



HAL
open science

Unsupervised Word Polysemy Quantification with Multiresolution Grids of Contextual Embeddings

Christos Xypolopoulos, Antoine Tixier, Michalis Vazirgiannis

► **To cite this version:**

Christos Xypolopoulos, Antoine Tixier, Michalis Vazirgiannis. Unsupervised Word Polysemy Quantification with Multiresolution Grids of Contextual Embeddings. Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, Apr 2021, Online, Ukraine. pp.3391-3401, 10.18653/v1/2021.eacl-main.297 . hal-04447218

HAL Id: hal-04447218

<https://hal.science/hal-04447218v1>

Submitted on 8 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Unsupervised Word Polysemy Quantification with Multiresolution Grids of Contextual Embeddings

Christos Xypolopoulos*
École Polytechnique

Antoine J.-P. Tixier*
École Polytechnique

Michalis Vazirgiannis
École Polytechnique & AUEB

Abstract

The number of senses of a given word, or polysemy, is a very subjective notion, which varies widely across annotators and resources. We propose a novel method to estimate polysemy based on simple geometry in the contextual embedding space. Our approach is fully unsupervised and purely data-driven. Through rigorous experiments, we show that our rankings are well correlated, with strong statistical significance, with 6 different rankings derived from famous human-constructed resources such as WordNet, OntoNotes, Oxford, Wikipedia, etc., for 6 different standard metrics. We also visualize and analyze the correlation between the human rankings and make interesting observations. A valuable by-product of our method is the ability to sample, at no extra cost, sentences containing different senses of a given word. Finally, the fully unsupervised nature of our approach makes it applicable to any language. Code and data are publicly available¹.

1 Introduction

Polysemy, the number of senses that a word has, is a very subjective notion, subject to individual biases. Word sense annotation has always been one of the tasks with the lowest values of inter-annotator agreement (Artstein and Poesio, 2008). Yet, creating high-quality, consistent word sense inventories is a critical pre-requisite to successful word sense disambiguation.

Towards creating word sense inventories, it can be helpful to have some reliable information about polysemy. That is, knowing which words have

many senses and which words have only a few senses. Such information can help in creating new inventories but also in validating and interpreting existing ones. It can also help select which words to include in a study (e.g., only highly polysemous words).

We propose a novel, fully unsupervised, and data-driven approach to quantify polysemy, based on basic geometry in the contextual embedding space.

Contextual word embeddings have emerged in the last few years as part of the NLP transfer learning revolution. Now, entire deep models are pre-trained on huge amounts of unannotated data and fine-tuned on much smaller annotated datasets. Some of the most famous examples include ULM-FiT (Howard and Ruder, 2018) and ELMo (Peters et al., 2018), both based on recurrent neural networks; and GPT (Radford et al., 2018) and BERT (Devlin et al., 2018), based on transformers (Vaswani et al., 2017). These models all are very deep *language models*. During pre-training on large-scale corpora, they learn to generate powerful internal representations, including fine-grained contextual word embeddings. For instance, in a well pre-trained model, the word *python* will have two very different embeddings depending on whether it occurs in a programming context (as in, e.g., “I love to write code in python”) or in an ecological context (“while hiking in the rainforest, I saw a python”).

Our approach capitalizes on the contextual embeddings previously described. It does not involve any tool and does not rely on any human input or judgment. Also, thanks to its unsupervised nature, it can be applied to any language (even those with limited resources), provided that contextual embeddings are available.

The remainder of this paper is organized as follows. We detail our approach in section 2. Then,

* Equal contribution. CX handled the data, generated the rankings, and sampled the examples of section 7. AJPT computed the results, plots, and wrote the paper. Both authors participated in the design of the study.

¹<https://github.com/ksipos/polysemy-assessment>




Figure 1: Illustration of the proposed approach with $D = 2$ and $L = 3$.

we present our experimental setup (sec. 3), evaluation metrics (sec. 4), and report and interpret our results (sec. 6). In section 7, we briefly touch on two other interesting applications of our method. One that allows the user to sample sentences containing different senses of a given word and one that goes towards word sense induction. Finally, related work is presented in section 8.

2 Proposed approach

2.1 Basic assumption

First, by passing diverse sentences containing a given word to a pre-trained language model, we construct a representative set of vectors for that word (one vector for each occurrence of the word). The basic and intuitive assumption we make is that *the volume covered by the cloud of points in the contextual embedding space is representative of the polysemy of the associated word*.

2.2 Main idea: multiresolution grids

As a proxy for the volume covered, we adopt a simple geometrical approach. As shown in Fig. 1, we construct a hierarchical discretization of the space, where, at each level, the same number of bins are drawn along each dimension. Each level corresponds to a different resolution. Our polysemy score is based on the proportion of bins covered by the vectors of a given word at each level.

