

Morphological vs. Lexical Antonyms in Italian: A Computational Study on Lexical Competition

Martina Saccomando^{1,*}, Andrea Zaninello² and Francesca Masini¹

¹Alma Mater Studiorum - University of Bologna, Bologna (Italia)

²Fondazione Bruno Kessler, Trento; Free University of Bozen-Bolzano (Italy)

Abstract

In this paper, we examine the competition between pairs of adjectives in Italian that are antonyms of the same term: one is a “morphological antonym” formed by negative prefixation, the other is a “lexical antonym” with no morphological relationship with the term in question. We consider pairs of adjectives that are reported as antonyms in lexicographic resources and extract the nouns that can be modified by both adjectives from a large corpus. We select a set of 8 nouns for each pair that present higher, lower, and comparable frequencies combined with each antonym respectively and then we perform two experiments with a LLM. Firstly, we perform experiments for masked-token prediction of the adjective, to study the correlation between prediction accuracy and the frequency of the noun-antonym pair. Secondly, we perform a polarity-flip experiment with a multilingual LLM, asking to change the adjective into its positive counterpart. Our results point to the conclusion that the lexical antonym seems to have a narrower lexical coverage and scope than the morphological antonym.

Keywords

Antonymy, Morphological antonyms, Lexical antonyms, Competition, Corpus analysis, Large language model, Token prediction, Polarity flip, Italian

1. Introduction

Antonymy is the semantic relationship between terms with opposite meanings. In their canonical form, two antonyms’ meanings can be represented as the poles of a semantic continuum [1] where one term has a “positive” semantic value, the other a “negative” one [2].

In Italian, given a word (e.g., the adjective *felice* ‘happy’), antonyms can either be realized via prefixation of that word (e.g., *infelice* ‘unhappy’) or through an independent lexeme (e.g., *triste* ‘sad’). In our work, we refer to these types of antonyms as *morphological antonym* and *lexical antonym*, respectively. A word in the lexicon may have both a morphological and a lexical antonym, only one of them, or neither. In this paper, we are interested in triplets of adjectives where a positive adjective (e.g., *felice*) presents two possible antonyms (or “co-antonyms”), one formed morphologically by prefixation (e.g., *infelice*), one morphologically independent (e.g., *triste*) (Figure 1).

In this paper, we are interested in studying the factors that govern the selection of the morphological antonym vs. the lexical one. These two types of antonyms express “negative” semantics with respect to the opposite, “positive” term in different ways: *implicitly* in the case of




Figure 1: Two possible antonyms, one morphological, one lexical, for the same word.

lexical antonyms; *explicitly* in the case of morphological antonyms, by adding a prefix with a negative, contradicting value. Considering their different morphological structure, one possible hypothesis on their lexical competition is that the morphological antonym has a more *restricted* semantics, representing the negation of the semantics of its adjectival base, while the lexical antonym has a broader semantic coverage, as it is morphologically independent from its positive counterpart (see Section 3).

To the best of our knowledge, despite the wealth of literature on antonyms (Section 2), there is no empirical in-depth study that investigates the competition between morphological and lexical antonyms in single languages, including Italian. Studies on antonyms do identify the two types of antonyms but generally do not address the factors influencing the preference for one type over the other intralinguistically.

This study investigates the competition between these two types of antonyms by firstly studying their distribution in corpora (Section 5.1); secondly, by testing the ability of a native-Italian language model to predict them in a masked-token prediction task (Section 5.2); and, fi-

CLiC-it 2024: Tenth Italian Conference on Computational Linguistics, Dec 04 – 06, 2024, Pisa, Italy

*Corresponding author.

✉ martina.saccomando@studio.unibo.it (M. Saccomando);
azaninello@fbk.eu (A. Zaninello); francesca.masini@unibo.it
(F. Masini)

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



nally, by performing a substitutability task within the same context by switching the polarity of the context sentence with a SOTA multilingual LLM, in order to study when the adjective is switched to the positive *un*-prefixed adjective or to another, positive but morphologically unrelated one (Section 5.3).

2. Background

Whereas the exploration of the competition between morphological and lexical antonyms, addressed in this paper, has not gained much attention so far, the literature on antonymy in general is abundant, especially in relation to the English language.

