

RAW-C: Relatedness of Ambiguous Words—in Context (A New Lexical Resource for English)

Sean Trott

University of California, San Diego
sttrott@ucsd.edu

Benjamin Bergen

University of California, San Diego
bkbergen@ucsd.edu

Abstract

Most words are *ambiguous*—they convey distinct meanings in different contexts—and even the meanings of unambiguous words are *context-dependent*. Both phenomena present a challenge for NLP. Recently, the advent of contextualized word embeddings has led to success on tasks involving lexical ambiguity, such as Word Sense Disambiguation. However, there are few tasks that directly evaluate how well these embeddings accommodate the continuous, dynamic nature of word meaning—particularly in a way that matches human intuitions. We introduce **RAW-C**, a dataset of graded, human relatedness judgments for 112 ambiguous words in context (with 672 sentence pairs total), as well as human estimates of sense dominance. The average inter-annotator agreement for the relatedness norms (assessed using a leave-one-annotator-out method) was 0.79. We then show that a measure of cosine distance, computed using contextualized embeddings from BERT and ELMo, correlates with human judgments, but that cosine distance also systematically *underestimates* how similar humans find uses of the same sense of a word to be, and systematically *overestimates* how similar humans find uses of different-sense homonyms. Finally, we propose a synthesis between psycholinguistic theories of the mental lexicon and computational models of lexical semantics.

1 Introduction

Words mean different things in different contexts. Sometimes these meanings appear to be distinct, a phenomenon known as *lexical ambiguity*. In English, approximately 7% of wordforms are *homonymous*, i.e., they have multiple, unrelated meanings¹ (e.g., “tree bark” vs. “dog bark”), and as many

¹Dautriche (2015) estimates the average rate of homonymy across languages to be 4%.

as 84% of wordforms are *polysemous*, i.e., they have multiple, related meanings (e.g., “pet chicken” vs. “roast chicken”) (Rodd et al., 2004). But even unambiguous words evoke subtly different interpretations depending on the context of use, i.e., their meanings are dynamic and *context-dependent* (Yee and Thompson-Schill, 2016; Li and Joannis, 2021). While the uses of *runs* in “the boy runs” vs. “the cheetah runs” may not be considered distinct meanings, a human comprehender will likely activate a different mental image when processing each sentence (Elman, 2009).

These facts present a challenge for computational models of lexical semantics. Any downstream task that involves meaning requires models capable of *disambiguating* among the multiple possible meanings of an ambiguous word in a given context. Further, the *graded* nature of human semantic representations can influence how comprehenders construe events and participants in those events (Elman, 2009; Li and Joannis, 2021). In turn, a number of Natural Language Processing (NLP) tasks could benefit from context-sensitive representations that go beyond discrete sense representations and capture the manner in which humans construe events—including sentiment analysis, bias detection, machine translation, and more (Trott et al., 2020). If an eventual goal of NLP is human-like language understanding, models must be equipped with semantic representations that are flexible enough to accommodate the dynamic, context-dependent nature of word meaning—as humans appear to do (Elman, 2009; Li and Joannis, 2021).

Yet a crucial prerequisite to developing better models is *evaluating* those models along the relevant dimensions of performance. Thus, at the minimum, we need metrics that evaluate a model along two critical dimensions:

1. **Disambiguation:** A model’s ability to distinguish between distinct meanings of a word.
2. **Contextual Gradation:** A model’s ability to modulate a given meaning in context, in ways that reflect the continuous nature of human judgments.

A promising development in recent years is the rise of contextualized word embeddings, produced using neural language models such as BERT (Devlin et al., 2018), ELMo (Peters et al., 2018), XLNet (Yang et al., 2019), and more. Advances in these models have yielded improved performance on a number of tasks, including Word Sense Disambiguation (WSD) (Boleda et al., 2019; Loureiro et al., 2020).

WSD satisfies the Disambiguation Criterion above, but not the Contextual Gradation Criterion. In fact, there is still a dearth of metrics for assessing the degree to which contextualized representations match human judgments about the way in which context shapes meaning.

In Section 2, we describe several related datasets that satisfy at least one of these criteria. In Section 3, we introduce and describe the dataset construction process for RAW-C: Relatedness of Ambiguous Words—in Context.² In Section 4, we describe the procedure we followed for collecting human relatedness norms for each sentence pair. In Section 5, we report the results of several analyses that probe how well contextualized embeddings from two neural language models (BERT and ELMo) predict these norms. Finally, in Section 6, we explore possible shortcomings in current models, and propose avenues for future work.

2 Related Work

Most existing datasets fulfill either the **Disambiguation** or the **Contextual Gradation** criterion, but few datasets fulfill both (see Haber and Poesio (2020a) for an exception).

Several datasets contain human relatedness and similarity judgments for *distinct* words in isolation (see Section 2.1). Others are used for Word Sense Disambiguation, and contain ambiguous words in different sentence contexts, along with annotated sense labels (see Section 2.2); as noted in the Introduction, WSD fulfills the Disambiguation Criterion, but not the Contextual Gradation Criterion. Several

²The dataset can be found on GitHub: <https://github.com/seantrott/raw-c>.

recent datasets contain graded relatedness judgments for words in different contexts (see Section 2.3). However, none focus specifically on graded relatedness judgments for *ambiguous* words, controlling both the inflection and part of speech of the target word in question. Finally, one dataset (Haber and Poesio, 2020a) contains similarity judgments for polysemous words in context, but is more limited in size and does not match the sentence frame across the two uses (see Section 2.4).