Grid vs. clustering. Using a binning strategy makes more sense than a clustering-based approach. Indeed, clusters do not partition the space equally and regularly. This is especially problematic since word representations are not uniformly distributed in the embedding space (Ethayarajh, 2019). Therefore, the vectors lying in the same dense area of the space will always belong to one single large cluster, while outliers lying in the same, but sparser,

area of the space, will be assigned to many different small clusters. Therefore, counting the number of clusters a given word belongs to is not a reliable indicator of how much of the space this word covers.

2.3 Scoring scheme

We quantify the polysemy degree of a word w as:

$$\text{score}(w) = \sum_{l=1}^L \frac{\text{coverage}_w^l}{2^{L-l}} \quad (1)$$

where coverage_w^l designates the proportion of bins covered by word w at level l , between 0 and 1. At each level, 2^l bins are drawn along each dimension (see the vertical and horizontal lines in Fig. 1). The hierarchy starts at $l = 1$ since there is only one bin covering all the space at $l = 0$ (so all words have equal coverage at this level). The total number of bins in the entire space, at a given level l , is equal to $(2^l)^D$.

Consider again the example of Fig. 1. In this example, each word is associated with a set of 10 contextualized embeddings in a space of dimension $D = 2$, and the hierarchy has $L = 3$ levels. First, we can clearly see that word 1 (blue circles) covers a large area of the space while all the vectors of word 2 (orange squares) are grouped in the same region. Intuitively, this can be interpreted as “word 1 occurs in more different contexts than word 2”, which per our assumption, is equivalent to saying that “word 1 is more polysemous than word 2”.

Let us now see how this is reflected by our scoring scheme. First, the penalization terms (denominators) for levels 1 to 3 are $[\frac{1}{2^2}, \frac{1}{2^1}, \frac{1}{2^0}] = [\frac{1}{4}, \frac{1}{2}, 1]$. Note that the higher the level, the exponentially more bins, and so the less penalized (or the more rewarded) coverage is, because getting good coverage becomes more and more difficult. Now, per

Eq. 1, the score of word 1 is computed as the dot product of its coverage vector $[\frac{3}{4}, \frac{7}{16}, \frac{10}{64}]$ (coverage at each level) with the penalization vector, which gives a score of 0.5625. Likewise, the score of word 2 is computed as $[\frac{1}{4}, \frac{1}{2}, 1] \cdot [\frac{1}{4}, \frac{4}{16}, \frac{7}{64}] = 0.297$. We can thus see that our scores reflect what can be observed in Fig. 1: word 1 covers a larger area of the space than word 2.

Note that the score of a given word is only meaningful compared to the scores of other words, i.e., in rankings, as will be seen in the next section.

3 Experiments

In this section, we describe the protocol we followed to test the extent to which our rankings match human rankings.

3.1 Word selection

The first step was to select words to include in our analysis. For this purpose, we downloaded and extracted all the text from the latest available English Wikipedia dump². We then performed tokenization, stopword, punctuation, and number removal and counted the occurrence of each token of at least 3 characters in size. Out of these tokens, we kept the 2000 most frequent.

3.2 Generating vector sets

For each word in the shortlist, we randomly selected 3000 sentences such that the corresponding word appeared exactly once within each sentence. The words that did not appear in at least 3000 sentences were removed from the analysis, reducing the shortlist’s size from 2000 to 1822. Then, for each word, the associated sentences were passed through a pre-trained ELMo model³ (Peters et al., 2018) in test mode, and the top layer representations corresponding to the word were harvested. The advantage of using ELMo’s top layer embeddings is that they are the most contextual, as shown by Ethayarajh (2019). We ended up with a set of exactly 3000 1024-dimensional contextual embeddings for each word.

3.3 Dimensionality reduction

Remember that the total number of bins in the entire space is equal to $(2^l)^D$ at a given level l , which would have given us an infinite number of bins even

²<https://dumps.wikimedia.org/>

³We used the implementation and pre-trained weights publicly released by the authors <https://allennlp.org/elmo>.

at the first level, since the ELMo representations have dimensionality $D = 1024$. To reduce the dimensionality of the contextual embedding space, we applied PCA, trying 19 different output dimensionalities, from 2 to 20 with steps of 1. Due to the quantity and high initial dimensionality of the vectors, we used the distributed⁴ version of PCA provided by the PySpark’s ML Library (Meng et al., 2016).

3.4 Score computation

We computed our scores for each PCA output dimensionality, trying with 18 different hierarchies whose numbers of levels L ranged from 2 to 19. So in total, we obtained $19 \times 18 = 342$ rankings.

3.5 Ground truth rankings and baselines

We evaluated the rankings generated by our approach against several ground truth rankings that we derived from human-constructed resources.

Since the number of senses of a word is a subjective, debatable notion, and thus may vary from source to source, we included 6 ground truth rankings in our analysis, in order to minimize source-specific bias as much as possible. For sanity checking purposes, we also added two basic baseline rankings (frequency and random). We provide more details about all rankings in what follows.

3.5.1 WordNet

We used WordNet (Miller, 1998) version 3.0 and counted the number of synonym sets or “synsets” of each word.