A term's **antonym** is related to it according to three main **characteristics**: *polarity*, *gradability* and *canonicity*. The first two characteristics refer to the positioning of the two antonymic terms w.r.t. the two poles (**polarity**) of a graded (**gradability**) scale [3], along which free positions can be occupied by other similar but differently graded terms. The scale is based on a specific property that the two terms share. For example, in the pair *long-short*, the two antonyms share the property of "length", defining the start and end of an axis whose poles are defined by two terms, with *long* representing the "unmarked" base term of the opposition (this is why we ask for "how long" rather than "how short" something would last [4]).

However, there are cases where antonymic pairs are formed with potentially competing antonyms, like *friendly-unfriendly* vs. *friendly-hostile*: *friendly-unfriendly* is placed on a scale that defines a greater or lesser degree (gradability) of a property, while *friendly-hostile* are certainly in opposition but belong to two scales of incompatible properties. In terms of their gradability, therefore, it seems that the morphological antonym is "more gradable" than the lexical one.

Canonicity, according to Paradis and Willners [5], defines two semantically related and conventionalized terms as a pair in the language. It is a gradual property and can be possessed to a greater or lesser extent. A high degree of canonicity translates into a high degree of semantic-lexical embedding in memory and leads to conventionalization in usage.

Psycholinguistic studies also suggest that canonical antonyms derive from the speakers' experience with the language: the two terms are inseparable, one elicits the other [6, 7, 8, 9, 10]. When a term has two structurally different but semantically similar antonyms (Figure 1), canonicity is influenced by factors such as learned preference for specific pairings, the speaker's familiarity due to exposure, and different nuances of meaning [11]. Out of context, the antonyms may appear equivalent, but within context, a specific meaning may be activated that

only one of the antonyms possesses, preventing their full synonymy and interchangeability.

Justeson and Katz [12] take a different approach. Using the Brown Corpus and Deese's antonym dataset, they were among the first to study antonymy based on a corpus. They found that antonymic terms co-occur more frequently than expected, confirming a syntagmatic relationship between them (in addition to the paradigmatic one). This syntagmatic relationship was confirmed by a more extensive work carried out in subsequent years: Jones [11] collected 56 antonym pairs, analyzing journalistic texts, to identify eight discursive functions of antonym co-occurrence.

A study that does address the competition between forms is Aina et al. [13], who studied syntactically negated adjectives and morphological antonyms (e.g., *not happy* vs *unhappy*, respectively). Using distributional semantics they found that a syntactically negated adjective is more similar to the positive adjective than to its lexical antonym. Additionally, they show that the morphological antonym is less similar to its lexical base compared to the similarity between a negated adjective and its non-negated counterpart.

Last but not least, a very recent typological study [14] examines 37 antonym pairs across 55 languages, focusing on antonym formation. When a derived form is attested, it typically applies to the member of the pair with lower valence or lesser magnitude. Antonyms related to core property concepts (dimension, age, value, color) tend to be expressed through distinct lexical forms (resulting in lexical antonyms), while those related to peripheral property concepts (physical property, human property, speed) are generally encoded using derived forms (namely morphological antonyms). Although the study is insightful and inspiring in a number of ways, the specific question of which reasons underlie the preference for lexical or morphological forms in a single language remains unanswered.

3. Morphological vs. Lexical Antonyms

A **morphological antonym** (e.g., *inattivo* 'inactive' from *attivo* 'active') is immediately recognizable as a negative term due to the presence of the negative prefix *un-* 'in-/un-', while the nature of the opposition is less immediate with a **lexical antonym** (e.g., *passivo* 'passive'), because one has to identify the property shared by the two opposing terms.

Moreover, these two ways of forming an antonym from the same term create an asymmetric system: while one of the two terms in lexical opposition (e.g., *attivo* 'active') has its morphological antonym (e.g., *inattivo* 'in-

active’), the other term (e.g., *passivo* ‘passive’) does not (e.g., **apassivo*, **impassivo*, ‘*unpassive’). This imbalance is due to the greater emphasis that language places on everything that requires more precise specifications [15]. The situation is not always perfectly equivalent interlinguistically: for example, a form like **invero* to indicate the opposite of *vero* ‘true’ is not attested in Italian, while it is possible in English (*untrue*).

