

# CoSimLex: A Resource for Evaluating Graded Word Similarity in Context

Carlos S. Armendariz\*, Matthew Purver\*<sup>†</sup>,  
Matej Ulčar<sup>‡</sup>, Senja Pollak<sup>†</sup>, Nikola Ljubešić<sup>†</sup>, Marko Robnik-Šikonja<sup>‡</sup>,  
Mark Granroth-Wilding<sup>◇</sup>, Kristiina Vaik<sup>§</sup>

\*Cognitive Science Research Group, Queen Mary University of London, London, UK  
{c.santosarmendariz, m.purver}@qmul.ac.uk

<sup>†</sup>Department of Knowledge Technologies, Jožef Stefan Institute, Ljubljana, Slovenia  
{senja.pollak, nikola.ljubesic}@ijs.si

<sup>‡</sup>University of Ljubljana, Faculty of Computer and Information Science, Slovenia  
{matej.ulcar, marko.robnik}@fri.uni-lj.si

<sup>◇</sup>Department of Computer Science, University of Helsinki, Finland  
mark.granroth-wilding@helsinki.fi

<sup>§</sup>Department of Data Analysis, Texta, Estonia  
kristiina.vaik@ut.ee

## Abstract

State of the art natural language processing tools are built on context-dependent word embeddings, but no direct method for evaluating these representations currently exists. Standard tasks and datasets for intrinsic evaluation of embeddings are based on judgements of similarity, but ignore context; standard tasks for word sense disambiguation take account of context but do not provide continuous measures of meaning similarity. This paper describes an effort to build a new dataset, CoSimLex, intended to fill this gap. Building on the standard pairwise similarity task of SimLex-999, it provides context-dependent similarity measures; covers not only discrete differences in word sense but more subtle, graded changes in meaning; and covers not only a well-resourced language (English) but a number of less-resourced languages. We define the task and evaluation metrics, outline the dataset collection methodology, and describe the status of the dataset so far.

**Keywords:** corpus, annotation, semantics, similarity, context, salience, context-dependence

## 1. Introduction

Recent work in language modelling and word embeddings has led to a sharp increase in use of context-dependent models such as ELMo (Peters et al., 2018) and BERT (Devlin et al., 2019). These models, by providing representations of words which depend on the surrounding context, allow us to take account of the effects not only of discrete differences in word sense but of the more graded effects of context. However, evaluation of these models has generally been in terms of either their performance as language models, or their effect on downstream tasks such as sentiment classification (Peters et al., 2018): there are few resources available which allow evaluation in terms of the properties of the embeddings themselves, or in terms of their ability to model human perceptions of meaning. There are established methods to evaluate word embedding models intrinsically via their ability to reflect human similarity judgements (see e.g. WordSim-353 (Finkelstein et al., 2002) and SimLex-999 (Hill et al., 2015)) or model analogies (Mikolov et al., 2013); however, these have generally ignored context and treated words in isolation. The few that do provide context (e.g. SCWS (Huang et al., 2012) and WiC (Pilehvar and Camacho-Collados, 2019)) focus on word sense and discrete effects, thus missing some of the effects that context has on words in general, and some of the benefits of context-dependent models. To evaluate current models, we need a way to evaluate their ability to reflect similarity judgements *in context*: how well do they model the effects that context has on word meaning?

In this paper we present our ongoing efforts to define and build a new dataset that tries to fill that gap: **CoSimLex** (Armendariz et al., 2020). CoSimLex builds on the familiar pairwise, graded similarity task of SimLex-999, but extends it to pairs of words as they occur in context, and specifically provides two different shared contexts for each pair of words. This will provide a dataset suitable for intrinsic evaluation of state-of-the-art contextual word embedding models, by testing their ability to reflect human judgements of word meaning similarity in context, and crucially, the way in which this varies as context is changed. It goes beyond other existing context-based datasets by taking the *gradedness* of human judgements into account, thus applying not only to polysemous words, or words with distinct senses, but to the phenomenon of context-dependency of word meaning in general. The dataset is also multilingual, and includes three less-resourced European languages: Croatian, Finnish and Slovene. It is to be used as the gold standard for evaluation of a task at SemEval2020: Task 3, Graded Word Similarity in Context.<sup>1</sup>

## 2. Background

From the outset, our main motivation for the development of this dataset came from an interest in the cognitive and psychological mechanisms by which context affects our perception of the meaning of words. There have been many different ways in the literature to look at this phenomenon,

<sup>1</sup><https://competitions.codalab.org/competitions/20905>

which lie in the intersection of several different fields of research, and a detailed discussion of the different approaches to this problem is out of the scope of this paper; here, we present two of the most prominent ideas that helped define what we were trying to capture, and made an impact in the design of the dataset and its annotation process. We then look at previous datasets that deal with similarity in context.