3.5.2 WordNet-Reduced

There are very subtle differences among the WordNet senses (“synsets”), making distinguishing between them difficult and even irrelevant in some applications (Palmer et al., 2004, 2007; Brown et al., 2010; Rumshisky, 2011; Jurgens, 2013). For instance, *call* has 41 senses in the original WordNet (28 as verb and 13 as noun). Even for other words with fewer senses, like *eating* (7 senses in total), the difference between senses can be very tiny. For instance, “take in solid food” and “eat a meal; take a meal” are really close in meaning. This very fine granularity of WordNet may somewhat artificially increase the polysemy of some words.

To reduce the granularity of the WordNet

⁴15 executors with 10 GB of RAM each.

synsets, we used their sense keys⁵. They follow the format `lemma%ss_type:lex_filenum:lex_id:head_word:head_id`, where `ss_type` represents the synset type (part-of-speech tag such as noun, verb, adjective) and `lex_filenum` represents the name of the lexicographer file containing the synset for the sense (`noun.animal`, `noun.event`, `verb.emotion`, etc.). We truncated the sense keys after `lex_filenum`.

For instance, “take in solid food” and “eat a meal; take a meal” initially correspond to two different senses with keys `eat%2:34:00::` and `eat%2:34:01::`, but after truncation, they both are mapped to the same sense: `eat%2:34`. However, coarse differences in senses are still captured. For instance, *bank* “sloping land” (`bank%1:17:01::`) and *bank* “financial institution” (`bank%1:14:00::`) are still mapped to two different senses after truncation, respectively `bank%1:17` and `bank%1:14`.

3.5.3 WordNet-Domains

WordNet Domains (Bentivogli et al., 2004; Magnini and Cavaglia, 2000) is a lexical resource created in a semi-automatic way to augment WordNet with domain labels. Instead of synsets, each word is associated with a number of semantic domains. The domains are areas of human knowledge (politics, economy, sports, etc.) exhibiting specific terminology and lexical coherence. As for the two previous WordNet ground truth rankings, we simply counted the number of domains associated with each word.

3.5.4 OntoNotes

OntoNotes (Hovy et al., 2006; Weischedel et al., 2011) is a large annotated corpus comprising various text genres (news, conversational telephone speech, weblogs, newsgroups, broadcast, talk shows) with structural information and shallow semantics.

We counted the senses in the sense inventory of each word. The senses in OntoNotes are groupings of the WordNet synsets, constructed by human annotators. As a result, the sense granularity of OntoNotes is coarser than that of WordNet (Brown et al., 2010).

3.5.5 Oxford

We counted the number of senses returned by the Oxford dictionary⁶, which was, at the time of this

⁵See ‘Sense Key Encoding’ here: <https://wordnet.princeton.edu/documentation/senseidx5wn>

⁶www.lexico.com




Figure 2: Average score distribution of the 5 ground truth rankings and frequency baseline (histogram) vs. average score distribution of the random baseline (blue curve).

study, the resource underlying the Google dictionary functionality.

3.5.6 Wikipedia

We capitalized on the Wikipedia disambiguation pages⁷. Such pages contain a list of the different categories under which one or more articles about the query word can be found. For example, the disambiguation page of the word *bank* includes categories such as geography, finance, computing (data bank), and science (blood bank). We counted the number of categories on the disambiguation page of each word to generate the ranking.

3.5.7 Frequency and random baselines

In the frequency baseline, we ranked words in decreasing order of their frequency in the entire Wikipedia dump (see subsection 3.1). The naive assumption made here is that words occurring the most have the most senses.

With the random baseline, on the other hand, we produced rankings by shuffling words. Further, we assigned them random scores by sampling from the Log Normal distribution⁸, to imitate the long-tail behavior of the other score distributions, as can be seen in Fig. 2. All distributions can be seen in Fig. 6. Note that to account for randomness, all results for the random baseline are averages over 30 runs. Not every of the 1822 words included in our analysis had an entry in each of the resources described above. The lengths of each ground truth ranking are shown in Table 1.

4 Evaluation metrics

We used 6 similarity, correlation and information retrieval standard metrics to compare among methods: cosine similarity, Spearman’s rho (Spearman,

⁷[www.wikipedia.org/wiki/word_\(disambiguation\)](http://www.wikipedia.org/wiki/word_(disambiguation))

⁸with mean and standard deviation 0 and 0.6 (resp.)

Ranking	# words
WN	1535
WN-reduced	1535
WN-Domains	1420
Oxford	1536
Wikipedia	1042
OntoNotes	723
Frequency & random	1822

Table 1: Length of the ground truth rankings.

1904), Kendall’s tau (Kendall, 1938), precision at k , Normalized Discounted Cumulative Gain (Järvelin and Kekäläinen, 2002), and Rank Biased Overlap (Webber et al., 2010).

To ensure a fair comparison, the scores in the rankings of all methods were normalized to be in the $[0, 100]$ range before proceeding.

