

Homonymy Information for English WordNet

Rowan Hall Maudslay Simone Teufel

Dept. of Computer Science and Technology

University of Cambridge

{rh635, sht25}@cam.ac.uk

Abstract

A widely acknowledged shortcoming of WordNet is that it lacks a distinction between word meanings which are systematically related (polysemy), and those which are coincidental (homonymy). Several previous works have attempted to fill this gap, by inferring this information using computational methods. We revisit this task, and exploit recent advances in language modelling to synthesise homonymy annotation for Princeton WordNet. Previous approaches treat the problem using clustering methods; by contrast, our method works by linking WordNet to the Oxford English Dictionary, which contains the information we need. To perform this alignment, we pair definitions based on their proximity in an embedding space produced by a Transformer model. Despite the simplicity of this approach, our best model attains an F1 of .97 on an evaluation set that we annotate. The outcome of our work is a high-quality homonymy annotation layer for Princeton WordNet, which we release.

Keywords: WordNet, Oxford English Dictionary, polysemy, homonymy

1. Introduction

Words have multiple meanings that are related to each other in different ways. Meanings which are systematically related are said to exhibit **polysemy**. One example of polysemy is the use of the same wordform to refer to a product or its producer (Pustejovsky, 1995):

- (1) a. John spilled coffee on the *newspaper*.
- b. The *newspaper* fired its editor.

Aside from such highly productive alternation patterns, polysemy also includes semi-productive metaphorical extensions (Lakoff and Johnson, 1980):

- (2) a. They *adopted* a child.
- b. The theory was rapidly *adopted*.

Polysemy exemplifies humans' ability to flexibly extend categories to cover new members, which is of significant interest to researchers in cognitive science (Lakoff, 1987). These extensions include figurative uses, like in example (2). The polysemisation of words also plays a key role in lexical evolution and semantic drift (e.g. Koch, 2016).

On the other hand, meanings of the same word which exhibit no systematic relation are described as instances of **homonymy**.¹ These associations are non-productive, and result instead from language change. Usually, this occurs when new word senses are borrowed from other languages, and can involve vowelshifts and similar transformations. For example, consider the English word *bank*:

- (3) a. I need to get money out from the *bank*.
- b. Let's sit by the river on the *bank*.

¹This is sometimes called 'incidental polysemy', which is contrasted with 'systematic polysemy' (e.g. Pustejovsky, 1995).

The financial sense has its origin in the romance languages, and the river-edge sense comes from Old Norse. Another example of homonymy happens when acronyms become conventionalised, and are ultimately lower cased (e.g. Personal Identification Number):

- (4) a. Put a *pin* in the hem of the fabric.
- b. Never share your credit card's *pin*.

Although homonymous meanings are not semantically related, their presence in a particular language is not random, and instead may serve a communicative function (Piantadosi et al., 2012).

WordNet (Miller, 1995) is a popular computational lexicon. In WordNet, concepts are represented as an equivalence class of wordforms associated with that concept, called synsets. WordNet makes no distinction between polysemy and homonymy. If it did, WordNet would have the potential to be an ideal repository for research into these phenomena.

Several researchers have acknowledged this shortcoming of WordNet, and have attempted to produce computational models to synthesise homonymy annotation for it (e.g. Utt and Padó, 2011; Veale, 2004; Freihat et al., 2013). We revisit this task using contemporary methods. By exploiting large language models, we synthesise a high-quality annotation layer for distinguishing between polysemy and homonymy in the English Princeton WordNet.

More specifically, to identify homonyms in WordNet, we align it with the Oxford English Dictionary, a historical dictionary of English. In this dictionary, as a general principle in lexicography, a lemma is defined as a wordform plus all its polysemous senses. Homonymous wordforms are associated with multiple lemmas. By aligning the senses in WordNet with corresponding senses in the Oxford English Dictionary, we can work out which lemma they belong to, and thus distinguish between senses which are related by polysemy