2.1 De-contextualized Word Similarity and Relatedness

Several datasets contain human judgments of the *similarity* or *relatedness* of (mostly English) word pairs, in isolation (see Taieb et al. (2020) for a review). This includes SimLex-999 (Hill et al., 2015), SimVerb-3500 (Gerz et al., 2016), WordSim-353 (Finkelstein et al., 2001), MTurk-771 (Halawi et al., 2012), MEN (Bruni et al., 2014), and more. These datasets are primarily used for evaluating the quality of static semantic representations, including distributed semantic models such as GloVe (Pennington et al., 2014), as well as representations that use knowledge bases like WordNet (Faruqui and Dyer, 2015).

However, these resources are (by definition, as decontextualized judgments) not directly amenable to evaluating how well a model incorporates *context* into its semantic representation of a given word.

2.2 Word Sense Disambiguation

In Word Sense Disambiguation (WSD), a classifier predicts the “sense” of an ambiguous word in a given context, often using a contextualized embedding. WSD relies on annotated sense labels, which in turn requires determining whether any given pair of word uses belong to the same or distinct senses—i.e., whether to “lump” or “split”. There is considerable debate about how *granular* word sense inventories should be (Hanks, 2000; Brown, 2008a);³ resources range in granularity from WordNet (Fellbaum, 1998) to the Coarse Sense Inventory, or CSI (Lacerra et al., 2020). Recent work using coarse-grained sense inventories has achieved success rates of 85% and beyond (Lacerra et al.,

³This also raises deeper philosophical issues about exactly what qualifies as a “sense” (Hanks, 2000; Tuggy, 1993; Geeraerts, 1993; Kilgarriff, 2007); answering these questions is beyond the scope of this paper, though see Section 6 for a brief discussion.

2020; Loureiro et al., 2020).

In terms of the criteria listed above, WSD satisfies the Disambiguation Criterion, but not the Contextual Gradation Criterion. WSD only captures a model’s ability to distinguish between distinct senses; it does not assess how meaning is modulated within a given sense category, e.g., that a human comprehender might consider the meaning of *runs* in “the cheetah runs” as more similar to “the jaguar runs” than to “the toddler runs”, or that some uses of a sense might be more prototypical than others.

2.3 Contextualized Word Similarity and Relatedness

There have been several recent efforts to address this gap in the literature:

The Stanford Contextual Word Similarity (SCWS) dataset (Huang et al., 2012) contains similarity judgments for 2,003 English word pairs in a sentence context. Approximately 12% of the pairs contain the same word (e.g., “pack his bags” vs. “pack of zombies”), though not always in the same part of speech; in most cases, the words compared are different (e.g., “left” vs. “abandon”). This dataset is a useful step towards contextualized similarity judgments, but because most pairs contain different words (or the same word in different parts of speech), static word embeddings such as Word2Vec can still perform quite well without considering the context at all (Pilehvar and Camacho-Collados, 2018).

The Word in Context (WiC) dataset (Pilehvar and Camacho-Collados, 2018) contains a set of over 7,000 sentence pairs with an overlapping English word, labeled according to the use of that word corresponds to same or different senses. As Pilehvar and Camacho-Collados (2018) note, the structure of the dataset requires some form of contextualized meaning representation to perform above a random baseline, which makes it more suitable for interrogating contextualized embeddings. However, the task is a binary classification task along the lines of WSD, making it harder to assess the Contextual Gradation Criterion.

The CoSimLex dataset (Armendariz et al., 2020), created with the Graded Word Similarity in Context (GWSC) task, contains graded similarity judgments for a number of word pairs across English (340), Croatian (112), Slovene (111), and Finnish (24). Each pair of words is rated in two sepa-

rate contexts, yielding 1174 scores in total. Sentence contexts were extracted from each language’s Wikipedia. Unlike WiC, the word pairs do not actually contain the same word—rather, for any given word pair (e.g., “beach” and “seashore”), there are at least two pairs of sentence contexts with associated similarity judgments. Thus, this dataset can be used to assess graded differences in contextualized meaning representations, but not directly for the *same* ambiguous word.

2.4 Contextualized Similarity of Ambiguous Words

Finally, one dataset (Haber and Poesio, 2020a,b) contains graded similarity judgments (as well as co-predication acceptability judgments) for a number of polysemous words in distinct sentential contexts, meeting both **Contextual Gradation** and the **Disambiguation** criteria.

The main limitations of this dataset are its size (it contains examples for only 10 polysemes), as well as the fact that the sentence frames are also not always controlled for each polysemous word.

2.5 Summary

Most datasets reviewed above allow practitioners to evaluate models on their ability to disambiguate (i.e., the Disambiguation Criterion) or their ability to capture graded differences in word relatedness (i.e., the Contextual Gradation Criterion); one dataset (Haber and Poesio, 2020a,b) meets both criteria.