Also, each method played in turn the role of candidate and ground truth. This allowed us to compute not only the similarity between our rankings and the ground truth rankings, but also the similarity among the ground truth rankings themselves, which was interesting for exploration purposes.

For each pair of evaluated and ground truth method, only the parts of the rankings corresponding to the words in common (intersection) were compared. Thus, the rankings in each (candidate, ground truth) pair had equal length.

5 Implementation details

To compute our scores, we built on the code of the pyramid match kernel from the GraKeL Python library (Siglidis et al., 2018). We used the base R (R Core Team, 2018) `cor()` function⁹ to compute the τ and ρ statistics. For RBO, we relied on a publicly available Python implementation¹⁰. For all other metrics, we wrote our own implementations. Full details about design choices, tokenizers, stopword lists etc., can also be found in our publicly available code repository: <https://github.com/ksipos/polysemy-assessment>.

6 Results and observations

Our rankings correlate well with human rankings. Results are shown in Fig. 3, as pairwise similarity matrices, for all six metrics. For readability, all scores are shown as percentages. For a given metric, our configuration that best matches, on average, all other methods (except random and

⁹<https://stat.ethz.ch/R-manual/R-patched/library/stats/html/cor.html>

¹⁰<https://github.com/changyaochen/rbo>

frequency) is always shown as the first column. Since all metrics except NDCG are symmetric, we only show the lower triangles of the other matrices. For NDCG, candidate methods are shown as columns and ground truths as rows.

For each of the six evaluation metrics, it can be seen that the ranking generated by our unsupervised, data-driven method is well correlated with all human-derived ground truth rankings. This means that our method is robust to how one defines and measures correlation or similarity.

In some cases, we even very closely reproduce the human rankings. For instance, our best configurations for cosine and NDCG get almost perfect scores of 86.5 and 99.72 when compared against Wikipedia. In terms of Kendall’s tau, Spearman’s rho, $p@k$, and RBO, we are also very close to OntoNotes (scores of 49.43, 35.23, 39.53, and 33.47, resp.).

Finally, the correlation between our rankings and the human rankings can also be observed to be, everywhere, much stronger than that between the baseline rankings (random and frequency) and the human rankings.

Statistical significance. We computed statistical significance for the Spearman’s rho and Kendall’s tau metrics. As shown in Fig. 3, the null hypothesis that there is no correlation between our rankings and the human-derived ground truth rankings was systematically rejected everywhere, with very high significance ($p \leq 0.0001$).

However, against the random baseline, the same null hypothesis (no correlation) was accepted everywhere. Against frequency, the null was rejected, but very weakly (only at the $p \leq 0.01$ level), and with very low correlation coefficients (6.53 for Spearman and 4.44 for Kendall).

Finally, the correlation between the random and frequency rankings and the ground truth rankings is never statistically significant, except for the pair frequency/OntoNotes, but again, at a weak level ($p \leq 0.01$).

Hyperparameters have a significant impact on performance, but optimal values are consistent across metrics. First, as can be observed from Fig. 4 and Fig. 5, there is a large variability in performance when D (number of PCA dimensions) and L (number of levels in the hierarchy) vary.

However, for all six evaluation metrics, the best configurations are very similar: $D2L10$, $D2L8$,