However, there are cases in both Italian and English where the two competitors have different nuances: for example, *infelice-triste* / *unhappy-sad* cover different contexts in that *triste* and *sad* convey a stronger emotional meaning, while *infelice* and *unhappy* encompass certainly strong but less intense feelings.[16]

Regarding **distribution in usage**, lexical antonyms are predominant for more basic meanings, supporting the Principle of Least Effort theorized by Zipf [17], which suggests that we expect the most used concepts to be coded with short and simple words: basic terms therefore tend to have structurally simple antonyms even when more complex morphological antonyms would be possible. For example, for a pair like *alto-basso* ‘tall-short’, there is no morphological counterpart that can be associated with either term: neither **inalto* ‘*untall’ nor **inbasso* ‘*unshort’ exist. These are canonical antonyms referring to basic language concepts: in cases like this, but not exclusively, preference for simpler and more immediate words blocks the potential formation of a morphological antonym.

According to Murphy [4], culturally salient concepts necessitate clear and concise linguistic expressions. For this reason, lexical antonyms (e.g., *passivo* ‘passive’) are the most frequent choice because they require less cognitive effort to understand. Although it is possible to create new opposite terms through derivational and morphological processes with speaker’s creativity, this option is less commonly employed in this context, as it is perceived as a deviation from the linguistic norm.

The **competition** between the two terms of the antonym pair, i.e., the situation in which the usage context of both terms is nearly the same and allowing for a certain degree of substitutability, is still debated.

According to one hypothesis, since the morphological antonym is the “perfect” negation of a specific lexical base, it should occur in more restricted contexts (i.e., a subset of the contexts of its positive counterpart) and should therefore have a narrower semantics (cf. [18, 19, 15]). So, **morphological antonyms would be less polysemous**. On the other hand, the lexical antonym, not sharing identical lexical properties with the opposed term, should occur in broader contexts and thus be more polysemous.

However, Murphy [4], examining the English triple *friendly-unfriendly-hostile*, notes that “The two antonyms are hardly equivalent, though, since *unfriendly* describes

a wider range of ways of not being friendly (such as being aloof) whereas *hostile* is fairly specific” (Murphy 2003: 202). So, **the morphological antonym would be more polysemous**, while the lexical counterpart would have a narrower scope.

Given these two competing hypotheses, we aim to empirically verify:

- whether the lexical antonym is more frequent than the morphological antonym;
- whether the morphological antonym is actually less polysemous than the lexical antonym.

In order to do this, we design a set of experiments. We first select antonym pairs; we calculate their frequency of co-occurrence with nouns to have a defined context; then, we perform two tasks: (i) masked-token prediction and (ii) polarity flip.

4. Dataset Construction

4.1. Antonym Pair Selection

For our study, we decided to focus on adjectives, as this class is the most suitable for investigating antonymy, given that it includes content words that normally express qualities. Moreover, adjectives are semantically simpler compared to other word classes as they usually describe a single property that may be or may not be present to a greater or lesser degree (Jones et al. [1]).

Starting from 1535 adjectives of the fundamental Italian lexicon extracted from the Italian dictionary *Zingarelli 2024*, we filtered 303 items marked as ‘contr.’.

We then selected, for this pilot study, 5 adjectives with the following properties: they needed to be adjectives with both a morphological and a lexical antonym; they needed to be maximally interchangeable in different contexts.¹

Finally, we created our initial dataset by pairing each adjective with its corresponding morphological antonym and a possible lexical antonym.

The morphological antonym was formed by using one of the three possible prefixes productively used in Italian to create antonyms, namely: *dis-*, *s-*, *in-* [20]².

The lexical antonym was chosen randomly among all possible options, taking into account synonymy with the morphological antonym and semantic neutrality. This means that the lexical antonyms were selected to be ideally substitutable with the morphological ones in as many contexts as possible, and roughly possessing the same

¹The admittedly limited size of the dataset is due to the exploratory nature of our study.