## 2.1. Contextual Modulation

Within the field of lexical semantics, Cruse (1986) proposed an interesting compromise between those linguists that saw words as associated with a number of discrete senses and those that thought that the perceived discreteness of lexical senses is just an illusion. He distinguishes two different manners in which sentential context modifies the meaning of a word. First, the context can select for different discrete senses; if that is the case, the word is described as *ambiguous*, and the process is referred as **contextual selection of senses**. This effect is well known, and is the basis of many word-sense disambiguation tasks.

1. We finally reached the bank.
2. At this point, the bank was covered with brambles.

In example (1), the word *bank* can have the *financial* or *riverbank* sense; and here, the context doesn't really help us select the correct sense. This creates some tension on the part of the reader: we need to select a sense in order for the sentence to properly work, and without this we may feel that the sentence has not been fully understood. This is an example of *ambiguity*. In example (2), in contrast, the context makes one of the senses more *normal* than the other. Cruse (1986) sees the evaluation of *contextual normality* as the main mechanism for sense selection.

The second way in which context can modify the meaning of a word works within the scope of a single sense, modifying it in an unlimited number of ways by *highlighting* certain semantic traits and *backgrounding* others. This process is called **contextual modulation of meaning**, and the word is said to be *general* with respect to the traits that are being modulated. This effect is by nature not discrete but continuous and fluid, and since every word is *general* to some extent: it can be argued that a word has a different meaning in every context in which it appears.

3. Sue is visiting her pregnant cousin.
4. Peter doesn't like his cousin.
5. Arthur poured the butter into a dish.

In example (3), the context tells us that the cousin is female. The meaning of *cousin* is being *modulated* by the context to promote the "female" trait. *Cousin* is a *general* word that includes male and female, but also tall, short, happy and sad cousins. However, as we can see in example (4), the absence of information about these traits doesn't produce the type of tension we saw in (1) above; there is a distinction between meaning modulation and sense selection. The last example (5) is another case of *contextual modulation* in which *poured* highlights the "liquid" trait for *butter*.

It is interesting to notice that in this case not only "liquid" is highlighted, related traits like "warm" can be highlighted as a consequence.

It seems clear that the contextual selection of senses would modify human judgements of similarity. For example, the word *bank*, when used in a context which selects its financial institution sense, should be scored as more similar to other kinds of financial institution (e.g. *building society*) than when in a context which selects the geographic sense of the word. However, we should also expect that a word like *butter*, when contextually modulated to highlight its "liquid", "hot" and "frying" traits, should score more similar to *vegetable oil* than when contextually modulated to highlight its "animal sourced", "dairy", and "creamy" traits. This kind of hypothesis would be testable given a new context-dependent similarity dataset.

Both *sense selection* and *meaning modulation* happen very commonly together, with the same context forcing a sense and then modulating its expression. Many different explanations have been proposed for the emergence of these discrete senses, and some may have their origins in very commonly modulated meaning but, according to Cruse, once a discrete sense is established it becomes something different and follows different rules:

6. John prefers bitches to dogs.
7. John prefers bitches to canines.
8. Mary likes mares better than horses.

Here example (6) works because one of the discrete senses associated to the word *dog* refers only to male dogs. This cannot be explained by *contextual modulation*: if that was the case, example (7), which replaces *dog* with *canine*, should also work, as *canine* could be modulated in the same way that *dog* was; and similarly example (8). However, both seem unnatural at best. The fact that neither *canine* nor *horse* can be modulated in this same way indicates that meaning modulation and sense selection are two, strongly interconnected, but distinctive mechanisms of contextual variability.

A final interesting point about Cruse's view is that he doesn't find the contrast between polysemy and homonymy particularly helpful, and dislikes the use of these terms because they promote the idea that the primary semantic unit is some common lexeme and each of the different senses are just variants of it. He instead believes the primary semantic unit should be the *lexical units*, a union of a single sense and a lexical form, and finds it more useful to look at the contrast between discrete and continuous semantic variability.

## 2.2. Saliency Manipulation

Until now we have looked at contextual variability as an exclusively linguistic phenomenon, a point of view rooted in lexical semantics. We looked at how the context of the sentence affects the meaning of the word. In contrast, cognitive linguistics, and the more specific cognitive semantics, look at language and meaning as a more general expression of human cognition (Evans and Green, 2018).

This approach champions concepts, more specifically *conceptual structures*, as the true recipient of meaning, replacing words or lexical units. These linguistic units no longer refer to objects in an external world but to concepts in the mind of the speaker. Words get their meaning only by association with *conceptual structures* in our minds. The process by which we construct meaning is called conceptualisation, an embodied phenomenon based in social interaction and sensory experience.