But to our knowledge, no datasets contain graded relatedness judgments for ambiguous words in tightly controlled sentence contexts, with inflection and part-of-speech controlled across each use. In Section 3 below, we describe the procedure we followed for constructing such a dataset.

3 RAW-C: Relatedness of Ambiguous Words, in Context

Items were adapted from stimuli used in past psycholinguistic studies, which contrasted behavioral responses to homonymous and polysemous words, either in isolated lexical decision tasks (Klepousniotou and Baum, 2007) or in a disambiguating context (Klepousniotou, 2002; Klepousniotou et al., 2008; Brown, 2008b). We selected 115 words in total. For each ambiguous word (e.g., “bat”), we created four sentences: two each for two distinct meanings of the word. We attempted to match

the sentence frames as closely as possible, in most cases altering only a single word⁴ across the four sentences to disambiguate the intended meaning:

- 1a. He saw a *fruit* bat.
- 1b. He saw a *furry* bat.
- 2a. He saw a *wooden* bat.
- 2b. He saw a *baseball* bat.

We also labeled each word according to whether the two distinct meanings were judged by lexicographers to be Polysemous or Homonymous. Distinguishing homonymy from polysemy is notoriously challenging (Valera, 2020); common tests include determining whether the two meanings share an etymology (polysemy) or not (homonymy), or determining whether the two meanings are conceptually related (polysemy) or not (homonymy). Both tests can be criticized on multiple grounds (Tuggy, 1993; Valera, 2020), and do not always point in the same direction (e.g., etymologically related words sometimes drift apart, resulting in apparent homonymy).

For our annotation, we consulted both the online Merriam-Webster Dictionary (<https://www.merriam-webster.com/>) and the Oxford English Dictionary, or OED (<https://www.oed.com/>), and identified whether each dictionary listed the two meanings in question in separate lexical entries (homonymy), or as different senses under the same lexical entry (polysemy).⁵ For example, both dictionaries list the *animal* and *meat* senses of the word “lamb” as different senses under the same lexical entry, whereas they list the *animal* and *artifact* senses of the word “bat” under different lexical entries. There was one word (“drill”) on which the two dictionaries did not agree; in this case, we labeled the two meanings (“electric drill” vs. “grueling drill”) as homonymy (as per the OED).

There were also three words for which neither dictionary distinguished the two meanings (either in terms of homonymy or polysemy). For example, “best-selling novel” and “thick novel” refer to *cultural* and *physical* artifacts, respectively, but are not listed as distinct senses. Again, this highlights the

⁴There were 13 words for which at least one of the four sentences used a different article (“a” vs. “an”), in addition to having a different disambiguating word.

⁵Our primary goal with this labelling was not to definitively distinguish homonymy from polysemy; as noted above, there is no single, universal criterion for doing so, and different criteria might be more or less relevant for different purposes. It was simply to specify how lexicographers treat the different words, in case that information is useful for users of the resource.

challenge of distinguishing outright *ambiguity* from *context-dependence*; these items were included in the annotation study described below, but were excluded from the final set of norms, thus resulting in 112 target words altogether.⁶ Each word was used in four sentences, for a total of six *sentence pairs* (see Table 1 for more details). 84 of the target words were nouns, and 28 were verbs (note that Part-of-Speech was always held constant *within* each word).

Ambiguity Type	#Words	#Sentence Pairs
Homonymy	38	228
Polysemy	74	444

Table 1: Number of words (and sentence pairs) for each type of ambiguity.

4 Human Annotation

4.1 Participants

81 participants were recruited through UC San Diego’s undergraduate subject pool for Psychology, Cognitive Science, and Linguistics students. Participants received class credit for participation. Three participants were removed for failing the bot checks at the beginning of the study, and one was removed for failing the catch trials embedded in the experiment, leaving 77 participants in total (59 Female, 16 Male, 2 Non-binary). The median age of participants was 20 (M = 20.22, SD = 2.7), with ages ranging from 18 to 38. 74 participants self-reported as being native speakers of English.

4.2 Materials

We used the original set of 115 words described in Section 3, i.e., including the three items labeled “Unsure”. Each word had four sentences; accounting for order, this resulted in twelve possible sentence pairs (six pairs, with two orders each) for each word, for a total 1380 items.

⁶The existence of these “Unsure” items, as well as items for which the two dictionaries disagreed on the issue of homonymy vs. polysemy, raises the question of whether empirical measurements such as relatedness judgments (or even cosine distance) could help inform lexicographic decisions. As a proof of concept, we trained a logistic regression classifier (using leave-one-out cross-validation) to predict whether two contexts of use belonged to the Same Sense, using Mean Relatedness. The classifier successfully categorized 86.76% of held-out test items as belonging to the same or different senses. Further, for different sense items only, a trained classifier successfully categorized 79% of held-out test items as polysemous or homonymous. While only a proof of concept, this demonstration suggests a promising avenue for future research.

4.3 Procedure

After giving consent, participants answered two questions designed to filter out bots (e.g., “Which of the following is not a place to swim?”, with the correct answer being “Chair”). They were then given instructions, which included a description of how the meaning of a word can change in different contexts.

On each page of the study, participants were shown a pair of sentences, with the target word bolded (see Figure 1 for an example). They were asked to indicate how **related** the uses of that word were across the two sentences, with a labeled Likert scale ranging from “totally unrelated” to “same meaning”.

