest LLMs struggle on this challenging dataset: GPT-4o achieves the highest score at 45.6, followed by gemma-3-27b-it at 44.1. Among the smaller models, gemma-2-9b-it (43.7) and Phi-3.5-mini-instruct (41.4) stand out as the top performers, coming close to GPT-4o. It is worth noting that gemma-2-9b-it consistently performs well across both the Maru et al. 2022 benchmark and hardEN, ranking among the top open-weight models in both settings. In contrast, Phi-3.5-mini-instruct performs moderately on the Maru et al. 2022 benchmark, but shows competitive performance in hardEN. While the performance of most LLMs is leveled around 40 F1 – compared to 0 for most of the supervised WSD systems – it is interesting to look at the complementary portion of the Maru et al. 2022 benchmark: if we exclude the hardEN instances and consider only the remaining 4582 items in Maru et al. 2022, ConSeC reaches an F1 score of 88.5, outperforming GPT-4o, which scores 85.0. Taken together, these results suggest that supervised systems like ConSeC excel on the "easier" cases but fail almost entirely on the hardest ones, whereas LLMs such as GPT-4o demonstrate a more balanced behavior across the board.

Model	zero-shot	more-context	shuffle
Baseline	65.2	65.2	65.2
Llama-3.2-1B-it	56.3	56.2	40.4
gemma-2-2b-it	72.5	72.4	68.8
Llama-3.2-3B-it	68.5	68.2	64.2
Phi-3.5-mini-4B-it	71.0	69.5	73.3
gemma-3-4b-it	<u>72.9</u>	<u>72.7</u>	<u>70.3</u>
Phi-3-small-128k-7B-it	77.8	77.7	76.3
Mistral-7B-it	74.3	74.2	71.7
Ministral-8B-it	76.5	76.6	73.9
Llama-3.1-8B-it	73.3	74.5	70.4
gemma-2-9b-it	<u>78.9</u>	<u>78.3</u>	<u>78.1</u>
gemma-3-12b-it	78.2	77.6	77.5
gemma-3-27b-it	78.8	78.5	78.3
Qwen2.5-32B-it	78.7	78.6	76.8
Llama-3.3-70B-it	<u>81.4</u>	<u>81.8</u>	<u>79.7</u>
DeepSeek-V3	81.9	82.9	80.3
GPT-4o	82.3	81.7	81.3
Mean	76.5	76.3	74.7

Table 5: F1 scores obtained by models when evaluated in the WSD task on the Maru et al. 2022 test set in three different settings: zero-shot, more-context, and shuffle. Best LLM in bold, best results per cluster underlined.

42D. We report results on 42D in the last column of Table 4. Compared to the other datasets, performance on 42D sits in the middle: scores are generally lower than on the Maru et al. 2022 benchmark, yet higher than on hardEN. GPT-4o again leads with an F1 score of 76.8, followed by DeepSeek-V3 at 74.6. Notably, Phi-3.5-mini-instruct and gemma-2-9b-it, which were already among the best-performing open-weight models on hardEN, continue to show good results, achieving 66.2 and 70.0, respectively. Interestingly, ConSeC and ESCHER show weaker performance compared to that achieved on the Maru et al. 2022 benchmark. This highlights a key strength of LLMs – namely, their wider coverage and better adaptability across domains.

Impact of context and candidate order. We carry out two additional experiments in a zero-shot setting. The first investigates the effect of adding more context by extending the input prompt to include three sentences, where the target word appears in the middle one. The second experiment examines whether the order of candidate definitions influences model performance by randomly shuffling them. Examples of the prompt used in both experiments are provided in Table 19.

Table 5 shows the impact of extended context (middle column) and shuffled definitions (last column). Interestingly, most models do not seem to

Type	Sentence	Correct	GPT-4o
Overgeneralization	In the tower, five men and women pull rhythmically on ropes attached to the same five bells that first sounded here in 1614.	an adult person who is male (as opposed to a woman)	the generic use of the word to refer to any human being
Metonymy	Ricardo Ulate said it’s not surprising that the major powers are fighting over who should bear the costs for curbing greenhouse gases, even as vulnerable countries have become more aggressive in seeking to hold the big emitters accountable for their actions.	a politically organized body of people under a single government	the people who live in a nation or country
Gross Misinterpretation	The stranger, his head seemingly sunk in thought, started to cross the street against the light just as a huge moving van roared through the intersection.	take the first step or steps in carrying out an action	move or jump suddenly, as if in surprise or alarm

Table 6: Examples of GPT-4o prediction errors.

Model	zero-shot	one-shot	few-shot
gemma-3-4b-it	78.2	77.8	78.2
gemma-2-9b-it	81.1	81.3	81.2
Llama-3.3-70B-it	81.5	81.2	81.8
GPT-4o	84.3	84.0	84.2
ConSeC	82.2	82.2	82.2

Table 7: F1 scores obtained by models when evaluated in the WSD task on our new WikiPortal WSD dataset. Best results in bold.

benefit from the additional context. For instance, GPT-4o sees a slight, non-statistically significant decrease from 82.3 to 81.7. Notable exceptions include DeepSeek-V3, which improves from 81.9 to 82.9, and Llama-3.1-8B-Instruct, which increases from 73.3 to 74.5 (both improvements are statistically significant, $p < 0.05$). More revealing are the results from the shuffled definition setting: when definitions are presented in random order, average performance across models drops from 76.5 to 74.7. The biggest drop occurs in the smallest model, Llama-3.2-1B-Instruct, showing a strong positional bias toward early options. This pattern highlights a potential limitation in some LLMs: rather than evaluating definitions purely on semantic appropriateness, they may rely on superficial heuristics such as positional order, which may compromise their reliability and generalization in unbiased WSD tasks.

Contamination check. Table 7 reports the results on our newly introduced WikiPortal WSD dataset. We compute scores for the best-performing LLMs identified in our previous evaluations, along with the top supervised system, ConSeC. Consistently with the results on 42D and hardEN, GPT-4o achieves the highest F1 scores, reaching 84.3 in zero-shot and outperforming ConSeC (a statisti-

cally significant difference). These results provide unequivocal evidence that high scores are not the result of test set contamination.

Systems vs. humans. Although top-performing LLMs approach the results of state-of-the-art WSD systems, they are still far from human performance. To quantify this gap, we conduct an experiment on a random subset of 400 items from the Maru et al. 2022 WSD test set, where we ask an expert annotator with strong background in computational lexical semantics, paid according to the standards of their European country, to select the most appropriate sense for each instance. The annotator achieves an F1 score of 91.25, compared to 82.5 for GPT-4o on the same subset. This confirms that LLMs are not yet on par with human-level WSD performance (see also our evaluation of generation tasks in Section 4).