²The prefix *in-* is the most productive and the most widespread. It often adapts phonetically to the bases it attaches to, forming the allomorphs *im-*, *ir-*, and *il-* via assimilation [20].

number of senses according to the dictionary (cf. Table 2). The synonymy between the two types of antonyms was further confirmed using *Il grande dizionario dei Sinonimi e dei Contrari, Zingarelli 2013* [21].

Summing up, the antonym pairs examined are:

- *infelice* ‘unhappy’ - *triste* ‘sad’
- *impreciso* ‘imprecise’ - *approssimativo* ‘approximative’
- *scorretto* ‘incorrect’ - *sbagliato* ‘wrong’
- *imprudente* ‘imprudent’ - *avventato* ‘reckless’
- *insufficiente* ‘insufficient’ - *scarso* ‘scant’

resulting in 5 triplets (base, morph_ant, lex_ant).

4.2. Corpus-based Analysis

We analyzed the occurrences of the selected adjectives with nouns in the *itTenTen20* corpus, a large web corpus of written Italian, searched through the SketchEngine platform <https://www.sketchengine.eu/>.

The analysis of the occurrences highlighted that the two types of antonyms display partially different collocational preferences (see Appendix A, Table 3).

Overall, we can split the antonymic adjective-noun couples in two groups according to whether the co-occurrence is:

- (i) **polarization towards one of the two adjectives**: in these cases, we can speak of fairly stable distributional preferences, falling within the realm of collocations or idiomaticity (e.g., *matrimonio* preferably selecting *infelice* rather than *triste*);
- (ii) **similar with the two adjectives**, indicating potential substitutability of the two antonyms in the same contexts (e.g., *donna* selecting both *infelice* and *triste* with similar relative frequencies).

Both groups are relevant to explore the context of use of the two types of antonyms, although, for our current purposes, we specifically targeted the nouns in group (ii), namely nouns that occur with both adjectives, suggesting a certain degree of competition between the two antonyms: see sentence 1, where *infelice* ‘unhappy’ can be replaced by the lexical antonym *triste* ‘sad’.

1. *Un ritratto preciso ma discontinuo che ci restituisce l'immagine di una donna infelice, umiliata, affranta, ma non distrutta, non arresa alla sorte*³

³A precise but fragmented portrait that gives us the image of an unhappy, humiliated, distraught woman, but not destroyed, not resigned to fate.

4.3. Lexical Context Definition and Example-Sentence Extraction

Subsequently, for each antonym pair, eight nouns with different co-occurrence frequencies were selected. Specifically, we considered both nouns that typically occur with one of the antonyms (e.g., *matrimonio infelice*), falling within the domain of collocations, and (more generic) nouns whose co-occurrence frequencies with the lexical and the morphological antonym are very similar (e.g., *donna infelice* and *donna triste*) (cf. Table 1).

Noun + Adjective	Frequency
<i>matrimonio infelice</i> ‘unhappy marriage’	886
<i>matrimonio triste</i> ‘sad marriage’	24
<i>donna infelice</i> ‘unhappy woman’	325
<i>donna triste</i> ‘sad marriage’	316

Table 1

Differences between high and low frequency name+adjective co-occurrences

The latter case is especially interesting for our current purposes, as it represents possible ground for “competition”, namely a situation where the context of use is nearly the same and allows for a certain substitutability between the two terms of the antonym pair.⁴

After defining the noun list, for each noun we randomly selected 10 sentences containing the noun followed by the morphological antonym and ten sentences containing the same noun followed by the corresponding lexical antonym from *itTenTen20*. This was done for all eight nouns and for all five antonym pairs, resulting in a 800-sentence dataset.

5. Experiments

In Section 3 we outlined two possible hypotheses regarding the competition between the two terms of the antonym pair and we selected the following as a working hypothesis: the morphological antonym, being formed by a negative prefix applied to a specific lexical base, would have more restricted usage contexts (possibly a subset of the contexts of the lexical base), and therefore be less polysemous than the lexical one; on the other hand, the lexical antonym, not sharing morphological structure with the opposing term, would semantically cover some or all of its meanings along with other independent meanings, resulting in broader usage contexts and greater polysemy.