Cognitive linguists gravitate to themes that focus on the flexibility and the ability of the interaction between language and conceptual structures to model continuous phenomena, like prototyping effects, categories, metaphor theory and new ways to look at polysemy. Within the cognitive tradition, the idea of *conceptual spaces*, characterised by *conceptual dimensions*, has been especially influential (Gärdenfors, 2000; Gärdenfors, 2014). These dimensions can range from concrete ones like weight, temperature and brightness, to very abstract ones like awkwardness or goodness. Once a domain, or selection of dimensions is established, a concept is defined as a region (usually a convex one) of the conceptual space. An example would be to define the colour *brown* as a region of a space made of the dimensions *Red*, *Green* and *Blue*. This geometric approach lends itself perfectly to model phenomena like prototyping (central point of the region), similarity (distance), metaphor (projection between different dimensions) and, more importantly for our concerns here, fluid changes in meaning due to the effects of context.

Warglien and Gärdenfors (2015) use conceptual spaces to look at *meaning negotiation* in conversation. They investigate the mechanisms, consciously or unconsciously, employed by the people involved in conversation to negotiate meaning of vague predicates, in order to satisfy the coordination needed for communication. These tools help them to decide areas in which they don't agree as well. All these processes work by manipulating the conceptual dimensions in which meaning is represented. We will refer to them as **saliency manipulation** because their main role is to dynamically rise or lower the perceived importance of certain conceptual dimensions.

The main mechanism by which speakers can modify saliency of conceptual dimensions are the automatic *priming* effects described by, for example, Pickering and Garrod (2004): mentioning specific words early in the conversation can make the dimensions associated with such words more relevant. Speakers can also explicitly try to remove dimensions from the domain in order to promote agreement, or bring in new dimensions by using *metaphoric projections*. Because metaphors can be understood as mappings that transfer structure from one domain to another, they can introduce new dimensions and meaning to the conversation.

The lion Ulysses emphasizes Ulysses' courage but hides his condition of a castaway in Ogiya. Thus metaphors act by orienting communication and selecting dimensions that may be more or less favorable to the speaker. By suggesting that a storm hit the financial markets, a bank manager can move the conversation away from di-

mensions pertaining to his own responsibilities and instead focus on dimensions over which he has no control. (Warglien and Gärdenfors, 2015)

From this perspective, then, the change in meaning is no longer a change in the meaning of a specific word, but a change in the mind of the hearer (or reader), a change in their *mental state* triggered by their interaction with the context. We saw an example of the meaning of the word "butter" being *contextually modulated* before, let's see some examples of *saliency manipulation* having an effect on the same word:

9. My muffins were a failure, I should have used butter or margarine instead of olive oil.
10. Vegan chefs replace animal fats, like butter, with plant based ones like olive oil or margarine.
11. Vegan influencers believe the consumption of animal products is cruel and unnecessary.

In example (9), in the context of a baking recipe, important dimensions are related to the physical properties of butter, margarine and olive oil. When focusing on these type of dimensions butter and margarine seem more similar because they are both solid while olive oil is liquid. In contrast, in the following example (10) we bring up ideas about veganism and the dimension of animal versus based plant products becomes very salient. This could bring margarine and olive oil closer together and distance both of them from butter, which is an animal product.

There are important differences between this *saliency manipulation* effect and the similarly "graded" *contextual modulation* effect. In the previous example (5) *poured* modulated the meaning of the word butter by promoting its "liquid" trait. This effect is limited to the word butter. On the contrary, if the context triggers changes in the saliency of conceptual dimensions, any word the annotator evaluates after the change takes place will be affected by it. Once the idea of animal vs plant based is introduced, the change takes place in the mind of the annotator and the perception of the meaning of not only butter, but margarine and olive oil is impacted as well. Our hypothesis is that, by using *saliency manipulation*, a context like example (11) can have an impact in the scoring of the similarity of butter, margarine and olive oil without these words even being present in the context. Something that would be impossible if we were looking only at the *contextual modulation* and *sense selection* effects.

The expectation that priming is the main mechanism for modifying saliency has its own implications: Branigan et al. (2000) found that priming effects are much stronger in the context of as natural dialog as possible, when speakers had no time constraints and could respond at their own pace. These results were taken into account when designing our dataset and annotation methodology: it is crucial for us to create an annotation process in which the annotator interacts with the context, and does so in as natural a way as possible, before they rate the similarity. Because priming is an automatic process, them knowing that they should be annotating similarity in context becomes a lot less important.

<b>Word1: bank</b> <b>Word2: money</b>
<b>Context1</b> Located downtown along the east <b>bank</b> of the Des Moines River ...
<b>Context2</b> This is the basis of all <b>money</b> laundering, a track record of depositing clean money before slipping through dirty money ..