It was a hostile **atmosphere**.

It was a gaseous **atmosphere**.

How **related** are the uses of this word across these two sentences?

○ totally unrelated ○ not very related ○ somewhat related ○ very related ○ same meaning

Continue

Figure 1: Example item from study.

We included two “catch” trials in the study to identify participants who did not pay attention. In one, the two sentences were identical, such that the correct answer is “same meaning”; the other featured a homonym with two different parts of speech (*rose.v* and *rose.n*), such that the correct answer was “totally unrelated”.

Excluding the catch trials, participants saw 115 sentence pairs total; no word was repeated twice across trials for the same participant. The comparisons any given subject saw for a given word were randomly sampled from the 12 possible sentence pairs, and the order of trials was randomized.⁷

5 Analysis and Results

The analyses run below were performed on the 112 target words (i.e., excluding the “Unsure” items).

⁷Based on the suggestion of an anonymous reviewer, we also ran a follow-up norming study to collect estimates of sense frequency bias (sometimes called *dominance*); sense dominance is known to play an important role in the processing of ambiguous words (Klepousniotou and Baum, 2007; Rayner et al., 1994; Binder and Rayner, 1998; Leininger and Rayner, 2013). These dominance norms are included in the final dataset.

Human annotations were assigned to a scale from 0 (“totally unrelated”) to 4 (“same meaning”).

5.1 Analysis of Sentence Pairs

Before analyzing the responses of human annotators, we first sought to characterize how well two neural language models captured the *categorical* structure in the dataset—i.e., whether their contextualized representations could be used to distinguish same-sense from different-sense uses of the same word, as well as words labelled as different-sense Homonyms from different-sense Polysemes.

We ran every sentence through two language models: ELMo, using the Python AllenNLP package (Gardner et al., 2017), and BERT, using the `bert-embedding` package.⁸ Then, for each sentence pair, we computed the Cosine Distance between the contextualized representations of the target wordform (e.g., *bat* in “He saw the furry bat” and “He saw the wooden bat”). The distribution of Cosine Distances is visualized in Figure 2.

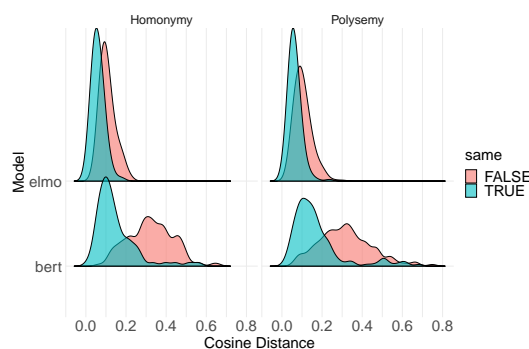


Figure 2: Cosine Distances between the target word’s contextualized embeddings for both language models, plotted by Same Sense (True vs. False) and Ambiguity Type (Homonymy vs. Polysemy).

We also performed several statistical analyses, using the `lme4` package in R (Bates et al., 2015). In each case, we compared a full model to a reduced model using a log-likelihood ratio test. All models had Cosine Distance as a dependent variable, and included Part-of-Speech as a fixed effect, random intercepts for Word and Language Model (i.e., ELMo vs. BERT), and by-Word random slopes for the effect of Same Sense.

Adding a fixed effect of Same Sense significantly improved model fit [$\chi^2(1) = 143.72, p < .001$], with same-sense uses significantly closer than different-sense uses [$\beta = -.099, SE = 0.005$].

⁸<https://pypi.org/project/bert-embedding/>

However, adding an interaction between Same Sense and Ambiguity Type (as well as fixed effects of both) did not significantly improve the fit above a model omitting the interaction [$\chi^2(1) = 2.19, p = 0.14$]. In other words, both language models could differentiate same-sense and different-sense uses of an ambiguous word, but their ability to discriminate between Homonymy and Polysemy was marginal at best.

5.2 Analysis of Human Annotations

Our primary goal was understanding the distribution of human relatedness annotations—both in terms of how it reflects the underlying categorical structure of the dataset (e.g., Homonymy vs. Polysemy), as well as the Cosine Distance measures from each language model. As in the section above, we constructed a series of linear mixed effects models and performed log-likelihood ratio tests for each model comparison; in each case, the dependent variable was Relatedness. All models included a fixed effect of Part-of-Speech, by-subject and by-word random slopes for the effect of Same Sense, by-subject random slopes for the effect of Ambiguity Type, and random intercepts for subjects and items.

First, we asked whether participants’ relatedness judgments varied across same-sense and different-sense sentence pairs. We added a fixed effect of Same Sense to the base model described above, along with fixed effects for the Cosine Distance measures from BERT and ELMo. This model explained significantly more variance than a model omitting only Same Sense [$\chi^2(1) = 207.11, p < .001$], with same-sense uses receiving higher relatedness judgments on average [$\beta = 1.94, SE = 0.1$]. The median relatedness judgment for same-sense uses was 4 ($M = 3.46, SD = 1.02$), while the median relatedness judgment for different-sense uses was 1 ($M = 1.31, SD = 1.45$). Second, we asked whether participants’ judgments were sensitive to the distinction between Homonymy and Polysemy. We added an interaction between Same Sense and Ambiguity Type (along with a fixed effect of Ambiguity Type) to the model described above. The interaction significantly improved model fit [$\chi^2(1) = 25.45, p < .001$]. The median relatedness for both same-sense homonyms and polysemes was 4, whereas the median relatedness for different-sense homonyms (0) was lower than that for different-sense polysemes

(2). Further, as depicted in Figure 3, there was considerably more variance across polysemous words than homonymous words—this makes sense, given that some polysemous meanings are highly related (e.g., “pet chicken” vs. “roast chicken”), while others are more distant (e.g., “desperate act” vs. “magic act”).