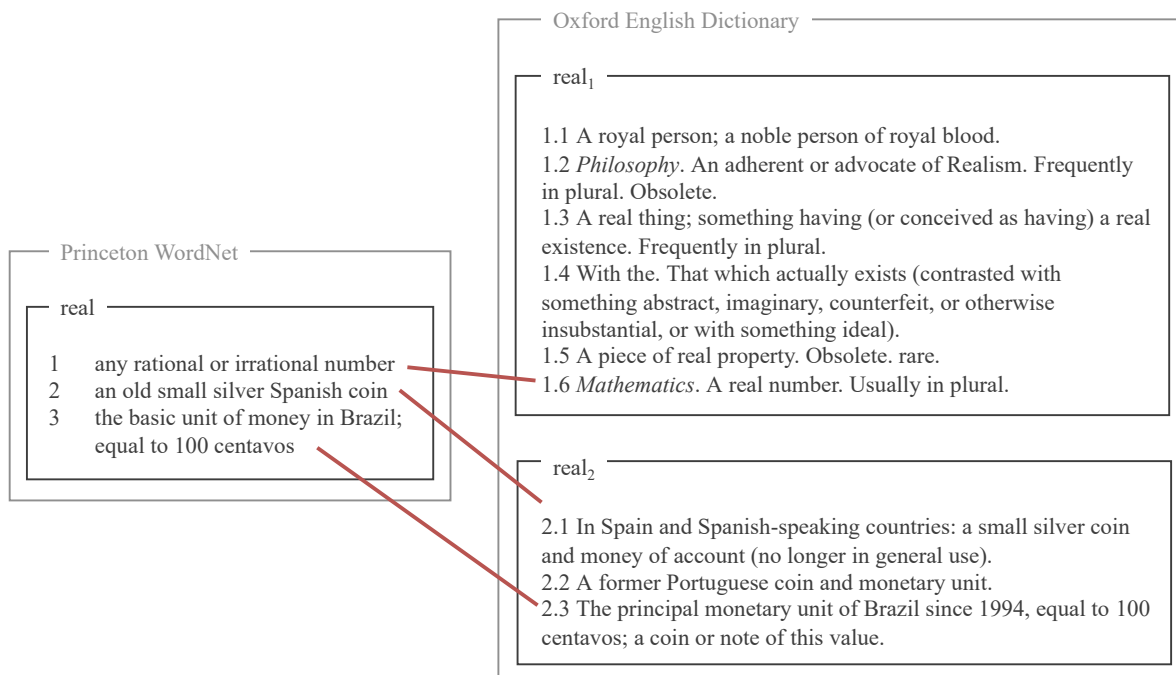


Figure 1: Noun definitions of the word *real* from the PWN (left) and the OED (right)

(same lemma), and those related by homonymy (different lemmas). Previous works that attempted to identify homonymy in WordNet did so by clustering senses. An advantage of our linking approach is that figurative senses can be correctly identified as instances of polysemy, even though their meaning might differ radically from the literal sense they extend.

To align the dictionaries, we compute the sentence embeddings of each definition using various Transformer models (Vaswani et al., 2017), and find the definition in the Oxford English Dictionary which is closest in embedding space to each WordNet definition. To evaluate the quality of the model, we annotate a small evaluation set of 196 words (554 senses). Despite the simplicity of our unsupervised method, it attains an F1-score of .97 on our evaluation set, indicating that our synthesised data is high quality.

2. Background

The **Princeton WordNet** (PWN) is an English computational lexicon, which maps wordforms to concepts, which are called synsets (Miller, 1995). Synsets are associated with a definition and often some example sentences, and are also linked to each other in a semantic network (consisting primarily of *is a* and *has a* relations). Since its creation, several works have added additional annotation layers to the PWN (e.g. Mendes and Chaves, 2001, Puşcaşu and Mititelu, 2008, Amaro et al., 2006). In research on polysemy and homonymy, we often want to build rich representations of each sense, and the PWN is associated with useful resources for that. One set of resources links synsets with textual examples, e.g. SemCor (Miller et al., 1994) and

the NTU-MC (Tan and Bond, 2011). Other resources link synsets to images depicting the synset, e.g. ImageNet (Deng et al., 2009) and BabelPic (Calabrese et al., 2020).

What the PWN lacks, however, is information which distinguishes homonymy from polysemy. Consider the word *real*, the noun senses of which are shown in Figure 1. In the PWN (left), the senses appear in a single group. In the **Oxford English Dictionary** (OED), however, the senses are divided into two separate lemmas, *real*₁ and *real*₂ (right).² The OED is an authoritative English historical dictionary: unlike the PWN, which is a contemporary lexicon that shows a snapshot of current English usage, the OED maps each word-form to all known senses that it has ever had. Senses in the same lemma have the same etymology and pronunciation, and are likely derived from each other, i.e. they are polysemous. Senses in different lemmas likely bare no systematic relation, i.e. they are homonymous. The word *real* exhibits homonymy, but the PWN does not encode this information.

The problem of separating homonymy from polysemy in the PWN has been recognised, and several works have attempted to address it. Because manually annotating this information for all of the PWN would be expensive, previous approaches have synthesised the data using computational methods (e.g. Utt and Padó, 2011, Veale, 2004, Freihart et al., 2013). These previous works all adopt a similarity-driven clustering ap-

²These are the lemmas that result following our homonymy identification procedure, which is detailed in §3.1.

proach to separate homonymy from polysemy. The problem with this approach is that some polysemous senses appear “further apart” in semantic space than homonyms. For example, two polysemous senses connected by metaphor are often extremely different on the surface (e.g. the *body* of a human v. the *body* of a guitar), and so are easily confused with homonymy even though they are related.

To ensure that instances figurative polysemy are not incorrectly labelled as homonymy, we use etymological information for the identification of homonyms. More specifically, we align WordNet with the OED (red lines in Figure 1). Our work is most similar to Navigli (2006), who also aligned the PWN with the OED to cluster PWN senses. However, while their clustering was produced for the purpose of Word Sense Disambiguation (WSD), we do so for the purpose of research into polysemy and homonymy; because of these research aims, we coarsen the OED lemmas, as outlined in §3.1.