Figure 1: Example from the SCWS dataset, the focus is in the different senses of the word **bank** and there is one independent context per word.

### 2.3. Existing Datasets

There are a few examples of datasets which take context into account. However, so far these have been motivated by discrete *sense disambiguation*, and therefore take a view of word meaning as discrete (taking one of a finite set of senses) rather than continuous; they are therefore not suited for the more graded effects we are interested to look into. The **Stanford Contextual Word Similarity (SCWS)** dataset (Huang et al., 2012) does contain graded similarity judgements of pairs of words in the context of organically occurring sentences (from Wikipedia). However it was designed to evaluate a discrete multi-prototype model, so the focus was on the contexts selecting for one of the word senses. This resulted in them presenting each of the two words of the pair in their own distinct context. From our point of view this approach has some drawbacks: First, even in the cases where they annotated the same pair twice, we find ourselves with four different contexts, each affecting the meaning of each of the instances of the words independently, and it is not possible to produce a systematic comparison of contextual effects on pairwise similarity. Second, beyond the independent lexical semantics of each word being affected by their independent *local context*, the annotator is being presented with two completely independently occurring contexts at the same time. Even if the two contexts did organically occur on their own, this combination of the two did not, and we have seen before how crucial we think keeping the interaction with the context as natural as possible is. There is no easy way to know how this newly assembled *global context* affects the cognitive state of the annotators and their perception of similarity. The same goes for the contextually-aware models trying to predict their results. Joining the contexts before feeding them to the model could create conflicting, difficult to predict effects, but feeding each context independently is fundamentally different to what humans annotators were presented with.

In addition to these limitations of the independent contexts approach, the scores found in SCWS show a worryingly low inter-rater agreement (IRA), measured as the Spearman correlation between different annotators. As pointed out by (Pilehvar and Camacho-Collados, 2019), the mean IRA between each annotator and the average of the rest, which is considered a human-level upper bound for model’s performance, is 0.52; while the performance of a simple context-independent model like word2vec (Mikolov et al., 2013) is 0.65. Examining the scores more in detail, we find that many scores show a very large standard deviation, with annotators rating the same pair very differently. One possible reason for this may lie in the annotation design: the task it-

self does not directly enforce engagement with the context, and the words were presented to annotators highlighted in boldface, making it easy to pick them out from the context without reading it; thus potentially leading to a lack of engagement of the annotators with the context.

A lot of these limitations were addressed by the more recent **Words-in-Context (WiC)** dataset (Pilehvar and Camacho-Collados, 2019). With a more direct and straightforward take on word sense disambiguation, each entry of the dataset is made of two lexicographer examples of the same word. The entry is completed with a positive value (T) if the word sense in the two examples/context is the same, or with a negative value (F) if the contexts point to different word senses. One advantage of this design is that it forces engagement with the context; another is that it creates a task in which context-independent models like word2vec “would perform no better than a random baseline”. Human annotators are shown to produce healthy inter-rater agreement scores for this dataset. However the dataset is again focused in looking at discrete word senses and cannot therefore capture continuous effects of context in the judgements of similarity between different words.

These datasets are also available only in English, and do not allow models to be evaluated across different languages.

### 3. Dataset and Task Design

CoSimLex will be based on pairs of words from SimLex-999 (Hill et al., 2015); the reliability and common use of this dataset makes it a good starting point and allows comparison of judgements and model outputs to the context-independent case. For Croatian and Finnish we use existing translations of Simlex-999 (Mrkšić et al., 2017; Venekoski and Vankka, 2017; Kittask, 2019). In the case of Slovene, we have produced our own new translation (Pollak et al., 2020), following the methodology used by (Mrkšić et al., 2017) for Croatian.

The English dataset consists of 333 pairs; the Croatian, Finnish and Slovene datasets of 111 pairs each. Each pair is rated within two different contexts, giving a total of 1554 scores of contextual similarity. This poses a difficult task: to find suitable, organically occurring contexts for each pair; this task is more pronounced for languages with less resources, and as a result the selection of pairs is different for each language.

Each line of CoSimLex will be made of a pair of words selected from Simlex-999; two different contexts extracted from Wikipedia in which these two words appear; two scores of similarity, each one related to one of the contexts; and two scores of standard deviation. Please see Figure 2 for an example from our English pilot.