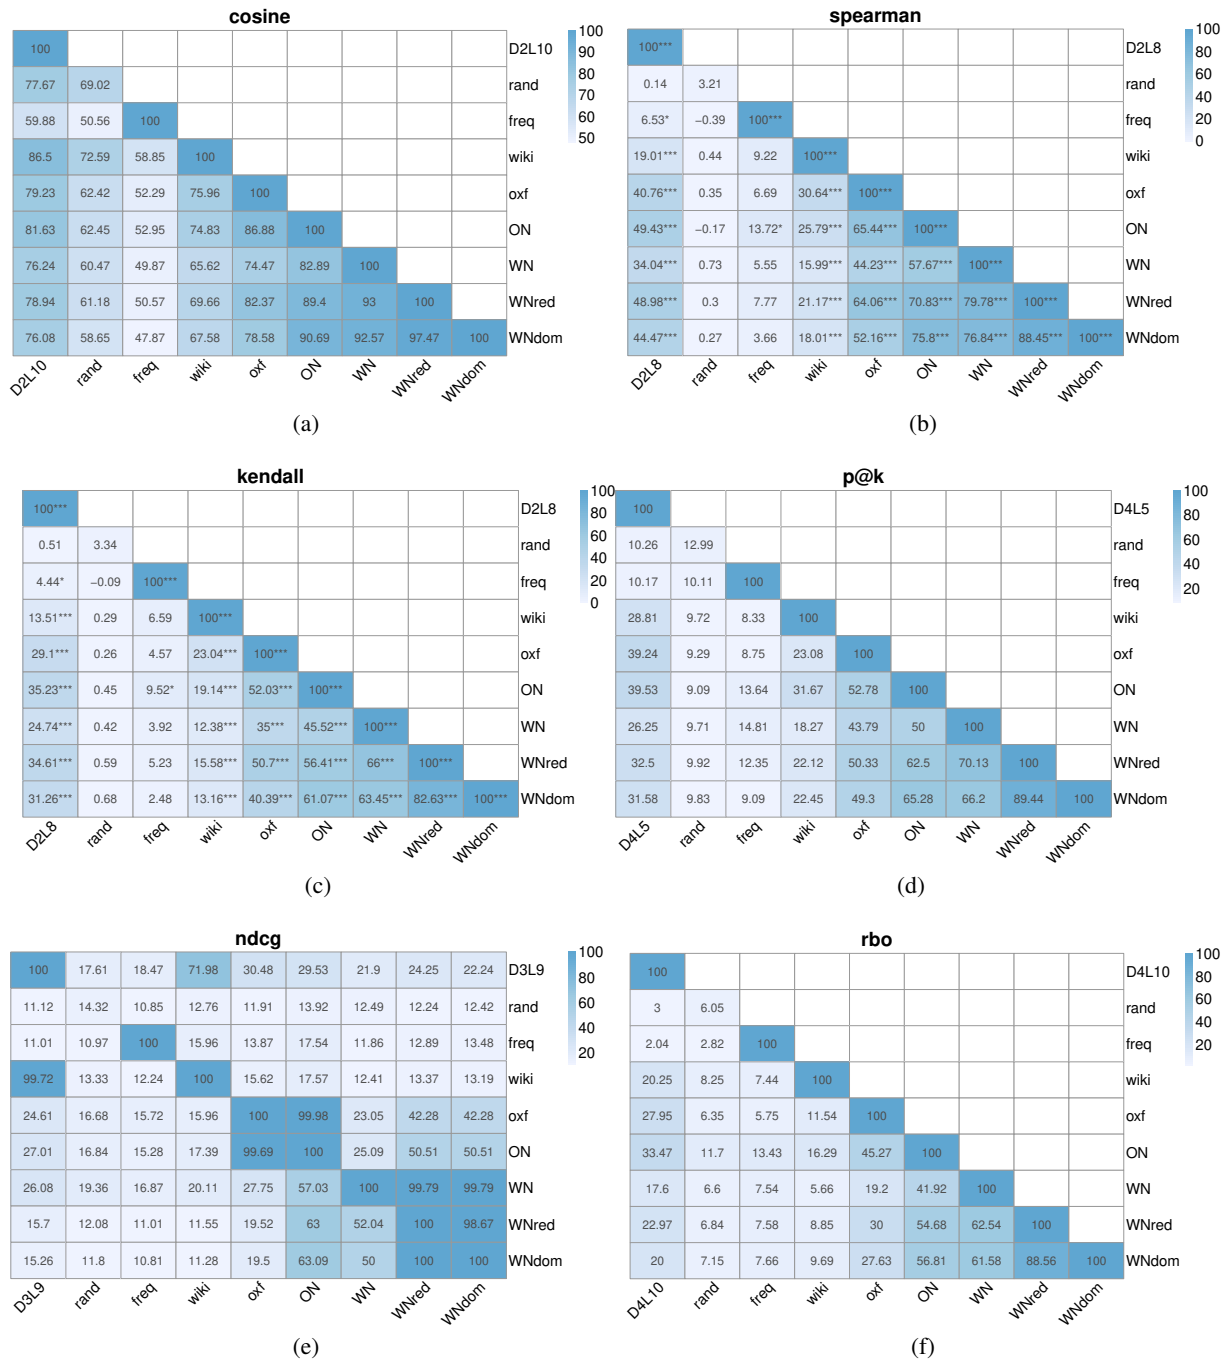


Figure 3: Pairwise similarity matrices between methods. For readability, all scores are shown as percentages. For Kendal and Spearman, * and *** mean statistical significance at $p \leq 0.01$ and $p \leq 0.0001$, respectively. For a given metric, our configuration that best matches (on average) all other methods (except random and frequency) is always shown first. *DDL*L means that the compressed contextual embedding space has D dimensions and that the hierarchy has L levels. Rand, freq, wiki, oxf, ON, WN, WNred, and WNdom are short for random, frequency, Wikipedia, Oxford, OntoNotes, WordNet, WordNet reduced, and WordNet domains. All metrics except NDCG are symmetric, hence we only show one triangle for them. For NDCG, candidate methods are shown as columns and ground truths as rows.

Sentences	Bin coordinates
it stars christopher lee as count dracula along with dennis waterman	(3, 5, 1)
the count of the new group is the sum of the separate counts of the two original groups	(4, 1, 3)
the first fight did not count towards the official record	(4, 5, 1)
five year old horatia came to live at merton in may 1805	(2, 5, 2)
it features various amounts of live and backstage footage while touring	(4, 2, 4)
first tax bills were used to pay taxes and to register bank deposits and bank credits	(4, 2, 4)
the ball nest is built on a bank tree stump or cavity	(5, 2, 3)

Table 2: Sentences containing different senses of the same word can be sampled by selecting from different bins.

