

# On the Distinctive Co-occurrence Characteristics of Antonymy

**Zhihan Cao**

Institute of Science Tokyo  
School of Computing  
cao.z.c8a7@m.isct.ac.jp

**Hiroaki Yamada**

Institute of Science Tokyo  
School of Computing  
yamada@comp.isct.ac.jp

**Takenobu Tokunaga**

Institute of Science Tokyo  
School of Computing  
take@c.titech.ac.jp

## Abstract

Antonymy has long received particular attention in lexical semantics. Previous studies have shown that antonym pairs frequently co-occur in text, across genres and parts of speech, more often than would be expected by chance. However, whether this co-occurrence pattern is distinctive of antonymy remains unclear, due to a lack of comparison with other semantic relations. This work fills the gap by comparing antonymy with three other relations across parts of speech using robust co-occurrence metrics. We find that antonymy is distinctive in three respects: antonym pairs co-occur with high strength, in a preferred linear order, and within short spans. All results are [available online](#).

## 1 Introduction

Among various semantic relations, antonymy has long received particular attention. [Charles and Miller \(1989\)](#) proposed *the co-occurrence hypothesis*, suggesting that adjectival antonymy pairs tend to co-occur within a sentence more often than expected by chance. Subsequent studies have provided empirical support for this hypothesis ([Justeson and Katz, 1991](#)). This co-occurrence tendency has also been observed in parts of speech (PoS) other than adjectives ([Fellbaum, 1995](#); [Mohammad et al., 2008](#)) and across various genres ([Jones and Murphy, 2005](#); [Jones, 2006](#)).

Prior studies have identified several characteristics of the co-occurrence of antonymy pairs. One is that co-occurrence is likely to be ordered. Each antonymy pair often appears in a specific linear order ([Kostić, 2015, 2017](#); [Wu and Zhang, 2022](#)). The more frequent antonym tends to precede its less frequent counterpart ([Jones, 2002](#)). Another characteristic is that antonymy pairs likely co-occur in close proximity. This proximity typically falls within five-word windows or appears in coordinated constructions ([Jones et al., 2007](#); [Mohammad](#)

[et al., 2008](#)). Therefore, antonymy pairs are characterised by an intra-sentential co-occurrence pattern, which might be distinctive of antonymy.

However, this assumption remains largely untested, as the intra-sentential co-occurrence characteristics of non-antonymy relations have been underexplored. Some indirect evidence does exist. For example, hypernymy or holonymy pairs can be retrieved using lexico-syntactic patterns ([Hearst, 1992](#); [Girju et al., 2006](#); [Roller et al., 2018](#)). Yet these approaches rely on specific constructions rather than reflecting the general co-occurrence patterns within sentences. As a result, existing work is insufficient to determine whether the co-occurrence pattern uniquely characterises antonymy.

There is also a methodological concern. Prior studies commonly relied on raw co-occurrence counts or pointwise mutual information (PMI) of word pairs. Such metrics are biased toward extremely low or high frequencies that are common due to the sparsity of word frequencies ([Schulte Im Walde and Melinger, 2008](#)). The sparsity also makes many statistical tests (such as the Chi-square test used by [Kostić \(2017\)](#)) unreliable ([Dunning, 1993](#)), raising concerns about whether the co-occurrence observed is truly significant.

If antonymy indeed stands out in its co-occurrence, this distinctiveness needs to be established rather than assumed. Understanding different semantic relations in natural text is foundational not only to linguistic theory but also to computational modeling. Recently, pretrained language models (PLMs) have exhibited a puzzling phenomenon, *antonymy bias*. Regardless of architecture and size, PLMs perform substantially better on predicting the antonym of a given word over other relations ([Pitarch et al., 2023](#); [Cao et al., 2025a](#)). Establishing the distinctive co-occurrence characteristics of antonymy will lay the groundwork for analysing the semantic behavior of PLMs trained on a plain corpus.