<b>Word1: population</b> <b>Word2: people</b>	<b>SimLex: <math>\mu</math> 7.68 <math>\sigma</math> 0.80</b>
<b>Context1</b>	<b>Context1: <math>\mu</math> 6.49 <math>\sigma</math> 1.40</b>
Disease also kills off a lot of the gazelle <b>population</b> . There are many <b>people</b> and domesticated animals that come onto their land. If they pick up a disease from one of these domesticated species they may not be able to fight it off and die. Also, a big reason for the decline of this gazelle population is habitat destruction.	
<b>Context2</b>	<b>Context2: <math>\mu</math> 7.73 <math>\sigma</math> 1.77</b>
But the discontent of the underprivileged, landless and the unemployed sections remained even after the reforms. The crumbling industries give rise to extreme unemployment, in addition to the rapidly growing <b>population</b> . These <b>people</b> mostly belong to the SC/ST or the OBC. In most cases, they join the extremist organizations, mentioned earlier, as an alternative to earn their livelihoods.	

Figure 2: Example from the English pilot, showing a word pair with two contexts, each with mean and standard deviation of human similarity judgements. The original SimLex values for the same word pair without context are shown for comparison.

**Evaluation Tasks and Metrics** The first practical use of CoSimLex will be as a gold standard for the public SemEval 2020 task 3: *Graded Word Similarity in Context*. The goal of this task is to evaluate how well modern context-dependent embeddings can predict the effect of context in human perception of similarity. In order to do so we define two subtasks and two metrics:

**Subtask 1 - Predicting Changes:** In subtask 1, participants must predict the *change* in similarity ratings between the two contexts. In order to evaluate it we calculate the difference between the scores produced by the model when the pair is rated within each one of the two contexts. We do the same with the average of the scores produced by the human annotators. Finally we calculate the uncentered Pearson correlation. A key property of this method is that any context-independent model will predict no change and get strongly penalised in this task.

**Subtask 2 - Predicting Ratings:** In subtask 2, participants must predict the absolute similarity rating for each pair in each context. This will be evaluated using Spearman correlation with gold-standard judgements, following the standard evaluation methodology for similarity datasets (Hill et al., 2015; Huang et al., 2012). Good context-independent models could theoretically give competitive results in this task, however we still expect context-dependent models to have a considerable advantage.

## 4. Annotation Methodology

As starting point for our annotation methodology, we adapted the annotation instructions used for SimLex-999. This way we benefit from its tested method of explaining how to focus on *similarity* rather than *relatedness* or *association* (Hill et al., 2015). As explained in their original paper, *cup* and *mug* are very similar, while *coffee* and *cup* are strongly related but not similar at all. For English we adopted a modified version of their crowd-sourcing process: we use *Amazon Mechanical Turk*, with the same scoring scale (0 to 6), the same post-processing and cleaning of the data (a necessary step when working with this kind of crowd-sourcing platform), and achieve similarly good inter-annotator agreement. For the less-resourced languages, crowdsourcing is not a viable option due to lack of available speakers, and we recruit annotators directly. This means fewer annotators (for Croatian, Finnish and Slovene,

12 annotators vs 27 in English), however the average quality of annotation is a lot higher and the data requires less post-processing - see Section 5 for details.

### 4.1. Finding Suitable Contexts

For each word pair we need to find two suitable contexts. These contexts are extracted from each language’s Wikipedia. They are made of three consecutive sentences and they need to contain the pair of words, appearing only once each. English is by far the easiest language to work with, not only because of the amount and quality of the text contained in the English version of Wikipedia but because the other four languages are highly inflected (Croatian, Finnish and Slovene). To overcome this, we work with data from (Ginter et al., 2017)<sup>2</sup> which contains tokenised and lemmatised versions of Wikipedia for 45 languages.

We first find all the possible candidate contexts for each word pair, and then select those candidates that are most likely to produce different ratings of similarity. The differences are expected to be small, especially in words that don’t present several senses and are not highly polysemous, so we need a process that has the most chances of finding contexts that make a difference. We use a dual process in which we use ELMo and BERT to rate the similarity between the target pair within each of the candidate contexts. Then we select the 2 contexts in which ELMo scored the pair as the most similar, and the 2 contexts in which it scored them as most different. We do the same using BERT scores. This gives us 4 contexts in which our target words are scored as very similar by the models and 4 contexts in which they are scored as very different.

The final selection of two contexts is made by expert human annotators, one per language. We construct online surveys with these 8 contexts and ask them to select the two in which they think the word pair is the most and the least similar, trying to maximise the potential contrast in similarity. In addition, we ask them how much potential for a difference they see in the contexts selected. This gives us not only the contexts we need, but a predicted performance and direction of change for use in later analysis.