The data synthesised by Navigli (2006), originally released for a 2007 shared task (Navigli et al., 2007), clusters WordNet 2.1 senses. Since the early time of this work, many new methodologies for dictionary alignment have emerged. Several works have aligned WordNet with other resources, for example Wiktionary and Wikipedia (Miller and Gurevych, 2014; Meyer and Gurevych, 2011; McCrae et al., 2012; Navigli et al., 2021). Recently, a shared task was held on supervised monolingual dictionary alignment (Kernerman et al., 2020); in the English subtask, models were tasked with aligning the PWN with a publicly accessible version of the Webster’s dictionary from 1913 (Ahmadi et al., 2020). All models participating in the subtask use a Transformer model (Vaswani et al., 2017) in some form. Transformer models are sentence encoders, which produce embeddings for each input token. In our work, we revisit Navigli (2006), and use Transformer models to produce a high quality alignment for WordNet 3.1.

Finally, we note that while a resource called ‘Etymological Wordnet’ already exists (de Melo, 2014), this resource is in fact unrelated to the WordNet project (Miller, 1995): it is an automatically extracted database of wordform derivations from Wiktionary.

3. Processing the OED

In this section, we describe how we extract homonymy data from the OED (§3.1), and then how we collect data to evaluate model performance (§3.2).

3.1. Extracting Homonyms from the OED

For every wordform with multiple senses in the PWN, we retrieve the corresponding lemmas from the OED.³ Lemmas in the OED have etymology data associated with them, in the form of the language family of origin. Depending on the records available, some lemmas

are annotated with more broad family information (e.g. Italic), while others have more fine grained information (e.g. French). Some have unknown origin. Because of this, sometimes it is ambiguous as to whether two lemmas are in fact related.

In these cases, we have to make a decision. We could either divide PWN senses into the lemmas as they are presented in the OED (and risk splitting polysemous senses into different lemmas), or we could merge lemmas together (and risk putting homonymous senses into the same lemma). We choose to do that latter, because for research in these areas it is preferable to overestimate polysemy and underestimate homonymy: if two polysemous senses were wrongly separated into different lemmas, this would provide a wrong gold standard for any model of polysemisation.

Our procedure for merging OED lemmas is as follows. Some lemmas are marked as being derived from others; in this case, we merge them with the lemma they are derived from. If there are multiple lemmas which have the same etymological derivation, we merge them. If one lemma’s derivation is a subclass of another’s (as with French v. Italic), we merge them. The exception to these merges is when a derivation is labelled as being the conventionalisation of an acronym; we leave these in their own lemma. Finally, if a lemma for a particular wordform has unknown etymology, we exclude that wordform (and thus assume that all its senses are polysemous).

3.2. Annotating an Evaluation Set

Sampling Data We sample wordform–part-of-speech combinations, which meet the following criteria:

- have at least two senses in the PWN;
- have at least two lemmas in the OED (following our coarsening procedure, §3.1), and further, that at least two of these lemmas have at least two senses (to avoid severely imbalanced lemmas);
- have a maximum of 15 senses overall in the OED (to reduce the cognitive load on annotators)

Following the above procedure, we sample 100 wordform–part-of-speech combinations. These combinations had an average of 2.18 lemma options in the OED, and yielded 286 PWN senses.

Annotation Procedure We need to collect a mapping of PWN senses to OED lemmas. However, as we will see in §4, the models we study work by aligning PWN senses to OED senses. Although this is not our primary concern, it would be interesting to also evaluate how well models perform at this finer granularity of analysis. Because of this, we decide to ask annotators to assign each PWN sense to a single OED sense, from which we can trivially recover the sense-to-lemma mapping which is our main interest. More specifically, we ask annotators to go through each

³Content provided by OED Researcher API, 2022.

word, and assign each PWN senses to a single OED sense. If there are multiple OED senses which would work, we ask them to select the best one. If there is no OED sense to align a PWN sense to, but there is an OED sense which is more broad and would include that PWN sense, we ask them to select that OED sense. If there is still not an appropriate OED sense, annotators have a choice. If they think the PWN sense is closely related to OED senses in a particular lemma, they assign the PWN sense to that lemma. Otherwise, if they think that the PWN sense is a different lemma, not contained in the OED, they leave it unassigned.

Recovering Lemma Assignments With the fine-grained sense-to-sense alignment which our annotators produce, we can reconstruct the sense-to-lemma mapping trivially. For each PWN sense that is aligned with an OED sense, we simply take the lemma that that OED sense is contained within in the OED.