To verify this hypothesis, we performed 2 sets of experiments: (i) **masked-token prediction**, to estimate the

⁴For a detailed view of the selected adjectives and nouns, as well as their co-occurrence frequencies, see Appendix A.

probability of occurrence of one antonym or the other according to a native Italian language model; (ii) **polarity flip**, to transform the collected sentences from a negative meaning to a positive meaning.

5.1. Word Senses and Lexical Variety

Our analysis started with the identification of adjectives and their possible antonyms, which, as mentioned (Section 4.1), have been chosen on the basis of their possible substitutability within the same context. For this reason, first of all, the various dictionary definitions of antonyms were taken into consideration. We counted how many senses are associated to each antonym in the *Zingarelli* dictionary [22], taking the number of senses reported as a first proxy of semantic broadness.

As a second proxy, the semantic coverage of each antonym was taken into account. We conducted an analysis of the lexical variety of each group’s context in the selected sentences, by calculating the token/type ratio for each group. We report the results in Table 4. As can be seen, no relevant differences were found according to these features.

5.2. Masked-token Prediction

According to our hypothesis, in this task we expect that the predicted antonym will have a higher prediction accuracy in the sentences with the highest occurrence of the adjective with the selected words (represented by high relative frequency). In contexts with similar frequencies, we expect that accuracy should be similar for both antonyms, showing a genuine competition between the two, as the language model should not have a specific preference.

We previously said that the words that form the antonym pairs can be considered synonymous. In fact, full synonyms are rather rare (Murphy [4], among others), also because languages tend to avoid synonymy by differentiation in terms of meaning or distribution [23]. Therefore, the two terms of the pair are better regarded as near-synonyms, meaning that one term can cover almost all the meanings of the other but not all of them.

To evaluate the factors that lead to the choice of one antonym over the other, we decided to observe how a native Italian language model pre-trained for masked-token prediction model behaves in terms of the **probability of occurrence** of an antonym in a given context. In this respect, see Niwa et al.[24], who used BERT to predict antonyms in specific contexts: experiments on Japanese slogans showed a top-1 accuracy of 29.3% and a top-10 accuracy of 53.8%, with human evaluations confirming that over 85% of predicted antonyms were appropriate, demonstrating the method’s effectiveness in capturing contextually relevant antonyms.

We used bert-base-italian-xxl-cased language model⁵ to perform a token prediction task by masking the antonym present in each sentence. The model was asked to predict the probability of occurrence between the two possible antonyms; then, we took the alternative with the highest probability according to the model as the model prediction.

5.3. Polarity-flip

In this task, we asked a SOTA LLM, GPT-4o, to transform the sentences extracted from the dataset, both those containing the morphological antonym and those containing the lexical antonym, into positive sentences.

We used the same prompt for all antonym pairs, parametrising the antonyms and sentences presented, by asking the model to flip the sentence from a negative sense to a positive one, always by changing the adjective accompanying the target noun.⁶

We then fetched the new adjective generated by the model, and calculated when the new, positive adjective coincided with the lexical base, and when not.

The rationale behind this is that the senses of an antonym can be separated through the various positive terms with which it can be changed. We expected that sentences containing the morphological antonym would be turned into positive using their lexical base more often than their lexical counterparts, indicating a narrower semantics.

6. Analysis and Results

As regards the number of meanings listed in the dictionary for the two terms of the antonym pair (Table 2), these are almost the same, indicating that the recognized senses of each antonym alone are not decisive to determine the selection between one or the other. As for lexical variability (see Appendix B, Table 4), token/type ratio also fails to reveal a significant difference between morphological and lexical antonyms; in only three cases does the token/type ratio of lexical antonyms exceed that of morphological antonyms. This seems to indicate that factors other than contextual variability underlie the preference for one or the other.