Analysis. To better understand the limitations of the best-performing model (i.e., GPT-4o) we manually analyze its incorrect predictions and identify six common error patterns. Table 6 reports three representative cases, while additional examples for the remaining patterns are provided in Appendix F. First, the model demonstrates **overgeneralization** errors, where it selects an overly broad or generic sense of a word. This is shown in the first row of Table 6, where GPT-4o opts for a generic "human being" definition instead of "an adult person who is male". Second, some errors arise from **metonymy**, where the model confuses closely related but distinct senses of a word. For example, for the word *countries* GPT-4o selects the sense "a politically organized body of people under a single government" instead of the intended meaning "the people who live in a nation or country". Finally, some errors are instances of **gross misinterpretation**, where the

model selects a sense that clearly diverges from the context. As shown in Table 6 (last row), the model interprets *started* as "move or jump suddenly, as if in surprise or alarm", demonstrating a misunderstanding of the context.

To validate the significance of these patterns and make a more meaningful comparison with humans (e.g. see the remarks by Tedeschi et al. (2023)), we manually categorize all 70 errors as either serious or borderline (the latter essentially due to the fine granularity of the sense inventory). 25 errors were tagged as **serious**, i.e., unlikely to be made by an expert human annotator. Interestingly, when we add these 25 to the 330 correctly classified items, the total (355) almost matches the performance of our human annotator in the same subset (88.75% vs. 91.25%). To cross-check these results, we also categorize the 35 errors made by our human annotator: only 2 were considered serious, due to oversight, while the remaining were marked as borderline cases (both validations were conducted blindly by a second expert annotator). This corroborates the gap between GPT-4o (82.5%) and the human annotator (91.25%). Moreover, only a third of the 45 borderline errors made by GPT-4o were also made by our annotator, suggesting that the ambiguity introduced by fine-grained sense distinctions leads to uncertainty, such that some items are correctly disambiguated by the model, and others by the human annotator.

Answer to RQ1. We find that top LLMs, such as GPT-4o and DeepSeek-V3, match or even surpass state-of-the-art supervised WSD systems like ConSeC and ESCHER. Overall, LLMs exhibit greater robustness across datasets compared to specialized WSD models, which are outperformed on hardEN and 42D by any 4B+ LLM. However, LLMs still lag behind humans and struggle with disambiguating hard and long-tail senses.

4 Lexical Understanding through Generation

While the above evaluation is an important step in positioning LLMs relative to state-of-the-art WSD systems in a classical setting, it reveals little about their ability to demonstrate an understanding of word senses in an unconstrained context, i.e., when they are free to answer without selecting from a predefined set of definitions, thus allowing them to express their full potential. This is the second research question of our paper (**RQ2**): to what ex-

tent can state-of-the-art LLMs explain the sense of words in context? Interestingly, this also helps address criticisms and limitations related to the discretization of sense distinctions in lexical semantics (Navigli, 2009).

4.1 Experimental Setup

Previous attempts at evaluating the quality of generated definitions (Bevilacqua et al., 2020) have relied on automatic metrics to assess the similarity between model outputs and gold definitions. In contrast, we aim to provide a more robust and reliable answer to the research question outlined above. In order to do this, we randomly sample 550 items from three different datasets: the Maru et al. 2022 WSD test set, 42D, and WikiPortal WSD. We task the two best-performing open-weight and closed-source models, Llama-3.3-70B-Instruct and GPT-4o, with answering three different types of prompts aimed at: i) **definition generation**, ii) **free-form explanation**, iii) **example generation** (Table 20 shows one example for each). Then, we request an expert annotator with strong background in computational lexical semantics to tag the outputs generated by the models in the three tasks.

For the dictionary-style definition generation task (Table 11, top) we ask the annotator to tag the definitions provided by the two LLMs in comparison with the WordNet definitions as selected in the gold standard dataset. For each item, the annotator has four options: *equally correct*, *better than gold*, *worse than gold*, *wrong*. According to the guidelines, we consider a generated definition better if it more clearly captures the correct word sense, and worse if it is less appropriate or overly specific. We count as evidence of correct sense understanding those outputs tagged as *equally correct*, *better than gold* or *worse because overly specific*.

In the free explanation task (Table 11, center) the annotator has to evaluate the quality of the unconstrained explanation produced by the models for the target word in context. The evaluation options include *good*, *bad* and *borderline*. Borderline is used for those cases where the system grasps the intended meaning, but its output is either imprecise or overly tied to the example. We consider that the model has correctly understood the target word’s sense when its output is tagged as *good*.

In the example generation task (Table 11, bottom), the annotator has to assess the ability of the LLMs to generate three distinct sentences that use the target word in the same sense as in the sentence

Model	Dataset											
	Maru et al. 2022				42D				WikiPortal WSD			
	equally correct	better	worse	wrong	equally correct	better	worse	wrong	equally correct	better	worse	wrong
Llama-3.3-70B-it	69.00	4.25	18.25	8.50	59.46	2.70	27.03	10.81	68.00	8.00	15.00	9.00
GPT-4o	73.75	5.75	14.50	6.00	67.57	2.70	21.62	8.11	68.00	11.00	13.00	8.00

Table 8: Results of Llama-3.3-70B-Instruct and GPT-4o on the Definition Generation task (% of outputs tagged as equally correct, better, worse, or wrong) across three datasets.

Model	Dataset								
	Maru et al. 2022			42D			WikiPortal WSD		
	good	bad	borderline	good	bad	borderline	good	bad	borderline
Llama-3.3-70B-it	95.75	2.50	1.75	83.78	8.11	8.11	98.00	2.00	0.00
GPT-4o	96.00	3.00	1.00	97.30	2.70	0.00	97.00	1.00	2.00

Table 9: Results of Llama-3.3-70B-Instruct and GPT-4o on the Free-form Explanation task (% of outputs tagged as good, bad or borderline) across three datasets.

under examination. In this case four options could be selected: *yes* (all three sentences are accurate), *yes but same pattern* (the sentences are correct but at least one contains a repetitive pattern that closely mirrors the input sentence), *no* (at least one sentence is clearly incorrect), and *borderline* (at least one sentence is ambiguous, leaving it open to multiple interpretations). We consider that the model has correctly captured the intended sense when its output is tagged as *yes* or *yes but same pattern*. Full annotation guidelines and examples are available in Appendix G.