Statistics and Agreement Two native British English speakers performed our annotation task. It is however not possible to report agreement in terms of chance-corrected Inter-Annotator Agreement (IAA) for a dictionary alignment task, because the number of possible categories that an item is assigned to varies depending on the wordform; we therefore report raw agreement. Both annotators gave the same lemma assignment 97.6% of the time, and the same sense assignment 80.4% of the time. 1.0% of the time, at least one annotator judged that no lemma existed for a PWN sense. 9.1% of the time, at least one annotator judged that none of the fine-grained senses was appropriate, but that an appropriate lemma existed. For comparability to similar tasks, we follow Ahmadi et al. (2020), and also compute IAA in terms of κ . Ahmadi et al. do this by treating each possible pair of senses (one from each dictionary) as a binary datapoint, which could be labelled 0 if they were not aligned, or 1 if they were. (However, we note that this method is problematic, as it overestimates agreement. This is because computations of κ assume that each datapoint is independent, and under this formulation many of the datapoints are counted as agreement although they are simply a consequence of other decisions.) Under these conditions, we find $\kappa=0.96$ ($N=909$, $k=2$, $n=2$) for the lemma assignments, and $\kappa=0.79$ ($N=3,396$, $k=2$, $n=2$) for the sense assignments. The high agreement is in line with previous work; Navigli (2006) found $\kappa = 0.85$ for sense-level alignment between the PWN and the OED (although it is unclear how they performed this computation).

Evaluation Data Having shown that our annotation procedure yielded high agreement, one annotator continued the annotation task for more examples, and labelled 96 more wordforms which met the above criteria. This yields a final annotated set consisting of 196 wordform-part-of-speech combinations covering 544 PWN senses, which we will use to evaluate model

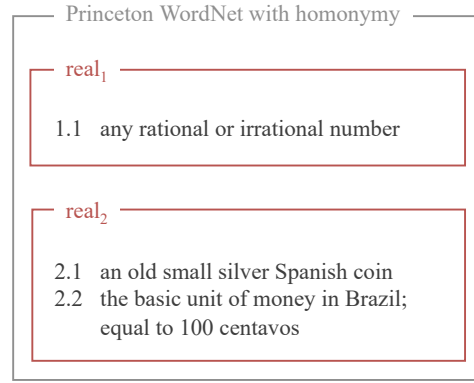


Figure 2: Our output annotation for the word *real*

performance, §5. In this final evaluation data, 1.3% of PWN senses are not assigned to an OED lemma.

4. Method

Our goal is to split homonymous PWN senses into separate lemmas (Figure 2). To achieve this, we align the PWN with the OED, in which senses are grouped according to their etymological derivation. Our method is a simple unsupervised approach, which pairs each definition from the PWN with the definition in the OED that it is closest to it in embedding space.

Let \mathcal{S} be a set of all senses, which we take as string definitions. Let $\mathcal{S}_{\text{OED}}^w \subseteq \mathcal{S}$ denote the set of sense definitions associated with a wordform w in the OED, and $\mathcal{S}_{\text{PWN}}^w \subseteq \mathcal{S}$ denote its senses in the PWN. Each sense in the OED is part of a lemma, $l \in \mathcal{L}$, which can be recovered trivially; we denote the function for doing so $\text{lemma}_{\text{OED}}^w : \mathcal{S}_{\text{OED}}^w \mapsto \mathcal{L}$. Our goal is to also map each sense from the PWN to a one of these lemmas, i.e. to construct a function, $\text{lemma}_{\text{PWN}}^w : \mathcal{S}_{\text{PWN}}^w \mapsto \mathcal{L}$.

No training data for this task exists, so we experiment with simple unsupervised methods. Let sim be a function which takes a pair of definitions, one from each dictionary, and returns a measure of their similarity, $\text{sim} : \mathcal{S}_{\text{PWN}}^w \times \mathcal{S}_{\text{OED}}^w \mapsto \mathbb{R}$. For a particular PWN sense, $s \in \mathcal{S}_{\text{PWN}}^w$, these unsupervised models assign the sense to the lemma of the most similar OED sense:

$$\text{lemma}_{\text{PWN}}^w(s) = \text{lemma}_{\text{OED}}^w\left(\underset{s' \in \mathcal{S}_{\text{OED}}^w}{\text{argmax}} \text{sim}(s, s')\right) \quad (1)$$

Our methods vary, then, in how they define sim . We experiment with very simple approaches, which compute similarity by comparing two definition embeddings. Let emb be a function that produces a d -dimensional sentence embedding of a given definition, $\text{emb} : \mathcal{S} \mapsto \mathbb{N}^d$. Additionally, let proximity be a function which compares two definition embeddings and returns a similarity rating, $\text{proximity} : \mathbb{N}^d \times \mathbb{N}^d \mapsto \mathbb{R}$. We can then express sim in terms of these functions:

$$\text{sim}(s, s') = \text{proximity}(\text{emb}(s), \text{emb}(s')) \quad (2)$$

This formulation allows us to experiment with a variety of different implementations of each of these functions, which we detail in §5.1.

5. Evaluation

All of our models are unsupervised, and parameter-free. Each model makes a prediction for each PWN sense in the evaluation data in terms of which lemma in the OED it belongs to. In this section, we evaluate how well they do so.

5.1. Experimental Setup

Data To evaluate our models, we use the data we collected in §3.2, which consists of 196 word–part-of-speech combinations, covering 554 PWN senses. When we evaluate the lemma assignments, we analyse all 554 senses, for an accurate idea of how the model will perform on the real data (and therefore include senses which were not assigned to a lemma, which the models will necessarily label incorrectly). When we evaluate the sense assignments, however, we filter out all the senses which were not assigned to a sense, leaving 497 senses.