Keywords	Bin coordinates
also, gas, used, system, protein, blood, new, steel, food, made	(20, 11, 16, 9)
first, new, one, second, later, world, national, olympic, team, games	(19, 13, 15, 13)
album music rock one labour number chart songs metal single	(21, 14, 14, 16)

Table 3: Towards automatic word sense induction: top 10 most frequent words for different bins containing the word *metal*. The bins correspond to 3 senses of metal: chemical element, Olympics (medals), and music.

$D2L8$, $D4L5$, $D3L9$, and $D4L10$ ¹¹. Given the rather large grid we explored ($[2, 20] \times [2, 19]$ for D and L , resp.), with 342 combinations in total, we can say that all these optimal values belong to the same small neighborhood. This interpretation is confirmed by inspecting Fig. 4, where it can clearly be seen that the optimal area of the hyperparameter space is robust to metric selection and consistently corresponds to small values of D (around 3), and values of L at least above 3 or 4, ideally around 8. For larger values of L , performance plateaus (keeping D fixed). In other words, it is necessary to have some levels in the hierarchy, but having very deep hierarchies is not required for our method to work well. A benefit of having such small optimal values of D and L is their affordability, from a computational standpoint.

All rankings derived from WordNet-based resources are highly correlated. It is interesting to note that the rankings generated from OntoNotes, WordNet, WordNet reduced, and WordNet domains, all are highly similar. And this, despite the very different sense granularities they have. This means that despite the apparent differences in these resources, they all tend to produce similar polysemy rankings. The Oxford rankings tend to be part of this high-similarity cluster as well, to a lesser extent.

Frequent words are not the most polysemous. Finally, one last interesting observation is that while the frequency ranking is much better than the random ones, it still is far away from the human rankings. In other words, the frequency of appearance of a word (excluding stopwords, of course) is not as good an indicator of its polysemy as one

could expect. Some words that follow this observation are "number", "population", and "war".

A note on ties. To assess the impact of ties on the reported results, we repeated all of our experiments multiple times with different tie-breaking strategies (e.g., random, alphabetical...). Results do not change: we find the same best parameter combinations, and the differences in the similarity matrices are minimal.

7 Other applications

Sampling diverse examples. An interesting by-product of our discretization strategy is that it can be used to select sentences containing different senses of the same word, as illustrated in Table 2. Provided a mapping, for a given word, between the sentences that were passed to the pre-trained language model and the vectors, we can sample vectors from different bins and retrieve the associated sentences. If the bins are distant enough, the sentences will contain different senses of the word. For instance, in Table 2, we can see that we are able to sample sentences containing three senses of the word *count*: (1) noble title, (2) determining the total number of, and (3) taking into account. This has many useful applications in practice, e.g., in information retrieval, NLG and conversational systems, dataset creation, etc.

Automatic word sense induction. A simple way of capitalizing on our binning strategy to create word sense inventories would consist in (1) selecting distant bins for a given word, and (2) labeling the selected bins with senses. Both steps can be performed automatically. While this will be investigated in future work, we still give, as a proof of concept, an example in Table 3. In this example,

¹¹for RBO, $D4L10$ and $D4L8$ had the same score.




Figure 4: Performance (color scale) vs. number of PCA dimensions (x axis) vs. number of levels in the hierarchy (y axis).

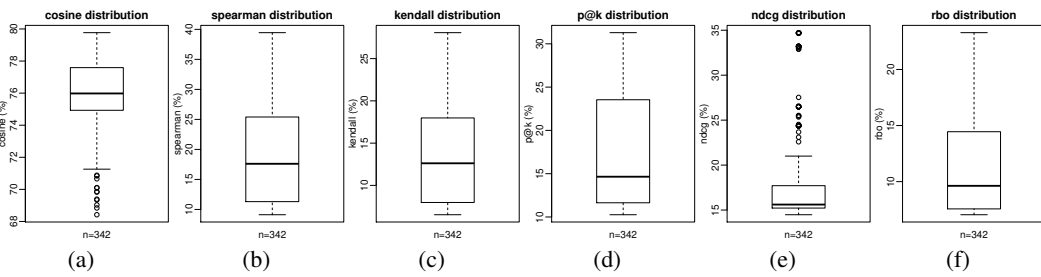


Figure 5: Performance distributions over the 342 values in the discrete hyperparameter space (grids of Fig. 4).