In the case of less resourced languages, the smaller size and lower quality of the Wikipedia text resources require some extra steps to ensure the quality of the final annota-

<sup>2</sup><http://hdl.handle.net/11234/1-1989>

Word1: <b>čovjek</b> (adult male)    Word2: <b>dijete</b> (child)	
<p><b>Context1</b></p> <p>Špinat ima dosta željeza, ali i oksalne kiseline. Oksalna kiselina veže kalcij i čini ga neupotrebljivim za ljudski organizam. Prema novijim istraživanjima, špinat se ne preporuča kao česta hrana mladim osobama i <b>djeci</b>, ali je izvrsna hrana za starije <b>ljude</b>.</p> <p>(Spinach has plenty of iron but also oxalic acid. Oxalic acid binds calcium and renders it unusable for the human body. According to recent research, spinach is not recommended as a common food for younger people and <b>children</b>, but it is an excellent food for older <b>people</b>.)</p>	<p><b>Context1: <math>\mu</math> 2.5 <math>\sigma</math> 1.76</b></p>
<p><b>Context2</b></p> <p>Nakon što su <b>ljudi</b> u selu saznali da je trudna, počinju sumnjati na dr. Richardsona jer je on proveo najviše vremena s njom. Kako vrijeme prolazi, pritisak glasina na kraju prisiljava liječnika da se preseli. Odluči se oženiti s Belindom i uzeti <b>dijete</b> sa sobom.</p> <p>(After <b>people</b> in the village find out she is pregnant, they begin to suspect Dr. Richardson because he spent the most time with her. As time goes on, the pressure of the rumors eventually forces the doctor to move. He decides to marry Belinda and take her <b>child</b> with him.)</p>	<p><b>Context2: <math>\mu</math> 4.25 <math>\sigma</math> 0.95</b></p>

Figure 3: Example from the Croatian pilot, showing the word pair with two contexts, each with mean and standard deviation of human similarity judgements. This example showed one of the most significant contextual effects in the pilot; it went in the opposite direction to the one predicted by the expert annotator. Note the effect of stemming: the target word *čovjek* appears in both cases via its irregular plural, *ljudi* (nominative) or *ljude* (accusative); and *dijete* appears in Context 1 in its dative plural form *djeci*. English translations (generated using Google Translate with manual post-correction) are shown here for exposition purposes but are not part of the dataset.

tion. For these languages we run the contexts through a set of heuristic filters to try to remove badly constructed ones. In addition we produce 16 candidates instead of 8 for the expert annotators to choose from, and we add the possibility for them to delete parts of the context in order to make them easier to read. Adding text is not allowed, in order to ensure that contexts are natural.

#### 4.2. Contextual Similarity Annotation

The next step is to obtain the contextualised similarity annotations. Our goal is to capture the kind of contextual phenomena discussed in Section 2, lexical meaning modulation and conceptual salience manipulation. In order to maximise our chances we define three goals:

- We want the interaction with the context to be as natural as possible, so as to maximise priming effects and capture the potential change in the salience of conceptual dimensions.
- We need a way in which annotators have the chance to account for lexical modulation within the sentence.
- We need to avoid the apparent lack of engagement we saw in the SCWS annotators.

With these goals in mind we designed a two-step mixed annotation process. Our online survey interface is composed of two pages per pair of words and context (each annotator scores only one of the contexts). In the first page the annotators are presented with the context, and asked to read it and come up with two words “inspired by it”. Once this is complete, the second page shown presents the context again, but with the pair of words now highlighted in bold; they are now asked to rate the similarity of the pair of words within the sentence.

The second page is the main scoring task; it is designed to capture changes in scores of similarity due both to lexical modulation and — because we hope the annotators are

still primed by their recent previous engagement with the context — the changes in the salience of conceptual dimensions. The separate task on the first page is intended to make annotators engage fully with the whole context, while maintaining a natural interaction with it to maximise any priming effects. One of the possible problems we identified in the the SCWS annotation process is the fact that the words were always highlighted in bold, making it easy for annotators (Amazon Mechanical Turk workers) to just look at the pair of words in isolation and to not read the rest of the contexts. Our initial task is designed to prevent this (the words are not bold in the first page).

In English, given the resources available, we follow SimLex-999 closely: we will use Amazon Mechanical Turk to get 27 annotators per pair and context. Annotators do not score the same pair twice: 27 annotators score the pair within one context and another 27 in the other. This means the whole dataset can be annotated at the same time. Reliability of annotations will be ensured by an adapted version of SimLex-999’s post-processing, which includes rating calibration and the filtering of annotators with very low correlation to the average rating. In addition, we will use responses to the first annotation question to check annotator engagement with the context text and thus filter low quality raters.

For Croatian, Finnish and Slovene we recruit annotators directly: this means we have less of them (12 vs 27) but we expect the quality of the annotation to be better (and pilots confirm this – see below). It also means, however, that we must use the same annotators to rate the two contexts of each pair. This has an advantage, because it controls for the variation in the particular judgement of different annotators, but means that we introduce a week’s delay in between annotations in order to make sure they don’t remember, and are influenced by, their own previous score.