Models Our model formulation centres around a similarity function, eq. (2), which has two main components, *emb* and *proximity*. For *emb*, we experiment with four different sentence embedding models. **GloVe** (Pennington et al., 2014) is a static embedding technique, which learns to approximate a collocation matrix. **RoBERTa** (Liu et al., 2019) is a variant of BERT (Devlin et al., 2019), a Transformer model (Vaswani et al., 2017) which was trained on a masked language modelling objective. For both of these embedding spaces, the sentence embedding is taken as the mean of all the token embeddings. The next two models, **MPNet** (Song et al., 2020) and **Sentence-T5** (Ni et al., 2021), however, were designed explicitly to produce quality sentence representations. MPNet was trained on a variety of tasks for all-round performance, while Sentence-T5 was trained on sentence similarity tasks in particular. For all of these sentence embedding models, we use the implementations in the Sentence Transformers Python library (Reimers and Gurevych, 2019); where multiple versions are present, we use the largest available. The dimensionalities (d) of these model’s representations are detailed in Table 1. Each of these embedding spaces might suit different similarity metrics, so for *proximity*, we experiment with dot product, cosine similarity, and Euclidean distance.⁴ Results presented are from whichever similarity metric attained the highest results (in all cases it was dot product).

Baselines We experiment with three baselines. As a lower bound for the task, the **random** baseline assigns each sense to a random lemma for a particular word with uniform probability. Because some lemmas have

⁴Since Euclidean distance is highest for two senses which are the least similar, we take its negation.

Name	d
GloVe	300
RoBERTa	1,024
MPNet	768
Sentence-T5	768

Table 1: Sentence embedding dimensionalities

more senses than others in the OED, we compute another baseline which assigns each sense to whichever lemma for the word has the **most** OED senses. Finally, following Navigli (2006), we reimplement the LESK algorithm (Lesk, 1986). The **LESK** baseline calculates the similarity between two definitions, s and s' , as the fraction of the shortest definition’s lemmas which are in both string definitions:

$$\text{sim}(s, s') = \frac{|\text{bow}(s) \cap \text{bow}(s')|}{\min(|\text{bow}(s)|, |\text{bow}(s')|)} \quad (3)$$

where *bow* (bag-of-words) returns the set of lemmas in a given definition. This implementation of *sim* is used to find lemma assignments using the same algorithm as the other models, eq. (1). To tokenise the definitions and to lemmatise the tokens, we use the word tokeniser and WordNet lemmatiser from NLTK (Bird et al., 2009). We additionally filter out stop words and punctuation, also using the NLTK list for stop words.

Metrics To evaluate the quality of the lemma assignments, we compute accuracy and the F1-score (macro-averaged over the lemmas). Finding the system that performs best at this level is the core interest in this paper. What is important is that a system maps each PWN sense to the correct lemma, which it can do successfully by mapping it to *any* OED sense of that lemma; even if it managed to additionally guess the finer-grained OED sense, this would only be of secondary interest to us. However, we are in a situation where we can report performance at a finer granularity because each model internally predicts a finer-grained OED sense. We therefore additionally report F1-score and accuracy of these sense assignments (macro-averaged over OED senses).

Significance Testing We use a two-tailed Monte Carlo permutation test at significance level $\alpha = 0.01$, with $r = 10,000$ permutations.

5.2. Results

Table 2 shows our results. Two of the baselines, random and majority, only make lemma assignments, and so we cannot evaluate them at the sense level.

The best performing model overall used the Sentence-T5 embedding space. Despite the simplicity of this approach, it attained an F1-score of 0.97 in the lemma assignment task, the main focus of this work. This was significantly better than all the baselines, and also significantly better than GloVe, the only non-Transformer

Model	Lemma Assignments		Sense Assignments	
	Accuracy	F1-score	Accuracy	F1-Score
GloVe	.94	.93	.71	.70
MPNet	.94	.95	.76	.75
RoBERTa	.95	.95	.72	.71
Sentence-T5	.97	.97	.84	.84
LESK	.88	.88	.65	.63
most	.73	.68	N/A	N/A
random	.47	.50	N/A	N/A

Table 2: Results

embedding space. Numerically, the difference in the lemma scores was small: GloVe embeddings achieved .93 F1, only .04 less than Sentence-T5.

In the evaluation data we collected, 1.3% of senses were not assigned to a lemma (see §3.2). Our model necessarily gets all of these wrong (it has no way of leaving senses unassigned), meaning the highest accuracy it could theoretically attain would be .98—only .01 higher than it achieves. For our purposes, that it erroneously assigns these senses is not an issue: as mentioned above (§3.1), because we are interested in research into polysemy and homonymy, we opt to overestimate polysemy and underestimate homonymy, rather than vice versa. This is the effect which this will have. The best model at predicting the sense-to-sense mapping also used the Sentence-T5 embedding space, but the quality of the mapping was not as high as its sense-to-lemma mapping, attaining an F1 of .84. This result is significantly better than not only GloVe, but also both other Transformer models. The numerical difference between the models is also more pronounced. GloVe attained .70 F1, which is .14 behind the best Transformer model, and only .07 above the LESK approach.