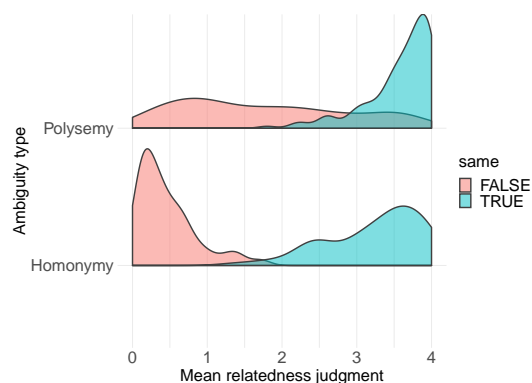


Figure 3: Mean relatedness judgments for each sentence pair, plotted by by Same Sense (True vs. False) and Ambiguity Type (Homonymy vs. Polysemy).

Third, we asked whether the Cosine Distance measures explained independent variance above and beyond that explained by Same Sense and Ambiguity Type. A full model including all factors explained more variance than a model excluding only the Cosine Distance measure from BERT [$\chi^2(1) = 36.19, p < .001$], as well as a model excluding only the Cosine Distance measure from ELMo [$\chi^2(1) = 16.92, p < .001$]. This indicates that Relatedness does not vary purely as a function of the categorical structure in the dataset—the graded relatedness judgments were sensitive to subtle differences in context.

5.3 Inter-Annotator Agreement

Inter-annotator agreement was assessed by calculating the average Spearman’s rank correlation between each participant’s responses and the Mean Relatedness for the set of 112 items observed by that participant—where Mean Relatedness was calculated after omitting responses by the participant in question. This answers the question: to what extent do each participant’s responses correlate with the consensus rating by the 76 other participants? Using this method, the average correlation was $\rho = 0.79$, with a median of $\rho = 0.81 (SD = .07)$. The lowest agreement was $\rho = 0.55$, and the highest was $\rho = 0.88$.

5.4 Evaluation of Language Models

To evaluate the language models, we collapsed across the single-trial data and computed the Mean and Median Relatedness for each unique sentence pair. The distribution of Mean Relatedness judgments is depicted in Figure 3.

As in past work (Hill et al., 2015), we computed the Spearman’s rank correlation between the distribution of Cosine Distance measures (from each model) and the Mean Relatedness for a given sentence pair. BERT performed slightly better than ELMo ($BERT_\rho = -0.58$, $ELMo_\rho = -0.53$).⁹ Putting this in context, both models performed considerably worse than the average inter-annotator agreement score ($\rho = 0.79$).

We also computed the R^2 of a linear regression including the Cosine Distance measures from both BERT and ELMo. Combined, both measures explained roughly 37% of the variance in Mean Relatedness judgments ($R^2 = 0.37$). Surprisingly, this was only slightly more than half the variance explained by a linear regression including only the interaction between Same Sense and Ambiguity Type ($R^2 = 0.66$), as well as a regression including all factors ($R^2 = 0.71$).

By visualizing the residuals from the linear regression with only BERT and ELMo (see Figure 4), we see that Cosine Distance appears to systematically *underestimate* how related participants find same-sense judgments to be (for both Polysemy and Homonymy). Further, we see that Cosine Distance systematically *overestimates* how related participants find different-sense Homonyms to be.

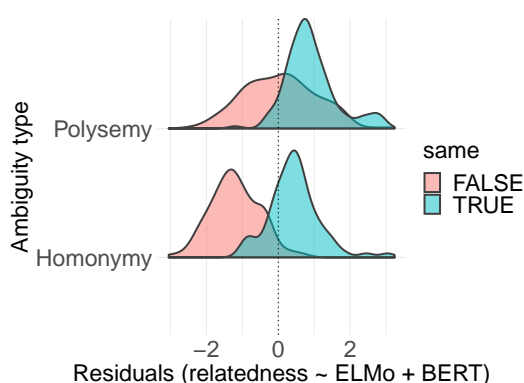


Figure 4: Residuals of a linear regression including Cosine Distance measures from both BERT and ELMo, plotted by Same Sense (True vs. False) and Ambiguity Type (Homonymy vs. Polysemy).

⁹Note that larger values of Cosine Distance indicate a larger distance between two vectors; thus, a negative correlation is expected between relatedness and Cosine Distance.

6 Discussion

Word meanings are dynamic, dependent on the contexts in which those words appear—and some words are even ambiguous, generating distinct, incompatible interpretations in different situations (e.g., “fruit bat” vs. “baseball bat”).