**Research Objective** In summary, although prior studies have suggested that antonymy exhibits distinctive co-occurrence characteristics, existing evidence remains insufficient to validate this assumption. The present work evaluates and compares the co-occurrence of antonymy and other semantic relations across PoS, using robust and statistically grounded co-occurrence metrics. We aim to establish whether antonymy presents a unique pattern of co-occurrence characteristics.

## 2 Metrics

Semantic relations theoretically hold between word lemmas that share the PoS. We then define metrics at the lemma level.

### 2.1 Strength of Co-occurrence

The first metric,  $G^2$  score (Dunning, 1993), measures the co-occurrence strength. It directly addresses the methodological concerns of prior studies because it is theoretically and empirically robust to data sparsity (Moore, 2004; Evert, 2009) and supports a reliable statistical test.

The  $G^2$  score of a word pair  $(w, v)$  is essentially the logarithm of the likelihood ratio test statistic. The underlying null hypothesis is that  $(w, v)$  co-occurs in the same sentence at a chance level<sup>1</sup>.  $G^2$  scores asymptotically follow a  $\chi^2$  distribution with one degree of freedom. This allows us to interpret the score in two ways: 1) whether the co-occurrence is statistically significant, and 2) if it is, a higher score means stronger co-occurrence.

Formally, the  $G^2$  score is defined as

$$G^2(w, v) = 2 \sum_{c \in C} O_c \log \frac{O_c}{E_c}, \quad (1)$$

where  $C$  denotes a set of the possible events regarding the co-occurrence of  $(w, v)$  in a sentence:  $\{wv, w\bar{v}, \bar{w}v, \bar{w}\bar{v}\}$ . Here,  $wv$  indicates that both  $w$  and  $v$  appear in the same sentence;  $w\bar{v}$  and  $\bar{w}v$  denote that only  $w$  or only  $v$  appears, respectively;  $\bar{w}\bar{v}$  denotes that neither occurs.  $O_c$  and  $E_c$  represent the observed and expected number of sentences for each event  $c$ . The expected count  $E_c$  is calculated under the assumption that  $w$  and  $v$  occur independently. For example, when  $c = wv$ , the expected count is computed as  $E_c = \frac{|w| \times |v|}{N}$  where  $|w|$  and

<sup>1</sup>The original null hypothesis is that  $(w, v)$  co-occurs adjacently by chance. We relax it according to our research objective.

$|v|$  are the number of sentences containing  $w$  and  $v$ , respectively, and  $N$  is the total sentence number.

The advantages of  $G^2$  scores are three-fold. First,  $G^2$  scores normalise the co-occurrence of two words by their individual frequency and are hence robust to extremely high frequencies. Second, the use of logarithm mitigates the effect of extremely low frequencies. Finally,  $G^2$  scores incorporate more information than PMI and raw counts by considering both how often a pair co-occurs and how often it does not. This enables an evaluation of the co-occurrence strength more comprehensively without being biased only by the co-occurred side.

### 2.2 Order of Co-occurrence

Prior studies (Kostić, 2015, etc.) suggest the order might distinguish antonymy from other relations. Our second metric evaluates the linear order of a word pair  $(w, v)$ . We let  $w$  always be the more frequent word and  $v$  the less frequent one. For each sentence where  $(w, v)$  co-occurs, we assign an order score of +1 if  $w$  precedes  $v$ , and  $-1$  otherwise. Then, for every word pair  $(w, v)$ , we obtain a list of order scores. We determine whether there is a preferred order of  $(w, v)$  by conducting a binomial test on whether the proportion of +1 is significantly greater or less than 0.5. The order of a pair is defined as the average order score if the pair has a preferred order, and zero otherwise.

### 2.3 Distance of Co-occurrence

Whether the co-occurrence takes place in close proximity may be a distinctive characteristic of antonymy (Mohammad et al., 2008, etc.). The last metric measures the distance of co-occurrence, defined by the average number of words separating the two co-occurring words in a sentence.