6. Final Annotation Layer

Having performed an evaluation of our approach on a small testset, we now present details for the entirety of the PWN. We use the highest-performing model from our evaluation, which was based on the Sentence-T5 (Ni et al., 2021) embedding space, and used the dot product to compare embeddings.

6.1. Between-POS v. Within-POS

We compute two distinct annotation layer variants, which we term between-POS and within-POS.

The OED is an etymological lexicon, and as such it can identify when two lemmas of the same wordform, but with different parts-of-speech, are derived from each other (this process is called zero-derivation). For example, as a verb, *tango* is to perform a particular dance, and as a noun, *tango* is that dance. In the **between-POS** homonymy annotation layer, we preserve this information, by applying our homonymy identification procedure (§3.1) to all the senses of a word at the same time, regardless of their part-of-speech.

This approach has one drawback. As mentioned above, the OED does not have complete information about all senses’ etymologies. Sometimes, a sense might be labelled with less specific information than another, or might have unknown etymology. When a wordform had a sense with unknown etymology, we assumed that no homonymy was present, i.e. that all the wordform’s senses were polysemous. This is to reduce the chance of erroneously labelling instances of polysemy as homonymy. However, in cases where a sense has unknown etymology, there is a chance that we incorrectly treat instances of homonymy as polysemy, an error which we would also like to minimise.

The more senses a wordform has, the more likely it is to have a sense with missing information, which may mean that it is incorrectly treated. In the **within-POS** layer, when applying our homonymy identification procedure, we treat the senses of each part-of-speech individually. This reduces the chance that a sense will be included which lacks etymology information, and so lowers the chance of missing instances of homonymy. However, this comes at the price of losing the alignment between different parts-of-speech.

In both the between-POS and within-POS variants, we exclude OED senses which were not part of the alignment. In other words, we first compute the alignment between the PWN and the OED, and then apply our homonymy identification procedure to only the OED senses which are part of the alignment. This is to minimise the unwanted effects of senses with unknown etymology as much as possible, for both variants.

6.2. Analysis

Statistics for the two variants of our annotation layer are presented in Table 3. We additionally report counts using out-of-the-box lemmas from the OED, without any of the processing in §3.1; reported as **raw**. This should give an idea of the number of exclusions resulting from our homonymy identification process. There are a total of 21,740 words which have multiple senses in the PWN.⁵ Of those, 20,169 (93%) have corresponding entries in the OED.

⁵We exclude all wordforms which are not lower case or which include spaces; this removes proper nouns and compound nouns, because these are not included in the OED.

POS	# Words in PWN	# Also in OED	# Homonymous in the OED			# Homonymous in the PWN		
			between-POS	within-POS	raw	between-POS	within-POS	raw
noun	15,019	14,228	806	849	2,830	237	244	794
verb	6,226	5,886	237	310	1,218	50	56	244
adj	6,661	6,115	75	88	303	17	17	54
adv	1,037	934	3	4	19	0	0	2
any	21,740	20,169	969	1,091	3,420	284	297	961

Table 3: Final annotation layer statistics

Using the within-POS variant, 1,091 wordforms are found to exhibit homonymy.⁶ As expected, fewer were found using the between-POS variant (969, a reduction of 11%). These numbers represent the maximum number of wordforms in the PWN which our method can identify as exhibiting homonymy. Of these, with the within-POS variant we identified 297 homonymous wordforms in the PWN (27% of those in the OED), which are associated with a total of 2,139 senses in the PWN (full list of words in App. A). With the between-POS variant we identified 284 wordforms. The fact that only a fraction of homonymous wordforms in the OED were also homonymous in the PWN is unsurprising. The OED is an etymological dictionary, which will contain senses which are no longer used. On the other hand, the PWN is a contemporary dictionary, which will not contain archaic instances of homonymy.

Clear-cut cases of homonymy are less numerable than we might expect (279 cases; ‘any’ under within-POS in Table 3). These are the cases where wordforms are associated with meanings which have distinct origins and are semantically unrelated. But then again, this number represents a lower-bound for the total amount of homonymy in the PWN, as a consequence of our decision to combine lemmas in ambiguous cases. An upper-bound (i.e. an overestimation of homonymy) is represented by the raw results (961 wordforms). This indicates that between 1.5% and 4.8% of wordforms in the PWN are homonymous (estimated using the wordforms that are in both dictionaries).

6.3. Release

We release our code and both variants of our homonymy annotation layer online.⁷ We additionally release a version based on the raw lemma assignments, which will be useful if overestimation of homonymy and underestimation of polysemy is preferred, but we caution that the quality of this data was not investigated in our annotation study.

⁶Note that for the within-POS variant, the ‘any’ part-of-speech row in Table 3 does not correspond to a simple summation of the statistics for each part-of-speech, because this would count any wordform which is homonymous in two or more different parts-of-speech multiple times.