RAW-C contains graded relatedness judgments (by human annotators) for ambiguous English words in distinct sentential contexts. Importantly, the ambiguous wordform (e.g., “bat”) is always matched for both part-of-speech and inflection across each sentence pair; 84 of the target words are nouns, and 28 are verbs. Each word has relatedness judgments for six different sentences pairs (four unique sentences): two same-sense pairs, and four different-sense pairs. Same sense pairs convey the same meaning, according to Merriam-Webster and the OED (e.g., “fruit bat” and “furry bat”), while different sense pairs correspond to meanings listed in either distinct lexical entries (e.g., “fruit bat” and “wooden bat”) or distinct sub-entries (e.g., “marinated lamb” and “baby lamb”). Furthermore, different-sense pairs are labeled according to whether they are related via homonymy or polysemy, a relevant distinction for both lexicographers and psycholinguists—recent evidence suggests that polysemous and homonymous meanings are represented differently in the mental lexicon (Klepousniotou, 2002; Klepousniotou and Baum, 2007). Finally, the sentential context is always tightly controlled; in most pairs, only one word differs across the two sentences (e.g., “fruit” vs. “furry”).

In Section 5, we reported several primary findings. First, contextualized representations from both BERT and ELMo capture the distinction between same-sense and different-sense uses of a word, but their ability to distinguish between homonymy and polysemy is marginal at best. This contrasts with other recent work (Nair et al., 2020), suggesting that BERT is able to differentiate between homonymy and polysemy. One possible explanation for this difference in results is that Nair et al. (2020) used naturally-occurring sentences from Semcor (Miller et al., 1993), whereas our sentence contexts were more tightly controlled. Our results indicate that even the presence of a single disambiguating word can trigger nuanced differences in semantic representation in humans, but not necessarily in current neural language models.

Second, we found that both BERT and ELMo

explain independent sources of variance in human relatedness judgments, above and beyond Same Sense and Ambiguity Type (i.e., homonymy vs. polysemy). This is encouraging, because it demonstrates a direct benefit of graded (rather than categorical) judgments; for example, among the broad category of different-sense polysemous pairs, some are closely related (e.g., “marinated lamb” and “baby lamb”), and others are considerably less closely related (e.g., “hostile atmosphere” and “gaseous atmosphere”). Overall, contextualized embeddings from BERT were better at predicting human relatedness judgments than those from ELMo—this is consistent with past work (Wiedemann et al., 2019) suggesting that BERT outperforms ELMo on tasks involving sense disambiguation.

Importantly, however, both BERT and ELMo failed to capture variance in relatedness judgments that is captured by Same Sense and Ambiguity Type. As depicted in Figure 4, Cosine Distance tended to *underestimate* how related humans find same-sense uses to be, and *overestimate* how related humans find different-senses to be. This is not entirely surprising, given that neither BERT nor ELMo are equipped with discrete sense representations—at most, they produce contextualized embeddings that are amenable to supervised classification or unsupervised clustering. Yet this also illustrates that—at least on this task—humans *do* appear to draw on some manner of (likely fuzzy) categorical representation, such that the difference between two contexts of use is *compressed* for same-sense meanings, and *exaggerated* for different-sense meanings (particularly for homonyms). This suggests several exciting avenues for future work: can neural language models such as BERT be augmented with semantic knowledge or representational schemes that improve their performance on RAW-C or similar tasks? Both possibilities are explored in Section 6.1 below.

6.1 Future Work

As Bender and Koller (2020) note, most language models are trained on linguistic form alone. In contrast, human language knowledge is *grounded* in our embodied experience of the world (Bisk et al., 2020). To the extent that human sense representations are guided by distinct sensorimotor or social-interactional associations, equipping language models with this information ought to fa-

cilitate their ability to distinguish between distinct meanings of a word (i.e., the Disambiguation Criterion) and modulate a given meaning in context (i.e., the Contextual Gradation Criterion).

Practitioners could also look to (and in turn, inform) models of the human mental lexicon (Nair et al., 2020). Several promising models attempt to address the *continuous* nature of word meaning, as well as the issue of apparent category boundaries (i.e., word senses) (Rodd et al., 2004; Elman, 2009); at this stage, the role of continuity vs. categorical structure in human sense representations remains an open question. Models such as SenseBERT (Levine et al., 2020) incorporate high-level sense knowledge into internal representations from the beginning, and find improvements on several WSD tasks—would this approach, or others like it, yield an improvement on RAW-C as well?

6.2 Limitations of Dataset

RAW-C has multiple limitations, some of which could also be addressed in future work. First, the broad category of “polysemy” is often subdivided into different mechanisms or manners of conceptual relation, such as metaphor and metonymy. This distinction is also believed to be cognitively relevant, with some evidence that metaphorically related senses are represented differently than metonymically related ones (Klepousniotou, 2002; Klepousniotou and Baum, 2007; Lopukhina et al., 2018; Yurchenko et al., 2020). Future work could annotate polysemous word pairs for whether they are related by metaphor, metonymy, or another class of semantic relation—annotations could even be made as granular as the specific semantic relation involved (e.g., Animal for Meat) (Srinivasan and Rabagliati, 2015). This finer-grained coding could be used to assess exactly which kinds of semantic relation correlate with the distributional profile of word tokens—i.e., are accessible from linguistic form alone—and which require some external module, whether in the form of grounded world knowledge or a structured knowledge base.