**Statistical Tests** In order to establish whether the antonymy is statistically different from other relations, for each metric, we conduct a Brunner-Munzel test on every possible relation pair. It generalises the Mann-Whitney U test by relaxing the assumption of equal variances between samples being compared, showing better empirical robustness (Brunner and Munzel, 2000). For statistical tests described above (test of individual  $G^2$  scores, binomial test, and Brunner-Munzel test), the significance level is set at 0.01.

pos relation NOUN HYP-HPO 11664

POS	ANT	HOL	HYP	SYN	UNR
NOUN	97	758	11,664	1,078	7,618
VERB	74	–	6,853	1,214	869
ADJ	192	–	–	309	1,409
ADV	20	–	–	140	104
Total	383	758	18,517	2,741	10,000

Table 1: Counts of lemma pairs. “–” means the relation is not defined for the PoS in WordNet.

### 3 Data

We evaluate four semantic relations: antonymy (ANT), synonymy (SYN), hyper-hyponymy (HYP), and holo-meronymy (HOL). The hyper-hyponymy is defined to hold if either word in a pair is a hypernym of the other; similarly, holo-meronymy holds if either is a holonym of the other. This treatment aligns the definitions of all relations with the symmetric nature of antonymy and synonymy, making all relations comparable.

We retrieve lemma pairs of nouns, verbs, adjectives, and adverbs in any of the four relations from WordNet (Miller, 1995). A series of filtering steps is applied in order to ensure that all lemma pairs are of the same linguistic unit level and are lexico-semantically interpretable.

We exclude: 1) pairs containing multi-word expressions, abbreviations, or named entities; 2) pairs where either word has a frequency of zero or one in WordNet; 3) pairs where multiple semantic relations exist between the two lemmas; 4) verb pairs involving linking verbs, auxiliary verbs, or light verbs; 5) hypernymy pairs with a path length more than two in WordNet hierarchy<sup>2</sup>. These procedures result in 25,115 lemma pairs.

We estimate the metrics on the Corpus of Contemporary American English (COCA, Davies, 2008). It is designed to reflect the general usage patterns of the English language across balanced genres.

Choosing COCA brings an additional benefit: it is annotated with both PoS tags and lemmas. The PoS and lemma annotations allow us to efficiently check semantic relations between two co-occurring lemmas using WordNet.

To reduce noise, we exclude sentences shorter than five words, which are often exclamatory or fragmentary. This leaves 17,718,403 sentences for analysis. Among the 25,115 lemma pairs retrieved

<sup>2</sup>Because hypernymy pairs with a path length of two or less are most intuitive to human judgments (Vulić et al., 2017).

PoS	ANT	HOL	HYP	SYN	UNR
NOUN	<b>11,144</b> 99%	1,819 92%	406 75%	628 88%	19 36%
VERB	<b>915</b> 91%	– –	83 52%	164 72%	7 24%
ADJ	<b>2,309</b> 97%	– –	– –	120 76%	9 36%
ADV	<b>1,471</b> 95%	– –	– –	128 71%	14 33%

Table 2: Average  $G^2$  scores (above) and percentages of significant cases (below).

from WordNet, 22,399 lemma pairs are observed in COCA.

As a control, we additionally randomly sample 10,000 lemma pairs from all lemma pairs that co-occur intra-sententially but are unrelated (UNR) in any relation defined in WordNet. The filtering procedures described above are applied to unrelated pairs as well.

We sort the two lemmas in each pair by descending frequency. Table 1 shows statistics of lemma pairs for each relation. The counts are highly imbalanced across relations, reflecting the sparsity of WordNet (Cao et al., 2025b).