⁷<https://github.com/rowanhm/wordnet-homonymy>

7. Conclusion

We present a new annotation layer for the Princeton WordNet, which splits senses into lemmas, making it possible to distinguish between polysemy and homonymy. We use a method which is conservative with respect to homonymy identification (we would rather erroneously label two homonymous senses as polysemous than vice versa, §3.1). Additionally, in contrast to previous work, we use an alignment-based method which will be able to correctly treat figurative polysemy. We create this annotation layer using a simple method that exploits recent advances in language modelling; although the annotation layer we produce is synthetic, the F1-score that our model attained on a small evaluation set that we produced was .97, indicating that it is of high quality.

In future work, we hope to enhance WordNet with more information. Lemmas in the OED are annotated with phonetic information; this could be used to infer homophony, which occurs when two unrelated meanings use the same phonetic form (even if they do not necessarily use the same orthographic form). An example is the word *base*, which is homophonous with the word *bass*. Additionally, if more complex models could be developed to produce a high quality sense-to-sense mapping to the OED, then we could leverage information the fine-grained senses in the OED contain about the dates of sense emergence, to make WordNet diachronic. This would be very useful in the study of language change.

Acknowledgements

We would like to thank the Oxford University Press (OUP) for giving us access to the OED research API, which made this work possible. In particular, we would like to thank Elinor Hawkes from the OUP for helping us with this.

8. Bibliographical References

Amaro, R., Chaves, R. P., Marrafa, P., and Mendes, S. (2006). Enriching WordNets with new relations and with event and argument structures. In *Proceedings of the 7th International Conference on Computational Linguistics and Intelligent Text Processing, CICLing’06*, page 28–40, Berlin, Heidelberg. Springer-Verlag.

- Bird, S., Klein, E., and Loper, E. (2009). *Natural language processing with Python: Analyzing text with the natural language toolkit*. O'Reilly Media, Inc.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. 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 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Freihat, A. A., Giunchiglia, F., and Dutta, B. (2013). Regular polysemy in WordNet and pattern based approach. *International Journal On Advances in Intelligent Systems*, 6.
- Ilan Kernerman, et al., editors. (2020). *Proceedings of the 2020 Globalex Workshop on Linked Lexicography*, Marseille, France, May. European Language Resources Association.
- Koch, P. (2016). Meaning change and semantic shifts. In Päivi Juvonen et al., editors, *The Lexical Typology of Semantic Shifts*, chapter 2, pages 21–66. De Gruyter Mouton.
- Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. University of Chicago Press.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. University of Chicago Press.
- Lesk, M. (1986). Automatic sense disambiguation using machine readable dictionaries: How to tell a pine cone from an ice cream cone. In *Proceedings of the 5th Annual International Conference on Systems Documentation, SIGDOC '86*, page 24–26, New York, NY, USA. Association for Computing Machinery.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Mendes, S. and Chaves, R. P. (2001). Enriching WordNet with qualia information. In *Proceedings of NAACL 2001 Workshop on WordNet and Other Lexical Resources*.
- Navigli, R., Litkowski, K. C., and Hargraves, O. (2007). SemEval-2007 task 07: Coarse-grained English all-words task. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 30–35, Prague, Czech Republic, June. Association for Computational Linguistics.
- Navigli, R. (2006). Meaningful clustering of senses helps boost word sense disambiguation performance. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 105–112, Sydney, Australia, July. Association for Computational Linguistics.
- Ni, J., Ábrege, G. H., Constant, N., Ma, J., Hall, K. B., Cer, D., and Yang, Y. (2021). Sentence-T5: Scalable sentence encoders from pre-trained text-to-text models. *CoRR*, abs/2108.08877.
- Pennington, J., Socher, R., and Manning, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Piantadosi, S. T., Tily, H., and Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, 122(3):280–291.
- Puşcaşu, G. and Mititelu, V. B. (2008). Annotation of WordNet verbs with TimeML event classes. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, May. European Language Resources Association (ELRA).
- Pustejovsky, J. (1995). *The generative lexicon*. MIT Press.
- Reimers, N. and Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China, November. Association for Computational Linguistics.
- Song, K., Tan, X., Qin, T., Lu, J., and Liu, T.-Y. (2020). MPNet: Masked and permuted pre-training for language understanding. In H. Larochelle, et al., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16857–16867. Curran Associates, Inc.
- Utt, J. and Padó, S. (2011). Ontology-based distinction between polysemy and homonymy. In *Proceedings of the Ninth International Conference on Computational Semantics (IWCS 2011)*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In I. Guyon, et al., editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Veale, T. (2004). Polysemy and category structure in WordNet: An evidential approach. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal, May. European Language Resources Association (ELRA).