Another possible limitation is the fact that RAW-C contains experimentally controlled minimal pairs, instead of naturally-occurring sentences (Nair et al., 2020; Haber and Poesio, 2020a,b). On the one hand, naturalistic sentences are useful for evaluating models on WSD “in the wild” (and indeed, there are a number of useful datasets for this purpose; see Section 2). On the other hand, controlled

datasets are useful if one’s goal is to better understand a particular model or linguistic phenomenon—especially if this also allows a direct comparison with human annotations. For example, our analyses suggest that human sense representations must involve some additional levels of processing or information beyond the statistical regularities in word co-occurrence captured by BERT and ELMo. Moving forward, we hope that experimentally controlled datasets such as RAW-C will serve as a useful complement to existing, more naturalistic datasets.

7 Conclusion

We have presented a novel dataset for evaluating contextualized language models: RAW-C (Relatedness of Ambiguous Words, in Context). This resource contains both categorical annotations, derived from expert lexicographers (Merriam-Webster and the OED), as well as graded relatedness judgments from human participants. We found that contextualized representations from BERT and ELMo captured some variance ($R^2 = .37$) in these relatedness judgments, but that the distinction between same-sense and different-sense uses, as well as between homonymy and polysemy, explains considerably more ($R^2 = .66$). Finally, we argued that this gap in performance represents an exciting opportunity for further development, and for cross-pollination between experimental psycholinguistics and NLP.

8 Ethical Considerations

All responses from human participants were anonymized before analyzing any data. Furthermore, the RAW-C dataset does not contain single-trial data—responses for a given sentence pair have been collapsed across all the human annotators who provided a rating for that pair. All participants provided informed consent, and were compensated in the form of SONA credits (to be applied to various Psychology, Cognitive Science, or Linguistics classes). The project was carried out with IRB approval.

Acknowledgments

We are grateful to Susan Windisch Brown and Eka-terini Klepousniotou for making their experimental stimuli available. We also thank the anonymous reviewers for their helpful suggestions, and Nathan Schneider for early feedback on the idea to publish

the dataset. Finally, we are grateful to other members of the Language and Cognition Lab (James Michaelov, Cameron Jones, and Tyler Chang) for valuable comments and discussion.

References

- Carlos Santos Armendariz, Matthew Purver, Senja Polak, Nikola Ljubešić, Matej Ulčar, Ivan Vulić, and Mohammad Taher Pilehvar. 2020. [SemEval-2020 Task 3: Graded word similarity in context](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 36–49.
- Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. [Fitting linear mixed-effects models using lme4](#). *Journal of Statistical Software*, 67(1):1–48.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198.
- Katherine S Binder and Keith Rayner. 1998. [Contextual strength does not modulate the subordinate bias effect: Evidence from eye fixations and self-paced reading](#). *Psychonomic Bulletin & Review*, 5(2):271–276.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, et al. 2020. [Experience grounds language](#). *arXiv preprint arXiv:2004.10151*.
- Gemma Boleda, Kristina Gulordava, and Laura Aina. 2019. [Putting words in context: LSTM language models and lexical ambiguity](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics; 2019 Jul 28-Aug 2; Florence, Italy. Stroudsburg (PA): ACL; 2019. p. 3342–8*. ACL (Association for Computational Linguistics).
- Susan Windisch Brown. 2008a. [Choosing sense distinctions for WSD: Psycholinguistic evidence](#). In *Proceedings of ACL-08: HLT, Short Papers*, pages 249–252.
- Susan Windisch Brown. 2008b. [Polysemy in the mental lexicon](#). *Colorado Research in Linguistics*, 21.
- Elia Bruni, Nam-Khanh Tran, and Marco Baroni. 2014. [Multimodal distributional semantics](#). *Journal of Artificial Intelligence Research*, 49:1–47.
- Isabelle Dautriche. 2015. [Weaving an ambiguous lexicon](#). Ph.D. thesis, Sorbonne Paris Cité.
- 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*.