Word1	Word2	Context1	Context2	STDev1	STDev2	P-Value
water	ice	2.57	8.13	2.60	1.82	2.18E-08**
friendly	generous	4.44	3.92	2.85	3.56	0.258
keep	protect	2.50	3.75	2.66	2.22	0.036**
pact	agreement	8.73	8.97	1.89	1.53	0.302
narrow	broad	0.42	1.97	1.19	2.60	0.012**
arm	neck	3.81	1.27	2.89	1.97	0.002**
cottage	cabin	8.07	9.56	2.37	0.94	0.003**
inform	notify	9.31	9.80	0.97	0.55	0.019**
mother	guardian	3.94	7.28	3.15	2.54	0.0001**
car	bicycle	4.12	4.85	2.58	2.46	0.169

Table 1: Example results from the English dataset showing the mean similarity, standard deviation and p-value calculated using the Mann-Whitney U test.

Word1	Word2	Context1	Context2	STDev1	STDev2	P-Value
nadbiskup	biskup	6.67	6.30	2.66	2.61	0.325
sretan	mlad	1.50	0.30	2.28	0.67	0.099*
kost	čeljust	6.00	3.33	2.38	2.24	0.013**
zvijer	životinja	9.44	6.30	1.09	3.09	0.004**
priča	tema	2.59	7.64	1.88	2.51	0.0004**

Table 2: Example results from the Croatian dataset showing the mean similarity, standard deviation and p-value calculated using the Mann-Whitney U test.

\*Read the following text and write down two words inspired by it:

Though for some reason often described as a farm boy, Pollard was 40 years old when he fell at Breed's Hill, reportedly beheaded by a cannon ball fired from the British ship the Somerset in Boston Harbor. Accounts of the circumstances of his death differ. A popular book Now We Are Enemies: The Story of Bunker Hill by Thomas Fleming (1960) relates an often told story that he was killed as he led other soldiers to water.

First word:

Second word:

Figure 4: First page shown for each word pair annotation task: annotators must read the context and come up with two words inspired by it. At this point, the word pair to be scored is not known to the annotator.

\*Read the sentences again and then score the similarity between the words **boy** and **soldier** when compared **within this specific text**:

Though for some reason often described as a farm **boy**, Pollard was 40 years old when he fell at Breed's Hill, reportedly beheaded by a cannon ball fired from the British ship the Somerset in Boston Harbor. Accounts of the circumstances of his death differ. A popular book Now We Are Enemies: The Story of Bunker Hill by Thomas Fleming (1960) relates an often told story that he was killed as he led other **soldiers** to water.

0 = Not similar at all

6 = Extremely similar

0    1    2    3    4    5    6

Figure 5: Second page shown for each word pair annotation task: the same context is now shown with the target words in bold, and annotators must give a similarity score for the word pair within that particular context.

## 5. Current Status

**Methodology prototyping** We have run three pilots with 13 pairs of words each to confirm the annotation design and methodology. Each study tested a slight variation: in the first pilot, annotators rated *relatedness* in addition to similarity; the second focused on similarity, and tested the use of contexts related to the target words but not containing them; the third experimented with marking the target words in the context paragraphs using boldface font.

The first pilot confirmed that (as with SimLex) similarity is a more useful metric for this task than relatedness, displaying a higher inter-annotator agreement and more variation between contexts; we therefore use similarity as the basis of our dataset, as described above.

The results of the second pilot saw significant contextual effects in many examples, including some in which the target words weren't included in the contexts. This indicates that our method seems suitable for capturing priming effects and salience manipulation, or at least some kind of cognitive effects different from lexical contextualisation.

The third pilot showed much lower agreement and lower difference between contexts: we take this as confirmation of our suspicion (from analysis of SCWS) that marking the target words makes it easy for annotators to ignore the rest of the context paragraph, and therefore use the two-stage annotation methodology described above, in which target words are *not* initially marked.

**Results** The dataset contains 341 entries in English, 113 in Croatian, 112 in Slovene and 25 in Finnish. Each of the entries contains a pair of words evaluated in two different contexts. Please see Table 1 for examples from the English results and Table 2 for Croatian.

Inter-rater agreement (IRA) is measured as Spearman correlation between each rating and the average values. After post-processing the data our dataset's IRA is surprisingly similar between the different languages: for English, Croatian and Slovene the mean is  $\rho = 0.77$ , while Finnish achieves  $\rho = 0.80$  (likely due to the small sample); these compare well to other related datasets (SimLex-999  $\rho = 0.78$ , SCWS  $\rho = 0.52$ ). The crowd-sourced nature of English data results in a higher percentage of annotations being dropped in the post-processing. However the fact that both methods converge to the same IRA is encouraging, and seems to indicate that both methods achieve comparable results.