9. Language Resource References

- Ahmadi, S., McCrae, J. P., Nimb, S., Khan, F., Monachini, M., Pedersen, B., Declerck, T., Wissik, T., Bellandi, A., Pisani, I., Troelsgård, T., Olsen, S., Krek, S., Lipp, V., Váradi, T., Simon, L., Gyorffy, A., Tiberius, C., Schoonheim, T., Ben Moshe,

- Y., Rudich, M., Abu Ahmad, R., Lonke, D., Kovalenko, K., Langemets, M., Kallas, J., Dereza, O., Franssen, T., Cillessen, D., Lindemann, D., Alonso, M., Salgado, A., Luis Sancho, J., Ureña-Ruiz, R.-J., Porta Zamorano, J., Simov, K., Osenova, P., Kancheva, Z., Radev, I., Stanković, R., Perdiš, A., and Gabrovsek, D. (2020). A multilingual evaluation dataset for monolingual word sense alignment. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3232–3242, Marseille, France, May. European Language Resources Association.
- Calabrese, A., Bevilacqua, M., and Navigli, R. (2020). Fatality killed the cat or: BabelPic, a multimodal dataset for non-concrete concepts. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4680–4686, Online, July. Association for Computational Linguistics.
- de Melo, G. (2014). Etymological Wordnet: Tracing the history of words. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1148–1154, Reykjavik, Iceland, May. European Language Resources Association (ELRA).
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255.
- McCrae, J., Montiel-Ponsoda, E., and Cimiano, P. (2012). *Integrating WordNet and Wiktionary with lemon*, pages 25–34. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Meyer, C. M. and Gurevych, I. (2011). What psycholinguists know about chemistry: Aligning Wiktionary and WordNet for increased domain coverage. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 883–892, Chiang Mai, Thailand, November. Asian Federation of Natural Language Processing.
- Miller, T. and Gurevych, I. (2014). WordNet—Wikipedia—Wiktionary: Construction of a three-way alignment. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 2094–2100, Reykjavik, Iceland, May. European Language Resources Association (ELRA).
- Miller, G. A., Chodorow, M., Landes, S., Leacock, C., and Thomas, R. G. (1994). Using a semantic concordance for sense identification. In *Proceedings of the Workshop on Human Language Technology, HLT '94*, page 240–243, USA. Association for Computational Linguistics.
- Miller, G. A. (1995). WordNet: A lexical database for English. *Commun. ACM*, 38(11):39–41, November.
- Navigli, R., Bevilacqua, M., Conia, S., Montagnini, D., and Ceconi, F. (2021). Ten years of BabelNet: A survey. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4559–4567. International Joint Conferences on Artificial Intelligence Organization, 8.
- Tan, L. and Bond, F. (2011). Building and annotating the linguistically diverse NTU-MC (NTU-multilingual corpus). In *Proceedings of the 25th Pacific Asia Conference on Language, Information and Computation*, pages 362–371, Singapore, December. Institute of Digital Enhancement of Cognitive Processing, Waseda University.

A. List of Homonyms in WordNet

The list below contains the 297 wordforms which are identified as exhibiting homonymy in the PWN. The 13 wordforms which appear in the within-POS variant but not the between-POS variant are marked with an asterisk:

adder, agora, alum, angle, apostrophe, armed, ass, ball, bank, bard, bark, bar, bath, batter, bat, beat, bill, birra, boil, bole, bongo, boom, boss*, bowl, boxer, boxing, box, bracer, buffer, buff, bumble, bust, butter, bye, calf, canon, caper, carbonado, castor, cheese, chela, chess, clove, coma, compact, compound, content, con, copper, corn, corona, cosmos, courser, cover, cramp*, croup, cube, curry, dam, deuce, dick, diet, ding, distemper, dock, don, dory, down, drill*, dub, excise, fag*, fan, fawn, feller, fen, file, filicide, filing, filler, flag*, flat, flicker, flop*, flounce, forte, fossa, full, fuse, gall, game, gauntlet, genial, gill, gin, gnarl, gnome, gobbler, gobble, go, grad, grate, grave, gray, gum, gutter, gyro, ha-ha, hack, hakim, hash, hatched, hatching, hatch, hawker, hobby, homer, hood, house, hypo, impress, indent, iridic, jack, jar*, jumper, junk, key, khan, kip, kit, krona, lame, launch, laver, letter, lien, limb, lime, ling, lister, lithic, lumber, lunger, man-akin, mandarin, mangle, mare, mark, match, matted, matting, mat, mean, meter, metric, mew, mil, miss, mogul, molar, mole, monstrance, mood, mould, mow, mummy, mush, must, nag, nanny, nap, net, nit, ore, paddle, pall, para, pass, patter, peewee, periwinkle, permit, phone, pile, pink, pipe, piping, pix, plantain, splash, plight, plonk, plump, poacher, poach, poise, poker, poke, poll, pom-pom, pool, pop, port, pot, psi, punch, punter, pyrene, pyrrhic, python, quack, quark, quid, quint, quiver, race, racy, rad, raft, raised, ramp, real, reef, rent, rest, retort, rip*, roach, rocket, rocky, rock, rook, root, round, router, rout, rue, rush, sack*, sake, salve, samba, sampler, sardine, scale, school, sconce, scope, scourer, scruple, scuffle, seal, seamy, secrete, set, sewer, shock, skipper, slug*, snarl, sod, sol, soma, sort, sound, spade, spanker, spell*, spike, stall*, stater, stay, stereo, still, stinger, stoop, strain, tack, talus, tanka, telluric, temple, test, tiller, timber, toot, topi, tower, tribune, tuck, tuna, unionized, verse, viola, yen, zip*