### 4 Results

Table 2 presents the average  $G^2$  scores and the percentage of pairs with significant scores. In this table and all subsequent tables, boldface indicates relations whose score differs significantly from all other relations within the same PoS for the corresponding metric. Antonymy pairs consistently yield both the highest  $G^2$  scores (ranging from 915 to 11,144) and the largest percentage (at least 91%) of significant co-occurring pairs across all PoS. Figure 1 shows the distribution of  $G^2$  scores per relation and PoS. The first and third quantiles of  $G^2$  scores for Antonymy are higher than those of all other relations across PoS. These results confirm that antonyms not only co-occur more frequently than expected by chance, but also with greater strength than other relations.

We acknowledge that antonymy tends to have extremely high  $G^2$  scores. For nominal antonymy, there are even two pairs, (“child”, “parent”) and (“man”, “woman”), whose  $G^2$  scores are greater than 100,000. Such extremely high outliers might have inflated the average  $G^2$  scores to 11,144 for nominal antonymy, making it the highest among all PoS and relations.




Figure 1: Distribution of  $G^2$  scores per PoS and relation with the y axis in logarithmic scale.

PoS	ANT	HOL	HYP	SYN	UNR
NOUN	<b>.13</b> 97%	.05 96%	.02 93%	.02 90%	.01 33%
VERB	<b>.11</b> 94%	–	.03 88%	.02 81%	.02 40%
ADJ	<b>.11</b> 91%	–	–	.04 61%	.00 21%
ADV	<b>.29</b> 95%	–	–	.05 87%	.03 29%

Table 3: Average order scores (above) and the percentages of pairs with a preferred order among all significantly co-occurring pairs (below).

Table 3 presents the results of the order preference. Antonymy pairs show a strong tendency to prefer a specific order. Across all PoS, more than 90% of the co-occurring antonymy pairs have a preferred order, typically with an average order score above 0.10. The average order scores for antonymy are overall significantly larger than other relations across PoS. This indicates that the more frequent antonym slightly tends to precede the less frequent counterpart. For other relations, a preferred order exists at the pair level but is not consistent across pairs, resulting in no clear pattern at the relation level.

Figure 2 shows the distribution of order scores per relation and PoS. For non-antonymy relations, the range of order scores tends to be around 0, confirming that there is no clear co-occurrence order



Figure 2: Distribution of order scores per PoS and relation.

PoS	ANT	HOL	HYP	SYN	UNR
NOUN	<b>15</b>	22	26	25	30
VERB	<b>18</b>	–	28	25	50
ADJ	<b>14</b>	–	–	17	25
ADV	<b>11</b>	–	–	20	27

Table 4: Average distances between significantly co-occurring pairs.

preference<sup>3</sup>.

Table 4 shows the co-occurrence distance. Antonyms co-occur more closely than other relations. On average, antonyms co-occur within 18 words, which is shorter than the span of pairs in other relations across PoS. Figure 3 shows the distribution of distances per relation and PoS. For antonymy, the third quantile and the maximum of co-occurrence distance are consistently the lowest among all relations across PoS, confirming that antonymous pairs tend to occur close in text.

In short, compared to other relations, antonymy pairs co-occur 1) with greater strength, 2) in the order where the more frequent word tends to precede, and 3) at a shorter distance. Our findings align with previous observations on antonyms, and extend them to other relations and PoS, revealing the distinctive co-occurrence characteristics of antonymy.

<sup>3</sup>Note that this finding is based on our definition of relations, in which holonymy and meronymy are not distinguished, nor are hypernymy and hyponymy. If these distinctions are made, holonyms tend to precede meronyms to a moderate degree. Detailed results are provided in Appendix A.




Figure 3: Distribution of distances per PoS and relation with the y axis in logarithmic scale.