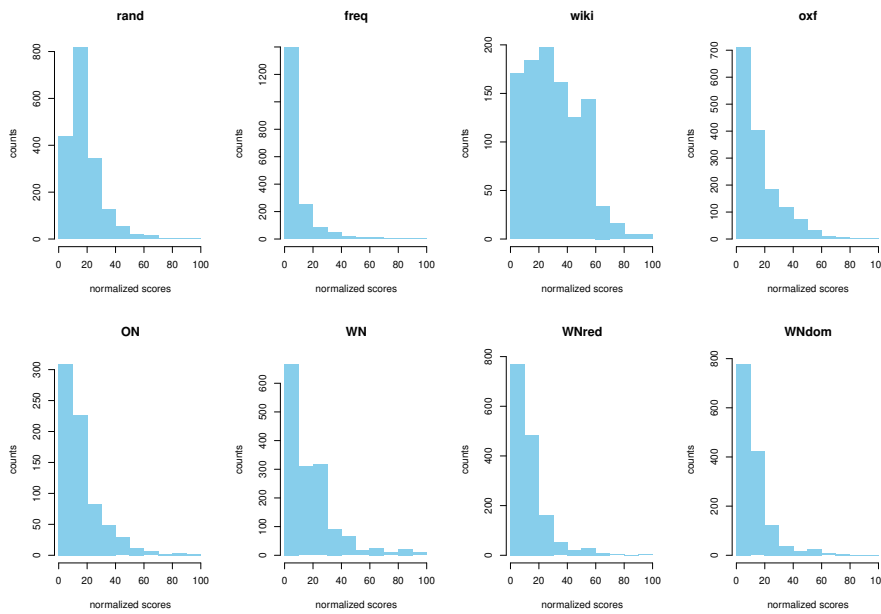


Figure 6: Normalized ranking score distributions for the random and frequency rankings and the human-derived ground truth rankings.

keywords are extracted from distant bins containing the word *metal*, and different senses are retrieved.

8 Related work

Task. Several previous efforts have interested themselves in creating sense inventories without human experts. As an example, in Rumshisky (2011); Rumshisky et al. (2012)¹², Amazon Mechanical Turk (AMT) workers are given a set of sentences containing the target word and one sentence that is randomly selected from this set as a target sentence. Workers are then asked to judge, for each sentence, whether the target word is used in the same way as in the target sentence. This creates an undirected graph of sentences where clustering can be applied to find senses. To label clusters with senses, one has to inspect the sentences in each cluster manually.

More recently, Jurgens (2013)¹³ compared three annotation methodologies for gathering word sense labels on AMT. The methods compared are Likert scales, two-stage select and rate, and the difference between counts of when senses were rated best/worst. Regardless of the strategy, inter-annotator agreement remains low (around 0.3).

Methodology. In the original ELMo paper, Peters et al. (2018) have shown that using contextual word representations (through nearest-neighbor matching) improves word sense disambiguation. Hadwinoto et al. (2019); Coenen et al. (2019) showed that this technique works well with BERT too. Pasini et al. (2020) uses a combination of BERT embeddings and a knowledge-based WSD model to generate word sense distributions, while Giulianelli et al. (2020) uses clustering over the embeddings to detect semantic shifts.

Our approach is also related in spirit to pyramid matching (Nikolentzos et al., 2017; Grauman and Darrell, 2007; Lazebnik et al., 2006). This kernel-based method originated in computer vision. It computes the similarity between objects by placing a sequence of increasingly coarser grids over the feature space and taking a weighted sum of the number of matches occurring at each level. Matches found at finer resolutions are weighted more than matches found at coarser resolutions.

¹²We asked the authors to share annotations with us to use as ground truth, but they were unable to do so.

¹³same as footnote 12.

9 Conclusion

We proposed a novel unsupervised, fully data-driven geometrical approach to estimate word polysemy. Our approach builds multiresolution grids in the contextual embedding space. Through rigorous experiments, we showed that our rankings are well correlated (with strong statistical significance) to 6 different human rankings, for 6 different metrics. Such fully data-driven rankings of words according to polysemy can help in creating new sense inventories, but also in validating and interpreting existing ones. Increasing the quality and consistency of sense inventories is a key first step of the word sense disambiguation pipeline. We also showed that our discretization could be used, at no extra cost, to sample contexts containing different senses of a given word, which has useful applications in practice. Finally, the unsupervised nature of our method makes it applicable to any language.

While our scores are a good proxy for polysemy, they are not equal to word sense counts. Moreover, we do not label each sense. Future work should address these challenges by, e.g., automatically selecting bins of interest and generating labels for them (see section 7).

Future work should also perform some sort of extrinsic evaluation. For instance, the Word-in-Context task (Pilehvar and Camacho-Collados, 2018) could be used, where two occurrences would be classified as having the same meaning if their two vectors fall in the same bin.

Another direction is investigating how different contextual embeddings (e.g., BERT, BART) impact our rankings, including in languages other than English (Eddine et al., 2020; Cañete et al., 2020), and low-resource languages.

Acknowledgments

We thank the three anonymous reviewers for their helpful feedback. We also thank Giannis Nikolentzos for helpful discussions about pyramid matching. The GPU used in this study was donated by the Nvidia corporation as part of their GPU grant program. This work was supported by the LinTo project.

References

Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.