4.2 Results

Definition generation. Table 8 reports results for both Llama-3.3-70B-Instruct and GPT-4o across the three evaluation datasets. Overall, most of the generated dictionary-style definitions are judged to be of the same quality as the WordNet gold standard. Importantly, when models produce worse definitions, these are most often considered to be overly specific rather than incorrect or less appropriate. GPT-4o consistently outperforms Llama-3.3-70B-Instruct, with a gap of about 5% in the Maru et al. 2022 and WikiPortal WSD datasets. Specifically, in the Maru et al. 2022 test set GPT-4o achieves 90.25% satisfactory outputs (73.75% equally correct, 5.75% better, and 10.75% overly specific), while Llama-3.3-70B-Instruct reaches 85% (69.00% + 4.25% + 11.75%). The gap increases to 8% on 42D, where GPT-4o attains 81% compared to 73% for Llama-3.3-70B-Instruct, which we assume is due to the fact that GPT-4o is more robust across domains.

Free-form explanation. As shown in Table 9,

both Llama-3.3-70B-Instruct and GPT-4o achieve excellent results (up to 98% answers rated as good) when let free to explain the word’s sense, being in the same ballpark on Maru et al. 2022 and WikiPortal (with no statistically significant differences). However, as in the definition generation task, a clear gap emerges on 42D, where GPT-4o outperforms Llama-3.3-70B-Instruct by 14%. These results show that placing LLMs in their most natural generative setting, i.e., completely free to explain, unlocks their full semantic potential.

Example generation. When tasked with the generation of same-sense usages of the word, the models exhibit mixed results (see Table 10). Across the three evaluation datasets, 60-70% of the outputs successfully contained three distinct and semantically appropriate sentences. However, a significant portion (10-36%) is tagged as relying on the same structural pattern as the input sentence, indicating a lack of syntactic variety while still preserving the intended word sense. As in the free-form explanation task, Llama-3.3-70B-Instruct and GPT-4o achieve similar performance on the Maru et al. 2022 and WikiPortal WSD datasets, whereas on 42D GPT-4o outperforms Llama-3.3-70B-Instruct by 10%. These results suggest that, while LLMs effectively capture the target word sense, they sometimes struggle to demonstrate flexibility in usage in various contexts.

Inter-annotator agreement. To corroborate the validity of our evaluation, we compute inter-annotator agreement with a second independent expert annotator on a subset of 100 items across all three generative tasks. We obtain Cohen’s κ

Model	Dataset											
	Maru et al. 2022				42D				WikiPortal WSD			
	yes	same pattern	no	borderline	yes	same pattern	no	borderline	yes	same pattern	no	borderline
Llama-3.3-70B-it	61.50	31.25	6.00	1.25	70.27	10.81	16.22	2.70	58.00	36.00	5.00	1.00
GPT-4o	59.75	33.50	4.50	2.25	75.68	16.22	8.11	0.00	57.00	36.00	6.00	1.00

Table 10: Results of Llama-3.3-70B-Instruct and GPT-4o on the Example Generation task (% of outputs tagged as yes, same pattern, no, or borderline) across three datasets.

Task	Context	Gold Definition	GPT-4o Generation
Definition Generation	Kagan served as solicitor general in the Obama administration when the first legal challenges to the law were brought at the trial court level.	legal document setting forth rules governing a particular kind of activity	a system of rules created and enforced through social or governmental institutions to regulate behavior.
Free-form Explanation	In this sense, a teacher can be compared to one’s family doctor.	a person whose occupation is teaching	In the given sentence, the word "teacher" is being compared to a "family doctor," suggesting that a teacher, like a family doctor, plays a critical and foundational role in someone’s life.
Example Generation	Ringers , she added, are ""filled with the solemn intoxication that comes of intricate ritual faultlessly performed.""	a person who rings church bells (as for summoning the congregation)	The chess grandmasters sat in silence around the board, ringers fully immersed in the complexity and precision of their strategic moves, each one leading to a crescendo of intellectual triumph.

Table 11: Examples of GPT-4o generation errors.

values of 0.791 for definition generation, 0.796 for free explanation, and 0.919 for example generation, indicating substantial to near-perfect agreement.

Analysis. As in Section 3, we focus our error analysis on the type of errors made by GPT-4o across the three tasks. In generating definitions, errors include gross hallucinations as in defining *law* as "a system of rules created and enforced through social or governmental institutions to regulate behavior" in a context in which "legal document setting forth rules governing a particular kind of activity" is the right sense (Table 11, top).

In the few errors observed in the free-form explanation task, GPT-4o tends not to provide an actual explanation, but rather to talk around the meaning, as in the example shown in Table 11, middle, for the word *teacher*.

When prompting the model to generate new examples, we still find gross hallucinations. For instance, considering the word *ringer* (Table 11, bottom), GPT-4o places the word in a context that does not align with its intended sense, failing to capture the correct sense. In other cases, it fails to follow the prompt instructions, generating sentences where the target word is missing altogether.

Answer to RQ2. Removing the constraint of selecting from a predefined set of definitions allows state-of-the-art LLMs to explain the sense of words

in context with considerably higher accuracy, above 90% across the three generative tasks in an all-words scenario. LLMs reach up to 96-98% in the free-form generation task, which aligns most naturally with their generative capabilities.

5 Conclusion

In this paper, we present a comprehensive study on the ability of instruction-tuned Large Language Models to understand word meanings in context. We first evaluate a wide range of LLMs – both open and closed-weight – on four different Word Sense Disambiguation (WSD) benchmarks, comparing their performance to state-of-the-art supervised systems. Top LLMs, such as GPT-4o and DeepSeek-V3, achieve performance comparable to specialized WSD models like ConSeC and ES-CHER. Moreover, LLMs exhibit greater robustness on more challenging datasets, confirming their adaptability and reliability in the face of lexical ambiguity. Beyond classical disambiguation, we assess whether an open- or closed-source state-of-the-art model can demonstrate its semantic understanding in an unconstrained setting. Our analysis indicates that, regardless of the model or dataset, removing the constraints of predefined inventories allows the model to more accurately express the contextual meaning of words up to 98% accuracy.