Let us now consider Table 3, Appendix A. In addition to the co-occurrences of nouns with the two antonyms, the accuracy for the two tasks for each noun-adjective pair is also provided (further divided into noun-morphological antonym and noun-lexical antonym). The distribution of nouns in this table follows a specific order: nouns at the extremes exhibit occurrence frequencies polarized

⁵dbmdz/bert-base-italian-xxl-uncased

⁶For detailed information on the prompt used in this study, please refer to the Appendix B.

	senses		rel.freq		accuracy mask token prediction	accuracy mask token prediction	polarity flip	polarity flip
Couple	Exp.morf	Exp.lex	Exp.morf	Exp.lex	Exp.morf	Exp.lex	Exp.morf	Exp.lex
infelice-triste	4	4	0.48	0.52	0.525	0.687	0.125	0.137
impreciso-approssimativo	2	2	0.57	0.43	0.256	0.833	0.081	0.222
scorretto-sbagliato	3	3	0.48	0.52	0.025	0.986	0.076	0.144
imprudente-avventato	2	1	0.43	0.57	0.437	0.637	0.012	0.075
insufficiente-scarso	3	3	0.40	0.60	0.743	0.662	0.012	0.037

Table 2

Final experiment results. Number of senses, average of relative frequency and average of the accuracy of the two tasks: mask token prediction and polarity flip.

towards either the lexical antonym or the morphological antonym. Some nouns form with the adjective a fairly stable collocation, while other nouns form freer expressions. For the purposes of this study, particularly in relation to the analysis of competition, the central nouns with similar frequencies are of greater importance.

Upon examining the occurrences and accuracy, we observe that the values are comparable.

As for masked-token prediction, the consistent higher values of the lexical antonym indicate higher predictability and/or higher degree of idiomaticity, which contradicts our working hypothesis that lexical antonyms display a broader semantic coverage.

Finally, as for the polarity flip, consistently with the masked-token experiments, the sentences containing the morphological antonym were turned into positive using the lexical base fewer times than their lexical counterparts, suggesting that the latter may have a more restricted semantic spectrum, contrary to our initial hypothesis.⁷

7. Conclusion and Future Work

Our study investigated the differences and competition between two types of antonyms, morphological and lexical, focusing on a computational account of their context of use. While a lexical analysis did not prove decisive, experiments on masked-token prediction and polarity flip, aimed at approximating their semantic coverage, indicate that, unlike what is suggested in some studies on antonymy, the lexical antonym seems to have a narrower lexical coverage and scope, supporting the view that it is actually the morphological antonym, despite its closer relationship with the lexical base, that displays a wider

⁷An anonymous reviewer observes that this result is even more remarkable given the potential purely morphological (rather than semantic) bias due to the derivational relatedness of the morphological antonym, that could be predicted to favour its replacement by the target positive adjective to some degree.

range of senses (see, e.g., Murphy [4]).

We believe that these results, that contradict our initial hypothesis, open up new avenues for future research in this area, despite the limitations of the present study, which has an exploratory nature and a narrow empirical coverage. Indeed, only 5 adjectives were analyzed, exclusively belonging to core vocabulary. Another shortcoming is that, unlike for English, there are no in-depth studies on antonyms in Italian. However, we want to stress the importance of conducting studies on languages other than English to avoid the well-known Anglocentric bias.

Hopefully, our results will be challenged by further studies in the future, which might even overturn our conclusions entirely, if a larger data set were considered. Furthermore, it would be interesting to investigate whether the results obtained for Italian are also found for other languages that present both lexical and morphological antonyms, including languages with a different morphological system. With a view of deepening the analysis methodologically, it would be interesting to focus on additional linguistic factors that might drive the choice between lexical and morphological antonyms, such as semantic networks or word frequency, and to expand the testing to the psycholinguistic dimension.

What is sure is that the relationship between morphological and lexical antonyms is more complex than previously thought and that the choice of one type of antonym over another depends on a variety of interconnected factors that are still to be fully unveiled.

References

- [1] S. Jones, M. L. Murphy, C. Paradis, C. Willners, *Antonyms in English: Construals, Constructions and Canonicity*, Cambridge University Press, 2012.
- [2] D. A. Cruse, *Lexical semantics*, Cambridge University Press, 1986.