The statistical significance of the difference in similarity evaluation between contexts was assessed using the Mann-Whitney U test. In English, from the 341 entries, 220 results showed a statistically significant difference at  $p < 0.1$  (ratio = 0.65), and 208 did so at  $p < 0.05$  (ratio = 0.61). The results again were quite similar for Croatian and Slovene with 73 and 65 statistically significant results at  $p < 0.1$  (ratio = 0.65 and 0.58). Finnish results showed a much smaller ratio of statistically significant results (8 entries, ratio = 0.33), which could be due to the small sample but may deserve further investigation. However, as in the case of the inter-rater agreement, the fact that English, Croatian and Slovene results are very similar is a good sign for both methodologies: the English crowd-sourced annotation and

the smaller sample of better quality annotation we used for Croatian and Slovene.

## 6. Conclusion

The growing use of context-dependent language models and representations in NLP motivates the need for a dataset against which they can be evaluated, and which can test their ability to reflect human perceptions of context-dependent meaning. CoSimLex will provide such a dataset, and do so across a number of less-resourced languages as well as English. The full dataset was provided for the evaluation stage of SemEval 2020 at the beginning of February 2020, and will be made public when the competition is over (before the LREC2020 conference).

## 7. Acknowledgements

This research is supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No 825153, project EMBEDDIA (Cross-Lingual Embeddings for Less-Represented Languages in European News Media). The results of this publication reflect only the authors' views and the Commission is not responsible for any use that may be made of the information it contains. The first author is also supported by the EPSRC and AHRC Centre for Doctoral Training in Media and Arts Technology (EP/L01632X/1); the second author is also supported by the EPSRC project Streamlining Social Decision Making for Improved Internet Standards (SoDeStream, EP/S033564/1).

## 8. Bibliographical References

- Branigan, H. P., Pickering, M. J., and Cleland, A. A. (2000). Syntactic co-ordination in dialogue. *Cognition*, 75(2):B13–B25.
- Cruse, D. A. (1986). *Lexical semantics*. Cambridge university press.
- 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.
- Evans, V. and Green, M. (2018). *Cognitive linguistics: An introduction*. Routledge.
- Finkelstein, L., Gabilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., and Ruppín, E. (2002). Placing search in context: The concept revisited. *ACM Transactions on information systems*, 20(1):116–131.
- Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. MIT Press.
- Gärdenfors, P. (2014). *The geometry of meaning: Semantics based on conceptual spaces*. MIT Press.
- Ginter, F., Hajič, J., Luotolahti, J., Straka, M., and Zeman, D. (2017). CoNLL 2017 shared task - automatically annotated raw texts and word embeddings. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Hill, F., Reichart, R., and Korhonen, A. (2015). Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695, December.
- Huang, E. H., Socher, R., Manning, C. D., and Ng, A. Y. (2012). Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 873–882. Association for Computational Linguistics.
- Kittask, C. (2019). *Computational Models of Concept Similarity for the Estonian Language*. Bachelor’s thesis, University of Tartu.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. In *Proceedings of Workshop at ICLR*.
- Mrkšić, N., Vulić, I., Ó Séaghdha, D., Leviant, I., Reichart, R., Gašić, M., Korhonen, A., and Young, S. (2017). Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *Transactions of the Association for Computational Linguistics*, 5:309–324.
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Pickering, M. J. and Garrod, S. (2004). Toward a mecha-

nistic psychology of dialogue. *Behavioral and brain sciences*, 27(2):169–190.

- Pilehvar, M. T. and Camacho-Collados, J. (2019). Wic: the word-in-context dataset for evaluating context-sensitive meaning representations. 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 1267–1273.
- Venekoski, V. and Vankka, J. (2017). Finnish resources for evaluating language model semantics. In *Proceedings of the 21st Nordic Conference on Computational Linguistics, NoDaLiDa, 22-24 May 2017, Gothenburg, Sweden*, number 131 in Linköping Electronic Conference Proceedings, pages 231–236. Linköping University Electronic Press, Linköpings universitet.
- Warglien, M. and Gärdenfors, P. (2015). Meaning negotiation. In *Applications of conceptual spaces*, pages 79–94. Springer.

## 9. Language Resource References

- Armendariz, Carlos Santos and Purver, Matthew and Ulčar, Matej and Pollak, Senja and Ljubešić, Nikola and Robnik-Šikonja, Marko and Granroth-Wilding, Mark and Vaik, Kristina. (2020). *CoSimLex*. EMBEDDIA project, ISLRN 613-489-674-355-0. Available via CLARIN <http://hdl.handle.net/11356/1308>.
- Pollak, Senja and Ljubešić, Nikola and Vulić, Ivan. (2020). *A Slovenian Translation for SimLex*. EMBEDDIA project, ISLRN 613-489-674-355-0. Available via CLARIN <http://hdl.handle.net/11356/1309>.