6 Limitations



All experiments were conducted exclusively in English. We leave the evaluation in different languages to future work. We note, however, that this goes beyond the scope of our paper, because for WSD very little data that is manually annotated exists in other languages. This pairs with the fact that in languages other than English sense inventories are incomplete, which would make the results less robust (e.g. due to missing senses not given as additional potential options).

More LLMs could be tested, which we leave to future work. However, we posit that we evaluated a wide range of LLMs by parameter size and that the key findings in our paper are not negatively affected by this.

The generation component was evaluated only manually, which – while avoiding well-known issues with the automated assessment of generated text – may limit reproducibility. However, we release all the human answers for verification purposes, and in order to enable testing of future automatic evaluation approaches.

Finally, while our prompt design study explored a range of prompting strategies, we observed that performance varied across different formulations. This suggests that prompt sensitivity may play a more significant role than initially assumed, and future work could benefit from a deeper exploration of adaptive and task-specific prompting techniques.

Acknowledgements

We gratefully acknowledge the support of the PNRR MUR project   PE0000013-FAIR.

We gratefully acknowledge the CREATIVE project (CRoss-modal understanding and gENERATION of Visual and tEXtual content), which is funded by the MUR Progetti di Ricerca di Rilevante Interesse Nazionale programme (PRIN 2020).

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.

Eneko Agirre and Aitor Soroa. 2007. [Semeval-2007 task 02: Evaluating word sense induction and discrimination systems](#). In *Proceedings of the fourth*

international workshop on semantic evaluations (semeval-2007), pages 7–12.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. [Qwen technical report](#). *arXiv preprint arXiv:2309.16609*.

Edoardo Barba, Tommaso Pasini, and Roberto Navigli. 2021a. [Esc: Redesigning wsd with extractive sense comprehension](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4661–4672.

Edoardo Barba, Luigi Procopio, and Roberto Navigli. 2021b. [Consec: Word sense disambiguation as continuous sense comprehension](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1492–1503.

Pierpaolo Basile, Elio Musacchio, and Lucia Siciliani. 2024. [Ita-sense-evaluate llms’ ability for italian word sense disambiguation: A calamita challenge](#). In *Proceedings of the 10th Italian Conference on Computational Linguistics (CLiC-it 2024)*, Pisa, Italy.

Pierpaolo Basile, Lucia Siciliani, Elio Musacchio, and Giovanni Semeraro. 2025. [Exploring the word sense disambiguation capabilities of large language models](#). *Preprint*, arXiv:2503.08662.

Michele Bevilacqua, Marco Maru, Roberto Navigli, et al. 2020. [Generatory or “how we went beyond word sense inventories and learned to gloss”](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7207–7221.

Michele Bevilacqua, Tommaso Pasini, Alessandro Raganato, and Roberto Navigli. 2021. [Recent trends in word sense disambiguation: A survey](#). In *International Joint Conference on Artificial Intelligence*, pages 4330–4338. International Joint Conference on Artificial Intelligence, Inc.

Terra Blevins, Mandar Joshi, and Luke Zettlemoyer. 2021. [FEWS: Large-scale, low-shot word sense disambiguation with the dictionary](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 455–465, Online. Association for Computational Linguistics.

Terra Blevins and Luke Zettlemoyer. 2020. [Moving down the long tail of word sense disambiguation with gloss informed bi-encoders](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1006–1017, Online. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.

- Niccolò Campolungo, Federico Martelli, Francesco Saina, and Roberto Navigli. 2022. [Dibimt: A novel benchmark for measuring word sense disambiguation biases in machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4331–4352.
- Luca Capone, Serena Auriemma, Martina Miliani, Alessando Bondielli, and Alessandro Lenci. 2024. [Lost in disambiguation: How instruction-tuned llms master lexical ambiguity](#).
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2024. [A survey on evaluation of large language models](#). *ACM Transactions on Intelligent Systems and Technology*.
- Yew Ken Chia, Pengfei Hong, Lidong Bing, and Soujanya Poria. 2023. [Instructeval: Towards holistic evaluation of instruction-tuned large language models](#). *arXiv preprint arXiv:2306.04757*.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2023. [A framework for few-shot language model evaluation](#).
- Nuno Miguel Guerreiro, Duarte M. Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre Colombo, and André F. T. Martins. 2023. [Hallucinations in large multilingual translation models](#). *Trans. Assoc. Comput. Linguistics*, 11:1500–1517.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. [Gpt-4o system card](#). *arXiv preprint arXiv:2410.21276*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Gaurav Kamath, Sebastian Schuster, Sowmya Vajjala, and Siva Reddy. 2024. [Scope ambiguities in large language models](#). *Transactions of the Association for Computational Linguistics*, 12:738–754.
- Raihan Kibria, Sheikh Intiser Uddin Dipta, and Muhammad Abdullah Adnan. 2024. [On functional competence of LLMs for linguistic disambiguation](#). In *Proceedings of the 28th Conference on Computational Natural Language Learning*, pages 143–160, Miami, FL, USA. Association for Computational Linguistics.
- Jan Kocoń, Igor Cichecki, Oliwier Kaszyca, Mateusz Kochanek, Dominika Szydło, Joanna Baran, Julita Bielaniewicz, Marcin Gruza, Arkadiusz Janz, Kamil Kanclerz, et al. 2023. [Chatgpt: Jack of all trades, master of none](#). *Information Fusion*, 99:101861.
- Cheng Li, Jindong Wang, Yixuan Zhang, Kaijie Zhu, Wenxin Hou, Jianxun Lian, Fang Luo, Qiang Yang, and Xing Xie. 2023. [Large language models understand and can be enhanced by emotional stimuli](#). *Preprint*, arXiv:2307.11760.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2023. [Holistic evaluation of language models](#). *Transactions on Machine Learning Research*.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. [Deepseek-v3 technical report](#). *arXiv preprint arXiv:2412.19437*.
- Alisa Liu, Zhaofeng Wu, Julian Michael, Alane Suhr, Peter West, Alexander Koller, Swabha Swayamdipta, Noah A Smith, and Yejin Choi. 2023. [We’re afraid language models aren’t modeling ambiguity](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 790–807.
- Federico Martelli, Stefano Perrella, Niccolo Campolungo, Tina Munda, Svetla Koeva, Carole Tiberius, and Roberto Navigli. 2025. [Dibimt: A gold evaluation benchmark for studying lexical ambiguity in machine translation](#). *Computational Linguistics*.
- Marco Maru, Simone Conia, Michele Bevilacqua, and Roberto Navigli. 2022. [Nibbling at the hard core of word sense disambiguation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4724–4737.
- George A Miller, Claudia Leacock, Randeep Tenji, and Ross T Bunker. 1993. [A semantic concordance](#). In *Human Language Technology: Proceedings of a Workshop Held at Plainsboro, New Jersey, March 21-24, 1993*.
- Roberto Navigli. 2009. [Word sense disambiguation: A survey](#). *ACM computing surveys (CSUR)*, 41(2):1–69.
- Roberto Navigli and Simone Paolo Ponzetto. 2012. [Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network](#). *Artificial intelligence*, 193:217–250.
- Tommaso Pasini, Alessandro Raganato, and Roberto Navigli. 2021. [XL-WSD: an extra-large and cross-lingual evaluation framework for word sense disambiguation](#). In *Proc. of Thirty-Fifth AAAI Conference*

on Artificial Intelligence, Virtual Event, February 2-9, 2021, pages 13648–13656. AAAI Press.

Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. [Wic: the word-in-context dataset for evaluating context-sensitive meaning representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1267–1273.

Alessandro Raganato, Jose Camacho-Collados, Roberto Navigli, et al. 2017. [Word sense disambiguation: a unified evaluation framework and empirical comparison](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 99–110.

Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.

T. G. D. K. Sumanathilaka, Nicholas Micallef, and Julian Hough. 2025. [Can LLMs assist with ambiguity? a quantitative evaluation of various large language models on word sense disambiguation](#). *Preprint*, arXiv:2411.18337.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. [Gemma 2: Improving open language models at a practical size](#). *arXiv preprint arXiv:2408.00118*.

Simone Tedeschi, Johan Bos, Thierry Declerck, Jan Hajič, Daniel Herscovich, Eduard Hovy, Alexander Koller, Simon Krek, Steven Schockaert, Rico Sennrich, Ekaterina Shutova, and Roberto Navigli. 2023. [What’s the meaning of superhuman performance in today’s NLU?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12471–12491, Toronto, Canada. Association for Computational Linguistics.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *Advances in neural information processing systems*, 30.

Xinpeng Wang, Bolei Ma, Chengzhi Hu, Leon Weber-Genzel, Paul Röttger, Frauke Kreuter, Dirk Hovy, and

Barbara Plank. 2024. ["My answer is C": First-token probabilities do not match text answers in instruction-tuned language models](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 7407–7416. Association for Computational Linguistics.

Jung H Yae, Nolan C Skelly, Neil C Ranly, and Phillip M LaCasse. 2024. [Leveraging large language models for word sense disambiguation](#). *Neural Computing and Applications*, pages 1–18.

A Prompts

We report below the 20 prompt templates used in our evaluation of instruction-tuned LLMs on the Word Sense Disambiguation task. These prompts are organized into six categories, each reflecting a different prompting strategy or style. For ease of reference, each group of prompts is presented in a dedicated table: instruction-based (Table 13), question-based (Table 14), conversational-based (Table 15), synonyms-based (Table 16), question-answer-based (Table 17) and negation-based (Table 18).

Table 19 reports an example for more-context and shuffled-definitions experiments.

Table 20 shows an example for each of the three generation tasks.

B Perplexity

In our evaluation, we primarily rely on lexical overlap between generated and gold definitions to extract model answers, as this method is straightforward and applicable to all models. However, to validate this choice, we performed additional experiments computing perplexity-based scores for a subset of open models where such access was possible.

In this approach, for each candidate definition, we compute the perplexity of the model generating that definition given the zero-shot prompt. The definition with the lowest perplexity is selected as the model’s answer. This method captures the model’s confidence in the candidate’s fit to the context.

Table 12 reports F1 scores for both extraction methods across several models in zero-shot setting. The results show that the scores based on lexical overlap and perplexity are very similar, with minimal differences for all models tested. This strong agreement supports the reliability of lexical overlap as a proxy for answer extraction.

C Implementation Details

For smaller open-source models, all experiments were conducted on a single Nvidia 4090 GPU with 40 GB of RAM, using the *transformers* library and its *pipeline* method provided by HuggingFace³. For larger models, including those with 12B parameters or more, we relied on one node of a high-performance computing (HPC) cluster equipped with four A100 GPUs. In this case, we used the vLLM library⁴ for efficient batched inference at scale. For GPT-4o and DeepSeek-V3, inference was performed using the OpenAI API⁵. We set the maximum number of output tokens to 75 and left all other hyperparameters, such as temperature and top-*k*, at their default values.

Due to resource constraints, all experiments were run once without repeated trials.

D WikiPortal WSD Dataset

To mitigate potential test set contamination, we created a new Word Sense Disambiguation (WSD) dataset using 1000 manually selected sentences, totaling 5549 instances, sourced from the Wikipedia Current Events Portal⁶. This portal aggregates recent news from a variety of international outlets, providing real-world, up-to-date content. The sources commonly used include Reuters, BBC News, Associated Press, Al Jazeera, The Guardian, and other international news organizations. Sentences were sampled from the last five years across domains and annotated with the appropriate sense using WordNet as the reference lexical resource. Annotations were carried out by an expert human annotator, paid according to standards in their European country. We also confirmed the quality of annotations by asking a second annotator to sense tag a random subsample of 200 items, computing substantial κ agreement, $p < 0.05$. For confidentiality and to prevent potential test set contamination in future research, we will not make this dataset publicly available. However, we remain available to support validation of the dataset upon request.

E WSD Results by Part of Speech

Table 21 reports the F1 scores obtained by all evaluated models on the WSD task on the Maru

³<https://huggingface.co/docs/transformers/index>

⁴<https://docs.vllm.ai/en/latest/>

⁵<https://openai.com/index/openai-api/>

⁶https://en.wikipedia.org/wiki/Portal:Current_events

et al. 2022 test set. Results are shown for three prompt settings: *zero-shot*, *one-shot*, and *few-shot*, and are broken down by part of speech (NOUN, ADJ, VERB, ADV) as well as overall performance (ALL).

F Examples of GPT-4o WSD Errors

Table 22 reports one representative example for each of the six error types identified in our analysis of the 70 incorrect predictions made by GPT-4o. Categories classified as serious errors, i.e., unlikely to be made by an expert human annotator, are: *Gross Misinterpretation*, *Literal vs. Metaphorical*, and *Metonymy*.

G Human Annotation Guidelines and Examples

To provide transparency into the evaluation process, we report here the full annotation guidelines used for each of the three generative tasks: definition generation, free-form explanation, and example generation.