## 5 Discussion

**Cross-PoS Analysis** The strong co-occurrence of antonymy pairs is often attributed to semantic contrast. Antonymy pairs denote two opposite ends of a scale or a dichotomy in a situation (Cruse, 1986), which might help highlight semantic contrast and hence lead to their significant co-occurrence (Jones, 2006). If the significant co-occurrence is indeed associated with semantic contrast, it should persist after derivation, which usually changes the PoS but not the core meaning. To test this, we examine whether co-occurrence significance is retained after derivation. Among the total of 192 adjectival antonymy pairs, we identify 18 pairs from which 19 non-adjectival antonymy pairs are derived (e.g. (“strong”, “weak”) to (“strongly”, “weakly”))<sup>4</sup>. The 18 original adjectival antonymy pairs co-occur significantly; 18 out of the 19 derived non-adjectival antonymy pairs also exhibit significant co-occurrence. For other PoS and relations, we see a similar pattern. Among all non-adjectival or non-antonymy pairs that significantly co-occur, 85% of their derived pairs sustain the significant co-occurrence. These findings suggest that significant co-occurrence is robust under derivational change, aligning with the idea that co-occurrence might stem from the semantic contrast.

<sup>4</sup>This is done using the information in WordNet. Details can be found in Appendix B.

**Cross-relation Analysis** Antonymy is considered to differ from hypernymy, holonymy, and synonymy in nature because it is not just semantic but also lexically constrained (Justeson and Katz, 1991). For example, “hot” is in antonymy with “cold” but not with “chill”. We verify to what extent each relation is lexically constrained. For each relation, we calculate how many lemmas the more frequent lemma is associated with on average<sup>5</sup>. For antonymy, the frequent lemma is associated with only one lemma on average, whereas for other relations, it ranges from 1.7 to 6.7. This result suggests that antonymy involves the strongest lexical pairing, which might enhance their suitability for parallel constructions and, in turn, contribute to their frequent co-occurrence in texts.

**Relating Language Models** Pretrained language models perform substantially better on antonymy across model architectures and sizes (Pitarch et al., 2023; Cao et al., 2025a). The distinctive co-occurrence characteristics of antonymy may provide signals that facilitate learning antonym pairs during pretraining. However, as prior studies commonly use sentence completion tasks in evaluation, models might take advantage of those intra-sentential co-occurrence characteristics, using them as a shortcut (Du et al., 2023) rather than relying on knowledge of antonymy. Given this, our findings highlight the need to disentangle the extent to which PLMs rely on such distributional clues from the extent to which they generalise beyond them.

## 6 Conclusion

This work presents an initial exploration of the intra-sentential co-occurrence of antonymy in comparison to other semantic relations across PoS. We find that antonymy is consistently distinctive from all other relations across all metrics used. Our findings establish a robust empirical foundation for the distributional nature of antonymy, offering a solid empirical basis for its future analysis.

## Limitations

We only characterise the co-occurrence of semantically related words in a quantitative manner and do not address the qualitative evaluation. We hence can not answer in which lexical and dependency

<sup>5</sup>The Appendix C presents the details.

constructions two words are likely to occur. Furthermore, as an exploratory study, we did not apply word sense disambiguation to a sentence but only used the lemma and PoS of each word as a shallow filtering. We chose not to apply word sense disambiguation in alignment with the previous studies which are done on lemma level. We only studied COCA, and hence, future work should focus on a specific pretraining corpus and examine how PLMs trained on it become biased toward antonymy.

## Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP25KJ1271. We would like to extend our sincere gratitude to Prof. Simone Teufel for valuable discussions. We would also like to express our appreciation to Wanjing Cao, Huaqing Qi, and Shuqin Zhang for their support throughout this work.