- Luisa Bentivogli, Pamela Forner, Bernardo Magnini, and Emanuele Pianta. 2004. Revising the wordnet domains hierarchy: semantics, coverage and balancing. In *Proceedings of the workshop on multilingual linguistic resources*, pages 94–101.
- Susan Windisch Brown, Travis Rood, and Martha Palmer. 2010. Number or nuance: Which factors restrict reliable word sense annotation? In *LREC*.
- José Cañete, Gabriel Chaperon, Rodrigo Fuentes, Joun-Hui Ho, Hojin Kang, and Jorge Pérez. 2020. Spanish pre-trained bert model and evaluation data. In *PMLADC at ICLR 2020*.
- Andy Coenen, Emily Reif, Ann Yuan, Been Kim, Adam Pearce, Fernanda Viégas, and Martin Wattenberg. 2019. Visualizing and measuring the geometry of bert. *arXiv preprint arXiv:1906.02715*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Moussa Kamal Eddine, Antoine J-P Tixier, and Michalis Vazirgiannis. 2020. Barthez: a skilled pre-trained french sequence-to-sequence model. *arXiv preprint arXiv:2010.12321*.
- Kawin Ethayarajh. 2019. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. *arXiv preprint arXiv:1909.00512*.
- Mario Giulianelli, Marco Del Tredici, and Raquel Fernández. 2020. [Analysing lexical semantic change with contextualised word representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3960–3973, Online. Association for Computational Linguistics.
- Kristen Grauman and Trevor Darrell. 2007. The pyramid match kernel: Efficient learning with sets of features. *Journal of Machine Learning Research*, 8(Apr):725–760.
- Christian Hadiwinoto, Hwee Tou Ng, and Wee Chung Gan. 2019. Improved word sense disambiguation using pre-trained contextualized word representations. *arXiv preprint arXiv:1910.00194*.
- Eduard Hovy, Mitch Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. Ontonotes: the 90% solution. In *Proceedings of the human language technology conference of the NAACL, Companion Volume: Short Papers*, pages 57–60.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.
- Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):422–446.
- David Jurgens. 2013. [Embracing ambiguity: A comparison of annotation methodologies for crowdsourcing word sense labels](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 556–562.
- Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.
- Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. 2006. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2169–2178. IEEE.
- Bernardo Magnini and Gabriela Cavaglia. 2000. Integrating subject field codes into wordnet. In *LREC*, pages 1413–1418.
- Xiangrui Meng, Joseph Bradley, Burak Yavuz, Evan Sparks, Shivaram Venkataraman, Davies Liu, Jeremy Freeman, DB Tsai, Manish Amde, Sean Owen, et al. 2016. Mllib: Machine learning in apache spark. *The Journal of Machine Learning Research*, 17(1):1235–1241.
- George A Miller. 1998. *WordNet: An electronic lexical database*. MIT press.
- Giannis Nikolentzos, Polykarpos Meladianos, and Michalis Vazirgiannis. 2017. Matching node embeddings for graph similarity. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Martha Palmer, Olga Babko-Malaya, and Hoa Trang Dang. 2004. Different sense granularities for different applications. In *Proceedings of the 2nd International Workshop on Scalable Natural Language Understanding (ScaNaLU 2004) at HLT-NAACL 2004*, pages 49–56.
- Martha Palmer, Hoa Trang Dang, and Christiane Fellbaum. 2007. Making fine-grained and coarse-grained sense distinctions, both manually and automatically. *Natural Language Engineering*, 13(2):137–163.
- Tommaso Pasini, Federico Scozzafava, and Bianca Scarlini. 2020. Clubert: A cluster-based approach for learning sense distributions in multiple languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4008–4018.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2018. Wic: the word-in-context dataset for evaluating context-sensitive meaning representations. *arXiv preprint arXiv:1808.09121*.

- R Core Team. 2018. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *URL* <https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/languageunderstandingpaper.pdf>.
- Anna Rumshisky. 2011. **Crowdsourcing word sense definition**. In *Proceedings of the 5th Linguistic Annotation Workshop*, pages 74–81. Association for Computational Linguistics.
- Anna Rumshisky, Nick Botchan, Sophie Kushkuley, and James Pustejovsky. 2012. Word sense inventories by non-experts. In *LREC*, pages 4055–4059.
- Giannis Siglidis, Giannis Nikolentzos, Stratis Limnios, Christos Giatsidis, Konstantinos Skianis, and Michalis Vazirgiannis. 2018. Grakel: A graph kernel library in python. *arXiv preprint arXiv:1806.02193*.
- Charles Spearman. 1904. The proof and measurement of association between two things.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- William Webber, Alistair Moffat, and Justin Zobel. 2010. A similarity measure for indefinite rankings. *ACM Transactions on Information Systems (TOIS)*, 28(4):1–38.
- Ralph Weischedel, Sameer Pradhan, Lance Ramshaw, Martha Palmer, Nianwen Xue, Mitchell Marcus, Ann Taylor, Craig Greenberg, Eduard Hovy, Robert Belvin, et al. 2011. Ontonotes release 4.0. *LDC2011T03, Philadelphia, Penn.: Linguistic Data Consortium*.