Definition Generation. Annotators are evaluating the output of a computational linguistics experiment on the ability of LLMs to understand word meanings. For each instance, annotators are asked to assess the quality of a generated definition relative to a gold (reference) definition provided for a target word in context. The task consists in determining how well the generated definition captures the intended meaning of the word as it appears in its original usage. Annotators must choose one of four possible options:

Equal: The generated definition is as appropriate as the gold definition. Both definitions capture the correct meaning of the word in context to a similar degree.

Better: The generated definition conveys the meaning of the word in context more clearly and accurately than the gold standard definition. It may also be more precise or informative.

Worse: The generated definition is less appropriate than the gold definition. It may be vague, overly specific, partially incorrect, or less clear in conveying the intended meaning of the word.

Wrong: The generated definition is incorrect, misleading, or not related to the meaning of the word in context.

Examples of annotations for each label are shown in Table 23.

Free-form Explanation. In this task, annotators are required to evaluate the quality of models’ unconstrained explanation of a target word as used in context. For each item, the annotator must determine how well the explanation conveys the correct meaning of the target word in its specific context. The following three options are available:

Good: This option must be selected when the explanation clearly and accurately conveys the correct meaning of the target word in context.

Bad: This option must be selected when the explanation does not capture the intended meaning of the word in context, or when the model fails to provide a clear definition of the word used in the given sentence.

Borderline: This option is appropriate when the explanation demonstrates partial understanding: the model appears to grasp the correct sense of the word, but the output is imprecise, vague, or unsatisfactory.

Examples of annotations for each label are shown in Table 24.

Example Generation. In this task, the annotator has to evaluate whether an LLM can generate three distinct example sentences that correctly use the target word with the same meaning it has in a provided reference sentence. The goal is to assess both the semantic accuracy of the examples and the variety of their syntactic or contextual formulations. The annotator must select one of the following four options:

Yes: This option has to be selected when all three generated sentences use the target word correctly, maintaining the same meaning as in the reference sentence.

Yes but same pattern: This option must be used when all three sentences are accurate in terms of word meaning, but follow a repetitive or overly similar pattern, closely resembling the structure or context of the original sentence.

No: This option must be used when at least one of the three sentences clearly uses the word with a different meaning from that in the reference, or the usage is incorrect, misleading or contextually inappropriate.

Borderline: This option has to be selected when at least one sentence is ambiguous, meaning that the use of the target word could be interpreted as reflecting a different sense, even if the sentence is grammatically correct.

Examples of annotations for each label are

Model	lexical-overlap	perplexity
Llama-3.2-1B-it	56.3	59.4
gemma-2-2b-it	72.5	72.4
Llama-3.2-3B-it	68.5	68.3
Phi-3.5-mini-4B-it	71.0	71.0
gemma-3-4b-it	72.9	72.7
Phi-3-small-128k-7B-it	77.8	77.9
Mistral-7B-it	74.3	74.4
Minstral-8B-it	76.5	75.6
Llama-3.1-8B-it	73.3	74.9
gemma-2-9b-it	78.9	79.0

Table 12: F1 scores obtained by models when evaluated in the WSD task on the Maru et al. 2022 test set using two answer extraction methods: lexical-overlap and perplexity

shown in Table 25.

Prompt template

- 1) Consider the following sentence: *<sentence>* Select the most suitable dictionary definition which identifies the meaning of "*<word>*" among the following definitions: *<definitions>*. Provide as output only the most suitable dictionary definition.
 - 2) Read the sentence: *<sentence>* Choose the correct dictionary definition of the word "*<word>*" from the options: *<definitions>*. Provide as output only the correct dictionary definition.
 - 3) In the sentence "*<sentence>*", determine which of the following definitions best represents the meaning of "*<word>*": *<definitions>*. Provide as output only the best definition.
 - 4) Given the sentence: "*<sentence>*", select the most appropriate meaning of the word "*<word>*" from the choices: *<definitions>*. Provide as output only the most appropriate definition.
 - 5) Consider the sentence: *<sentence>* Identify the most accurate definition for "*<word>*" from the following options: *<definitions>*. Provide as output only the best definition.
-

Table 13: Prompt templates formulated as explicit instructions.

Prompt template

- 6) What is the correct definition of "*<word>*" in the sentence: "*<sentence>*"? Choose from: *<definitions>*. Provide as output only the correct definition.
 - 7) In the sentence "*<sentence>*", which of the following definitions best describes the meaning of "*<word>*"? *<definitions>*. Provide as output only the best definition.
 - 8) Which of the following definitions correctly matches the meaning of "*<word>*" in "*<sentence>*"? *<definitions>*. Provide as output only the best definition.
 - 9) Based on the context of "*<sentence>*", what does "*<word>*" mean? Select from: *<definitions>*. Provide as output only the best definition.
-

Table 14: Prompt templates phrased as direct questions.

Prompt template

- 10) If I ask you what "*<word>*" means in "*<sentence>*", how would you answer? Choose from: *<definitions>*. Provide as output only the answer.
 - 11) You are teaching vocabulary to a student. How would you explain the meaning of "*<word>*" in "*<sentence>*"? Select from: *<definitions>*. Provide as output only the best definition.
 - 12) Imagine you are explaining the meaning of "*<word>*" in the sentence "*<sentence>*", which definition would you choose? *<definitions>*. Provide as output only the chosen definition.
-

Table 15: Prompt templates designed in a conversational or instructional-dialogue style.

Prompt template

13) In the sentence "<sentence>", the word "<word>" has a specific meaning. First, think about some synonyms that match its usage in this context. Then, choose the most appropriate definition from the following: <definitions>. Provide as output only the best definition.

14) Consider the sentence <sentence> Before selecting the correct meaning of "<word>", think of a few synonyms that fit the context. Then, choose the best definition from: <definitions>. Provide as output only the best definition.

15) Analyze the sentence <sentence> Before determining the best definition for "<word>", think of a few synonyms that could replace it in this context. Finally, choose the correct meaning from: <definitions>. Provide as output only the best definition.

Table 16: Prompt templates that encourage the model to consider synonyms before selecting the most appropriate definition.

Prompt template

16) Question: What is the correct meaning of "<word>" in the sentence: "<sentence>"? Answers: <definitions>. Provide as output only the correct definition.

17) Question: In the sentence "<sentence>", which definition best represents the meaning of "<word>"? Answers: <definitions>. Provide as output only the best definition.