## References

- Edgar Brunner and Ullrich Munzel. 2000. [The nonparametric behrens–fisher problem: Asymptotic theory and a small-sample approximation](#). *Biometrical Journal*, 42(1):17–25.
- Zhihan Cao, Hiroaki Yamada, Simone Teufel, and Takenobu Tokunaga. 2025a. [A comprehensive evaluation of semantic relation knowledge of pretrained language models and humans](#). *Language Resources and Evaluation*.
- Zhihan Cao, Hiroaki Yamada, Simone Teufel, and Takenobu Tokunaga. 2025b. [Misalignment of semantic relation knowledge between WordNet and human intuition](#). In *Global WordNet Conference 2025*.
- Walter G. Charles and George A. Miller. 1989. [Contexts of antonymous adjectives](#). *Applied Psycholinguistics*, 10:357–375.
- D. Alan Cruse. 1986. *Lexical Semantics*. Cambridge University Press, New York.
- Mark Davies. 2008. [The corpus of contemporary american english \(coca\)](#). Data available online at <https://www.english-corpora.org/coca/> (Accessed at 04 May 2025).
- Mengnan Du, Fengxiang He, Na Zou, Dacheng Tao, and Xia Hu. 2023. [Shortcut learning of large language models in natural language understanding](#). *Commun. ACM*, 67(1):110–120.
- Ted Dunning. 1993. [Accurate methods for the statistics of surprise and coincidence](#). *Computational Linguistics*, 19(1):61–74.
- Stefan Evert. 2009. [58. Corpora and collocations](#), volume 2, pages 1212–1248. De Gruyter Mouton, Berlin, New York.
- Christiane Fellbaum. 1995. [Co-occurrence and antonymy](#). *International Journal of Lexicography*, 8:281–303.
- Roxana Girju, Dan Moldovan, and Adriana Badulescu. 2006. [Automatic discovery of part-whole relations](#). *Computational Linguistics*, 32:83–135.
- Marti A Hearst. 1992. [Automatic acquisition of hyponyms from large text corpora](#). In *COLING 1992 Volume 2: The 14th International Conference on Computational Linguistics*, pages 539–545.
- Steven Jones. 2006. [A lexico-syntactic analysis of antonym co-occurrence in spoken English](#). *Text & Talk*, 26(2):191–216.
- Steven Jones and M. Lynne Murphy. 2005. [Using corpora to investigate antonym acquisition](#). *International Journal of Corpus Linguistics*, 10(3):401–422.
- Steven Jones, Carita Paradis, M. Lynne Murphy, and Caroline Willners. 2007. [Googling for ‘opposites’: a web-based study of antonym canonicity](#). *Corpora*, 2:129–155.
- Steven Jeffrey Jones. 2002. *Antonymy: A Corpus-Based Perspective*. Routledge.
- John S. Justeson and Slava M. Katz. 1991. [Co-occurrences of antonymous adjectives and their contexts](#). *Computational Linguistics*, 17:1–20.
- Nataša Kostić. 2015. [Antonym sequence in written discourse: a corpus-based study](#). *Language Sciences*, 47:18–31.
- Nataša Kostić. 2017. [The distributional asymmetries of English antonyms in language use](#). *Brno studies in English*, 1:[5]–32.
- George A. Miller. 1995. [Wordnet](#). *Communications of the ACM*, 38:39–41.
- Saif Mohammad, Bonnie Dorr, and Graeme Hirst. 2008. [Computing word-pair antonymy](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 982–991, Honolulu, Hawaii. Association for Computational Linguistics.
- Robert C. Moore. 2004. [On log-likelihood-ratios and the significance of rare events](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 333–340, Barcelona, Spain. Association for Computational Linguistics.
- Lucia Pitarch, Jordi Bernad, Lacramioara Dranca, Carlos Bobed Lisbona, and Jorge Gracia. 2023. [No clues good clues: out of context lexical relation classification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5607–5625, Toronto, Canada. Association for Computational Linguistics.

Stephen Roller, Douwe Kiela, and Maximilian Nickel. 2018. [Hearst patterns revisited: Automatic hypernym detection from large text corpora](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 358–363. Association for Computational Linguistics.

Sabine Schulte Im Walde and Alissa Melinger. 2008. An in-depth look into the co-occurrence distribution of semantic associates. *Italian Journal of Linguistics*, 20:89–128.