- [3] E. Sapir, Grading, a study in semantics., *Philosophy of Science* 11(2) (1944) 93–116.
- [4] M. L. Murphy, *Semantic relations and the lexicon*, Cambridge University Press, 2003.
- [5] C. Paradis, C. Willners, Antonymy and negation - the boundedness hypothesis, *Journal of Pragmatics* 38 (2006) 1051–1080.
- [6] D. S. Palermo, J. J. Jenkins, *Word association norms: grade school through college*, Minneapolis: University of Minnesota Press, 1964.
- [7] J. Deese, *The structure of associations in language and thought*, Baltimore: Johns Hopkins University Press, 1965.
- [8] W. G. Charles, G. A. Miller, Contexts of antonymous adjectives, *Applied Psycholinguistics* 10 (1989) 357–375.
- [9] D. J. Hermann, G. Conti, D. Peters, P. H. Robbins, R. J. S. Chaffin, Comprehension of antonymy and the generality of categorization models, *Journal of Experimental Psychology: Human Learning and Memory* 5 (1979) 585–597.
- [10] D. Gross, U. Fischer, G. A. Miller, The organization of adjectival meanings, *Journal of Memory and Language* 28 (1989) 92–106.
- [11] S. Jones, *Antonymy: a corpus-based approach*, Routledge, 2002.
- [12] J. S. Justeson, S. M. Katz, Co-occurrences of antonymous adjectives and their contexts, *Computational Linguistics* 17 (1991) 1–20. URL: <https://aclanthology.org/J91-1001>.
- [13] L. Aina, R. Bernardi, R. Fernández, Negated adjectives and antonyms in distributional semantics: not similar, in: *Proceedings of the Fifth Italian Conference on Computational Linguistics, IJCoL*, 2019, pp. 57–71. URL: <http://journals.openedition.org/ijcol/457>. doi:<https://doi.org/10.4000/ijcol.457>.
- [14] M. Koptjevskaja-Tamm, M. Miestamo, C. Börstell, A cross-linguistic study of lexical and derived antonymy, *Linguistics* (2024). URL: <https://doi.org/10.1515/ling-2023-0140>. doi:doi: 10.1515/ling-2023-0140.
- [15] F. Vicario, Note sull'ordine degli elementi in coppie di verbi antonimi., *Linguistica* XLIII 43 (2003) 3–12.
- [16] V. Muehleisen, *Antonymy and Semantic range in English*, Ph.D. thesis, Northwestern University, Evanston, 1997.
- [17] G. K. Zipf, *Human behavior and the principle of least effort*, Cambridge Addison-Wesley, 1949.
- [18] R. von Jhering, *Der Zweck im Recht*, Leipzig: Breitkopf and Härtel, 1883.
- [19] K. E. Zimmer, Affixal negation in english and other languages: an investigation of restricted productivity, *Word. Journal of the Linguistic Circle of New York* 20 (1964).
- [20] C. Iacobini, Prefissazione, in: M. G. e F. Rainer (Ed.), *La formazione delle parole in italiano*, Tübingen: Niemeyer, 2004, pp. 97–163.
- [21] G. Pittàno, *Il grande dizionario dei Sinonimi e dei Contrari*, Zanichelli, 2013.
- [22] N. Zingarelli, *Zingarelli 2024. Vocabolario della lingua italiana*, Zanichelli, 2023.
- [23] M. Aronoff, Competition and the lexicon, in: E. Annibale, C. Iacobini, M. Voghera (Eds.), *Livelli di analisi e fenomeni di interfaccia*, Roma: Bulzoni, 2016.
- [24] A. Niwa, K. Nishiguchi, N. Okazaki, Predicting antonyms in context using BERT, in: A. Belz, A. Fan, E. Reiter, Y. Sripada (Eds.), *Proceedings of the 14th International Conference on Natural Language Generation*, Association for Computational Linguistics, Aberdeen, Scotland, UK, 2021, pp. 48–54. URL: <https://aclanthology.org/2021.inlg-1.6>. doi:10.18653/v1/2021.inlg-1.6.