18) Question: Which of the following definitions correctly explains the meaning of "<word>" in "<sentence>"? Answers: <definitions>. Provide as output only the best definition.

Table 17: Prompt templates structured in a labeled question–answer format.

Prompt template

19) In the sentence "<sentence>", the word "<word>" has a specific meaning. First filter out definitions that do not fit this context, and then select the correct one from the remaining choices: <definitions>. Provide as output only the best definition.

20) Consider the word "<word>" in the sentence: <sentence> Identify and exclude the incorrect definitions and then choose the best one: <definitions>. Provide as output only the best definition.

Table 18: Prompt templates that ask the model to first eliminate inappropriate definitions before identifying the correct meaning of the word in context.

Setting	Example Prompt
more-context	Read the sentence: <i>I rode it out on the second floor of Leo’s at 55th and Telegraph in Oakland. I heard parts of the building above my head cracking. I actually thought that I might die.</i> Choose the correct dictionary definition of the word “building” from the options: 1) A structure that has a roof and walls and stands more or less permanently in one place. 2) The act of constructing something. 3) The commercial activity involved in repairing old structures or constructing new ones. 4) The occupants of a building.
shuffled-definitions	Read the sentence: <i>I heard parts of the building above my head cracking.</i> Choose the correct dictionary definition of the word “building” from the options: 1) The commercial activity involved in repairing old structures or constructing new ones. 2) A structure that has a roof and walls and stands more or less permanently in one place. 3) The act of constructing something. 4) The occupants of a building.

Table 19: Prompt examples (*zero-shot*) for the WSD task with more context and shuffled definitions.

Task	Example Prompt
Definition Generation	Consider the following sentence: <i>they needed rugs to cover the bare floors.</i> Provide a dictionary-style definition which identifies the meaning of “floor” in the above sentence. Do not motivate your answer and do not refer to the sentence context.
Free Explanation	Explain in your own words the meaning of the word “floor” in the following sentence: <i>they needed rugs to cover the bare floors.</i>
Example Generation	Provide three diverse examples that show me your understanding of the word “floor” used in the same meaning as the following sentence: <i>they needed rugs to cover the bare floors.</i> List only the sentences.

Table 20: Prompt example for the three generation tasks.

Model	Definition Selection														
	<i>zero-shot</i>					<i>one-shot</i>					<i>few-shot</i>				
	NOUN	ADJ	VERB	ADV	ALL	NOUN	ADJ	VERB	ADV	ALL	NOUN	ADJ	VERB	ADV	ALL
Llama-3.2-1B-Instruct	58.92	65.97	44.51	53.06	56.34	56.08	55.22	44.25	38.27	52.53	64.26	69.70	47.35	43.88	60.30
gemma-2-2b-it	75.21	79.10	61.77	<u>70.92</u>	72.48	75.45	<u>81.04</u>	59.82	<u>70.92</u>	72.44	<u>76.82</u>	81.79	59.20	<u>67.86</u>	73.09
Llama-3.2-3B-Instruct	72.78	75.52	55.58	54.08	68.46	73.47	80.60	58.14	55.61	70.21	72.37	75.22	56.46	52.04	68.29
Phi-3.5-mini-instruct	75.25	77.46	57.79	62.76	71.04	76.51	80.00	57.96	66.33	72.32	76.41	<u>82.24</u>	61.59	66.84	73.42
gemma-3-4b-it	<u>76.21</u>	<u>79.40</u>	<u>62.04</u>	62.76	<u>72.85</u>	<u>77.85</u>	80.15	<u>61.86</u>	67.35	<u>74.07</u>	76.65	81.64	<u>61.68</u>	65.82	<u>73.46</u>
Phi-3-small-128k-instruct	81.38	<u>85.22</u>	65.31	69.90	77.75	<u>82.88</u>	<u>87.76</u>	<u>69.91</u>	71.94	<u>80.13</u>	<u>82.16</u>	86.57	68.58	73.47	79.30
Mistral-7B-Instruct-v0.3	76.96	80.45	64.87	67.86	74.29	78.60	83.58	65.40	70.41	75.92	77.47	84.93	64.42	71.94	75.27
Minstral-8B-Instruct-2410	78.64	84.33	66.73	<u>72.96</u>	76.45	78.12	82.69	65.75	74.49	75.76	77.82	82.39	65.13	74.49	75.39
Llama-3.1-8B-Instruct	75.73	81.94	62.04	71.43	73.26	76.65	83.43	61.77	68.88	73.85	76.99	83.58	60.80	70.92	73.93
gemma-2-9b-it	<u>81.55</u>	84.18	<u>70.35</u>	71.43	<u>78.93</u>	80.90	83.88	<u>69.91</u>	72.45	78.44	81.44	85.82	<u>71.59</u>	<u>75.00</u>	<u>79.52</u>
gemma-3-12b-it	81.21	83.88	68.32	70.92	78.20	80.18	85.37	68.14	<u>76.02</u>	77.95	80.55	<u>86.87</u>	67.70	73.98	78.20
gemma-3-27b-it	80.83	86.72	70.88	67.35	78.81	81.10	<u>86.57</u>	69.47	66.33	78.58	81.24	84.93	70.09	71.94	78.81
Qwen2.5-32B-Instruct	81.65	84.18	70.53	64.80	78.77	81.79	82.09	69.47	69.39	78.50	82.78	82.99	70.44	69.90	79.46
Llama-3.3-70B-Instruct	<u>83.84</u>	<u>87.31</u>	<u>72.57</u>	<u>76.02</u>	<u>81.41</u>	<u>83.02</u>	85.07	<u>71.95</u>	<u>73.98</u>	<u>80.39</u>	<u>82.99</u>	<u>86.42</u>	<u>73.54</u>	<u>80.61</u>	<u>81.19</u>
DeepSeek-V3	84.39	87.76	72.83	<u>76.53</u>	81.88	84.32	87.61	73.27	<u>77.55</u>	81.96	84.46	86.72	73.45	<u>77.55</u>	81.96
GPT-4o	85.59	84.78	<u>74.51</u>	69.90	<u>82.31</u>	85.66	85.82	75.13	70.41	<u>82.65</u>	85.72	86.57	76.19	73.47	83.16
ESCHER	84.18	85.82	72.92	78.06	81.57	84.18	85.82	72.92	78.06	81.57	84.18	85.82	72.92	78.06	81.57
ConSec	<u>85.41</u>	<u>86.56</u>	75.13	80.61	83.01	<u>85.41</u>	<u>86.56</u>	75.13	80.61	83.01	<u>85.41</u>	<u>86.56</u>	<u>75.13</u>	80.61	<u>83.01</u>

Table 21: F1 scores obtained by models, broken down by part of speech (POS), when evaluated on the WSD task on the Maru et al. 2022 test set. The best result for each POS is shown in bold, and best result within each cluster is underlined.