Ivan Vulić, Daniela Gerz, Douwe Kiela, Felix Hill, and Anna Korhonen. 2017. [Hyperlex: A large-scale evaluation of graded lexical entailment](#). *Computational Linguistics*, 43:781–835.

Shuqiong Wu and Jie Zhang. 2022. [Antonym order in english and chinese coordinate structures](#). *Review of Cognitive Linguistics*, 20(2):530–557.

## Appendix

### A Ordered Scores for Asymmetric Relations

Figure 4 presents the order scores of asymmetric relations. We redefine the order score for asymmetric relations as follows. A score approaching 1 indicates that a holonym or hypernym precedes its meronym or hyponym, respectively, while a score close to -1 indicates the reverse. For meronymy, more than 50% of the order scores are positive. The results show that holonyms generally precede their meronyms. In contrast, for hypernymy and hyponymy, the order score distribution centers around 0, indicating that the ordering tendency is less clear in the hypernymy/hyponymy relations.




Figure 4: Distribution of order scores of asymmetry relations per PoS.

PoS		Rel.		Count
Orig.	Derv.	Orig.	Derv.	
ADJ	ADV	ANT	ANT	13 ( 12)
ADJ	ADV	SYN	SYN	8 ( 7)
ADJ	NOUN	ANT	ANT	5 ( 5)
ADJ	NOUN	SYN	SYN	1 ( 1)
ADJ	NOUN	SYN	HYP	8 ( 7)
ADJ	VERB	ANT	ANT	1 ( 1)
ADJ	VERB	SYN	SYN	1 ( 1)
NOUN	VERB	ANT	ANT	2 ( 2)
NOUN	VERB	SYN	SYN	15 ( 12)
NOUN	VERB	HYP	HYP	111 ( 90)
NOUN	VERB	SYN	HYP	46 ( 44)
TOTAL				211 (182)

Table 5: Counts of original (Orig.) and derived (Derv.) POS pairs and their corresponding relations. The number of pairs that sustain significance after derivation are in parenthesis.

POS	ANT	HOL	HYP	SYN	UNR
NOUN	1.0	1.9	5.8	1.5	2.3
VERB	1.1	–	6.9	1.9	1.5
ADJ	1.0	–	–	1.4	1.6
ADV	1.0	–	–	1.5	1.2
Micro AVG	1.0	1.9	6.7	1.7	2.2

Table 6: The average number of lemmas that are associated with the more frequent lemmas per relation and PoS.

### B Derived Pairs Statistics

In WordNet, some lemmas are linked to their derived lemmas. For each lemma pair  $(w, v)$  that is semantically related and significantly co-occurs, we retrieve all derived lemmas  $w_d$  of  $w$  and  $v_d$  of  $v$  using the information in WordNet. For every  $(w_d, v_d)$ , we check whether it forms a semantic relation and appears in our data. If so, we then examine whether it sustains the co-occurrence significance.

Table 5 presents the results. The semantic usually persists after derivation, particularly for antonymy and hyper-hyponymy. Synonymy pairs sometimes become hyper-hyponymy, reflecting the established similarity between synonymy and hypernymy (Vulić et al., 2017; Cao et al., 2025b).

### C Associated Lemma Counts

For each relation, we count how many lemmas are associated with the more frequent lemma on average. In other words, for all pairs  $(w, v)$ , we retrieve pairs that share  $w$  and compute the average number of associated  $vs$  per  $w$ . Table 6 presents the results.

Antonymy pairs usually have a unique associated antonym for each more frequent lemma  $w$ , as the average number of associated lemmas ranges only from 1.0 to 1.1. In holo-meronymy and synonymy, the frequent lemma is associated with slightly more lemmas than in antonymy, with the averages ranging from 1.4 to 2.0. In hyper-hyponymy, the frequent lemma is typically associated with around 6 lemmas. Hence, antonymy is the most lexically constrained among the other relations studied here.