APPENDIX A

Noun	infelice	triste	accuracy mtp (M)	accuracy mtp (L)	accuracy pf (M)	accuracy pf (L)
matrimonio	886	24	0.9	0.5	0.2	0.4
scelta	821	39	0.8	0.5	0	0.1
adolescenza	33	16	1	0.2	0	0.2
donna	325	316	0.5	0.8	0.2	0.1
uomo	300	440	0.5	0.6	0.4	0.1
situazione	120	339	0.4	0.9	0	0
momento	145	2.686	0.1	1	0	0.2
pagina	19	544	0	1	0.2	0
Noun	impreciso	approssimativo	accuracy mtp (M)	accuracy mtp (L)	accuracy pf (M)	accuracy pf (L)
affermazione	64	11	0.2	0.8	0.1	0.1
notizia	223	58	0.4	0.8	0	0.1
terminologia	23	17	0.3	0.7	0.1	0.2
ricezione	50	18	0.7	0.3	0.2	0.1
misurazione	31	52	0.4	0.9	0.1	0.5
traduzione	54	235	0.2	0.7	0	0
conoscenza	33	226	0.1	0.9	0	0
calcolo	23	1.305	0.2	1	0.1	0.6
Noun	scorretto	sbagliato	accuracy mtp (M)	accuracy mtp (L)	accuracy pf (M)	accuracy pf (L)
gioco	1.124	72	0.1	0.7	0	0
uso	2.901	276	0	1	0	0.3
alimentazione	2.926	1.228	0	0.9	0	0.1
posizione	1.317	910	0.1	1	0.3	0.2
abitudine	648	1.077	0.1	1	0.1	0
informazione	720	1.903	0	1	0.2	0.3
mossa	90	741	0	1	0	0.1
messaggio	88	1.214	0.1	0.9	0	0.1
Noun	imprudente	avventato	accuracy mtp (M)	accuracy mtp (L)	accuracy pf (M)	accuracy pf (L)
condotta	394	14	0.2	0.8	0	0.2
comportamento	680	77	0.5	0.8	0	0
parola	49	46	0.3	0.5	0	0
manovra	24	35	0.5	0.6	0.1	0.4
gesto	54	195	0.7	0.4	0	0
azione	63	245	0.3	0.8	0	0
scelta	73	373	0.5	0.7	0	0
decisione	24	478	0.5	0.5	0	0
Noun	insufficiente	scarso	accuracy mtp (M)	accuracy mtp (L)	accuracy pf (M)	accuracy pf (L)
apporto	263	19	0.9	0.4	0	0.2
quantità	1.097	162	1	0.4	0.1	0
alimentazione	300	120	1	0.5	0	0
produzione	208	158	0.9	0.4	0	0
utilizzo	16	27	0.9	0.9	0	0
pulizia	63	148	0.4	0.8	0	0
partecipazione	13	37	0.5	0.9	0	0
visibilità	14	274	0.4	1	0	0.1

Table 3

Co-occurrence frequencies of noun + morphological antonym and noun + lexical antonym.

Accuracy of the two task: mtp (Masked-Token Prediction) and pf (Polarity Flip) related to Morphological Antonyms (M) and Lexical Antonyms (L).

morphological antonym	TTR	lexical antonym	TTR
infelice	0.4694864048	triste	0.4504979496
impreciso	0.4726656991	approssimativo	0.4814814815
scorretto	0.4496086106	sbagliato	0.4500775996
imprudente	0.476119403	avventato	0.4644572526
insufficiente	0.4582118562	scarso	0.4805725971

Table 4
Token Type Ratio of 5 antonym pair from sentences extracted from itTenTen20

APPENDIX B

system_message = '''In una frase l'aggettivo originale è stato sostituito da un token [MASK]. Tu devi riscrivere la frase facendo minimi cambiamenti e sostituire l'aggettivo mascherato con un altro aggettivo, in modo che la frase risulti volta al positivo.

Il tuo output deve essere SOLO un json nel seguente formato e con i seguenti campi:

```
{
  "new_sentence": "<tua nuova frase>",
  "new_adj": "<l'aggettivo con cui hai sostituito [MASK] nella nuova frase>"
}
```

user_message = f'''Frase originale: "{masked_sent}" (aggettivo originale: {agg})'''

{system_message = f'''In a sentence, the original adjective has been replaced by a [MASK] token.

You need to rewrite the sentence making minimal changes and replace the masked adjective with another adjective, so that the sentence is positively oriented.

Your output must be ONLY a json in the following format and with the following fields:\n''' + '''{"new_sentence": "< your new sentence>", "new_adj": "<the adjective with which you replaced [MASK] in the new sentence>"}'''

user_message = f'''Original sentence: "{masked_sent}" (original adjective: {agg})'''