- Jeffrey L Elman. 2009. [On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon](#). *Cognitive science*, 33(4):547–582.
- Manaal Faruqui and Chris Dyer. 2015. [Non-distributional word vector representations](#). *arXiv preprint arXiv:1506.05230*.
- Christine Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. Cambridge, MA: MIT Press.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. [Placing search in context: The concept revisited](#). In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Taffjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2017. [Allennlp: A deep semantic natural language processing platform](#).
- Dirk Geeraerts. 1993. [Vagueness’s puzzles, polysemy’s vagaries](#). *Cognitive Linguistics (includes Cognitive Linguistic Bibliography)*, 4(3):223–272.
- Daniela Gerz, Ivan Vulić, Felix Hill, Roi Reichart, and Anna Korhonen. 2016. [Simverb-3500: A large-scale evaluation set of verb similarity](#). *arXiv preprint arXiv:1608.00869*.
- Janosch Haber and Massimo Poesio. 2020a. [Assessing polyseme sense similarity through co-predication acceptability and contextualised embedding distance](#). In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 114–124.
- Janosch Haber and Massimo Poesio. 2020b. [Word sense distance in human similarity judgements and contextualised word embeddings](#). In *Proceedings of the Probability and Meaning Conference (PaM 2020)*, pages 128–145.
- Guy Halawi, Gideon Dror, Evgeniy Gabrilovich, and Yehuda Koren. 2012. [Large-scale learning of word relatedness with constraints](#). In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1406–1414.
- Patrick Hanks. 2000. [Do word meanings exist?](#) *Computers and the Humanities*, 34(1/2):205–215.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. [Simlex-999: Evaluating semantic models with \(genuine\) similarity estimation](#). *Computational Linguistics*, 41(4):665–695.
- Eric H Huang, Richard Socher, Christopher D Manning, and Andrew Y Ng. 2012. [Improving word representations via global context and multiple word prototypes](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 873–882.
- Adam Kilgarriff. 2007. [Word senses](#). In *Word Sense Disambiguation*, pages 29–46. Springer.
- Ekaterini Klepousniotou. 2002. [The processing of lexical ambiguity: Homonymy and polysemy in the mental lexicon](#). *Brain and language*, 81(1-3):205–223.
- Ekaterini Klepousniotou and Shari R Baum. 2007. [Disambiguating the ambiguity advantage effect in word recognition: An advantage for polysemous but not homonymous words](#). *Journal of Neurolinguistics*, 20(1):1–24.
- Ekaterini Klepousniotou, Debra Titone, and Carolina Romero. 2008. [Making sense of word senses: The comprehension of polysemy depends on sense overlap](#). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(6):1534.
- Caterina Lacerra, Michele Bevilacqua, Tommaso Pasini, and Roberto Navigli. 2020. [CSI: A coarse sense inventory for 85% word sense disambiguation](#). In *AAAI*, pages 8123–8130.
- Mallorie Leinenger and Keith Rayner. 2013. [Eye movements while reading biased homographs: Effects of prior encounter and biasing context on reducing the subordinate bias effect](#). *Journal of Cognitive Psychology*, 25(6):665–681.
- Yoav Levine, Barak Lenz, Or Dagan, Ori Ram, Dan Padnos, Or Sharir, Shai Shalev-Shwartz, Amnon Shashua, and Yoav Shoham. 2020. [SenseBERT: Driving some sense into BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4656–4667. Online. Association for Computational Linguistics.
- Jiangtian Li and Marc F Joannis. 2021. [Word senses as clusters of meaning modulations: A computational model of polysemy](#). *Cognitive Science*, 45(4):e12955.
- Anastasiya Lopukhina, Anna Laurinavichyute, Konstantin Lopukhin, and Olga Dragoy. 2018. [The mental representation of polysemy across word classes](#). *Frontiers in psychology*, 9:192.
- Daniel Loureiro, Kiamehr Rezaee, Mohammad Taher Pilehvar, and Jose Camacho-Collados. 2020. [Language models and word sense disambiguation: An overview and analysis](#). *arXiv preprint arXiv:2008.11608*.
- George A Miller, Claudia Leacock, Randee Teng, and Ross T Bunker. 1993. [A semantic concordance](#). In *Human Language Technology: Proceedings of a Workshop Held at Plainsboro, New Jersey, March 21-24, 1993*.
- Sathvik Nair, Mahesh Srinivasan, and Stephan Meylan. 2020. [Contextualized word embeddings encode aspects of human-like word sense knowledge](#). *arXiv preprint arXiv:2010.13057*.

- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- 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*.
- Keith Rayner, Jeremy M Pacht, and Susan A Duffy. 1994. [Effects of prior encounter and global discourse bias on the processing of lexically ambiguous words: Evidence from eye fixations](#). *Journal of memory and language*, 33(4):527–544.
- Jennifer M Rodd, M Gareth Gaskell, and William D Marslen-Wilson. 2004. [Modelling the effects of semantic ambiguity in word recognition](#). *Cognitive science*, 28(1):89–104.
- Mahesh Srinivasan and Hugh Rabagliati. 2015. [How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy](#). *Lingua*, 157:124–152.
- Mohamed Ali Hadj Taieb, Torsten Zesch, and Mohamed Ben Aouicha. 2020. [A survey of semantic relatedness evaluation datasets and procedures](#). *Artificial Intelligence Review*, 53(6):4407–4448.
- Sean Trott, Tiago Timponi Torrent, Nancy Chang, and Nathan Schneider. 2020. [\(Re\) construing Meaning in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*.
- David Tuggy. 1993. [Ambiguity, polysemy, and vagueness](#). *Cognitive linguistics*, 4(3):273–290.
- Salvador Valera. 2020. [Polysemy versus homonymy](#). In *Oxford Research Encyclopedia of Linguistics*.
- Gregor Wiedemann, Steffen Remus, Avi Chawla, and Chris Biemann. 2019. [Does BERT make any sense? Interpretable Word Sense Disambiguation with contextualized embeddings](#). *arXiv preprint arXiv:1909.10430*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in neural information processing systems*, pages 5753–5763.
- Eiling Yee and Sharon L Thompson-Schill. 2016. [Putting concepts into context](#). *Psychonomic bulletin & review*, 23(4):1015–1027.
- Anna Yurchenko, Anastasiya Lopukhina, and Olga Dragoy. 2020. [Metaphor is between metonymy and homonymy: Evidence from event-related potentials](#). *Frontiers in Psychology*, 11:2113.