Type	Sentence	Correct	GPT-4o
Overgeneralization	When functioning normally, they make proteins that hold a cell's growth in check.	(biology) the process of an individual organism growing organically;	a process of becoming larger or longer or more numerous or more important
Over specificity	To the extent that Democratic legislators from the South have held a disproportionate share of power in Congress since 1932 and have been able to translate such clout into relatively more local benefits for their respective constituencies ...	have or possess, either in a concrete or an abstract sense	have rightfully; of rights, titles, and offices
Borderline	The damage brought about to the principles of the sport by these actions is not as much ethical as aesthetic.	a basic truth or law or assumption	a rule or standard especially of good behavior
Gross Misinterpretation	If they put a Republican into office, not only will they acquire less in terms of local benefits but their selected legislator will be relatively powerless to prevent other legislators from "bringing home the bacon" to their respective constituencies.	financial assistance in time of need	something that aids or promotes well-being
Literal vs. Metaphorical	He theorized that in the eye cancer, an infant inherited a damaged copy of a gene from one parent and a normal copy from the other.	to believe especially on uncertain or tentative grounds	form or construct theories
Metonymy	Ricardo Ulate, a Costa Rican delegate, said it 's not surprising that the major powers are fighting over who should bear the costs for curbing greenhouse gases, even as vulnerable countries have become more aggressive in seeking to hold the big emitters accountable for their actions.	a politically organized body of people under a single government	the people who live in a nation or country

Table 22: Examples of GPT-4o prediction errors.

Annotation	Context	Gold Definition	GPT-4o Generation
Equal	In the tower , five men and women pull rhythmically on ropes attached to the same five bells that first sounded here in 1614 .	an adult person who is male (as opposed to a woman)	man (noun): an adult male human being.
Better	I was just sitting down to meet with some new therapy clients, a couple , and the building started shaking like crazy.	a pair of people who live together	couple (noun): two people who are married, engaged, or otherwise romantically involved with each other.
Worse	Alimta is used together with cisplatin (another anticancer medicine) when the cancer is unresectable (cannot be removed by surgery alone) and malignant (has spread, or is likely to spread easily, to other parts of the body), in patients who have not received chemotherapy (medicines for cancer) before advanced or metastatic non-small cell lung cancer that is not affecting the squamous cells.	comparatively late in a course of development	advanced (adj.): Referring to a stage of disease, particularly cancer, that is more progressed, having spread significantly or becoming more serious in its development.
Wrong	If they put a Republican into office, not only will they acquire less in terms of local benefits but their selected legislator will be relatively powerless to prevent other legislators from "bringing home the bacon" to their respective constituencies.	financial assistance in time of need	benefit (noun): An advantage or profit gained from something

Table 23: Examples of human annotation on the Definition Generation task.

Annotation	Context	Gold Definition	GPT-4o Generation
Good	Ten shirt-sleeved ringers stand in a circle, one foot ahead of the other in a prize-fighter's stance, each pulling a rope that disappears through a small hole in the high ceiling of the ringing chamber.	a strong line	In this sentence, the word "rope" refers to the thick, strong cord or line that each of the ten ringers is pulling on. These ropes are part of the mechanism used to ring bells, typically in a church or bell tower. The ropes extend upwards and disappear through a hole in the ceiling, suggesting they are connected to the bells above. The ringers pull on these ropes to create the movement needed to sound the bells, with each person handling one rope to control one bell.
Bad	In this sense, a teacher can be compared to one's family doctor.	a person whose occupation is teaching	In the given sentence, the word "teacher" is being compared to a "family doctor", suggesting that a teacher, like a family doctor, plays a critical and foundational role in someone's life. Just as a family doctor provides care, guidance, and support for a person's health over time, a teacher provides knowledge, guidance, and support for a person's educational and personal development. This comparison highlights the importance and responsibility of a teacher in nurturing and shaping an individual's growth and well-being.
Borderline	Haney hoped that nobody noticed his sudden pallor , as he felt the blood drain from his cheeks .	notice or perceive	In this sentence, "noticed" means that Haney hoped no one saw or became aware of his sudden pallor, or the change in color of his face, which became pale as he felt the blood draining from his cheeks. He did not want anyone to observe or pay attention to this change, possibly because it might reveal his nervousness or fear.

Table 24: Examples of human annotation on the Free-form Explanation task.

Annotation	Context	Gold Definition	GPT-4o Generation
Yes	Nobody witnessed the fall – just the sickening impact when his body smashed on the pavement just outside the basement delivery entrance .	something that provides access (to get in or get out)	<ol style="list-style-type: none"> 1. The grand villa had an elaborate iron gate that served as the entrance to the sprawling estate. 2. She waited nervously at the entrance of the conference hall, gathering her thoughts before heading inside to give her presentation. 3. The amusement park’s entrance was adorned with colorful lights and banners, inviting families to a day of fun and excitement.
Yes but same pattern	One of the reasons American students do so poorly in math tests, as compared with British, French, German or Japanese students, is the influence of the “New Math” on American textbooks and teaching methods.	a justification for something existing or happening	<ol style="list-style-type: none"> 1. One of the reasons the restaurant consistently receives poor reviews is the inconsistent quality of its food and service. 2. One of the reasons the city experiences frequent traffic congestion is the outdated infrastructure and lack of public transportation options. 3. One of the reasons many people struggle to maintain a healthy lifestyle is the prevalence of fast food and sedentary work environments.
No	So you thought I didn’t mean what I said.	expect, believe, or suppose	<ol style="list-style-type: none"> 1. You assumed I was joking, didn’t you? 2. You believed I wasn’t serious about my statement? 3. Did you suppose I was speaking in jest?
Borderline	They may not be among the “best and brightest ” of their generation – there are very few such people, by definition.	all the people living at the same time or of approximately the same age	<ol style="list-style-type: none"> 1. The technological achievements of this generation will be studied by scholars for decades to come. 2. Artists from her generation faced unique challenges but also had unprecedented opportunities to reach global audiences. 3. Environmental concerns have become a defining feature of his generation, shaping policies and personal choices alike.

Table 25: Examples of human annotation on the Example Generation task.