

# The DUREl Annotation Tool: Human and Computational Measurement of Semantic Proximity, Sense Clusters and Semantic Change

Dominik Schlechtweg<sup>1</sup>, Shafqat Mumtaz Virk<sup>2</sup>, Pauline Sander<sup>1</sup>,  
Emma Sköldbberg<sup>2</sup>, Lukas Theuer Linke<sup>1</sup>, Tuo Zhang<sup>1</sup>,  
Nina Tahmasebi<sup>2</sup>, Jonas Kuhn<sup>1</sup>, Sabine Schulte im Walde<sup>1</sup>  
<sup>1</sup>University of Stuttgart, <sup>2</sup>University of Gothenburg

## Abstract

We present the DUREl tool that implements the annotation of semantic proximity between uses of words into an online, open source interface. The tool supports standardized human annotation as well as computational annotation, building on recent advances with Word-in-Context models. Annotator judgments are clustered with automatic graph clustering techniques and visualized for analysis. This allows to measure word senses with simple and intuitive micro-task judgments between use pairs, requiring minimal preparation efforts. The tool offers additional functionalities to compare the agreement between annotators to guarantee the inter-subjectivity of the obtained judgments and to calculate summary statistics giving insights into sense frequency distributions, semantic variation or changes of senses over time.

## 1 Introduction

The concept of **semantic proximity** between word uses has a rather long tradition in Cognitive Semantics (Blank, 1997) and is also acknowledged by notable lexicographers (Kilgarriff, 1997). Semantic proximity quantifies how much the meanings of two word uses “have in common” (Schlechtweg, 2023, cf. p. 25). The concept serves as (often vague) criterion in the lexicographic **clustering** process (Kilgarriff, 2007) and is thus essential to the process of creating dictionary entries of **word senses**. Being essential to the identification of word senses, semantic proximity has further relevance to research building on senses such as lexical semantic change or semantic variation (Schlechtweg, 2023).

In lexical semantics, multiple approaches operationalized semantic proximity in **human annotation** studies, showing that the concept can be practically implemented with reasonable agreement between annotators and correspondence to alternative annotation procedures (Soares da Silva, 1992; Erk et al., 2013; Schlechtweg et al., 2018).

Recently, there has been an upsurge on research in computational modeling of semantic proximity between word uses under the name of **Word-in-Context** models (Pilehvar and Camacho-Collados, 2019; Armendariz et al., 2020), resulting from advances in modeling the meaning of word uses with contextualized embeddings (Peters et al., 2018; Devlin et al., 2019). These models reach high performance (He et al., 2020; Raffel et al., 2020) and thus serve as an excellent starting point for any practical task building on semantic proximity like creating dictionary entries, finding novel/non-recorded senses or identifying words that change their meaning.

We present the **DUREl tool** combining the above-described lines of research into a user-friendly, online annotation interface with open source code.<sup>1</sup> The basic annotation data gathered in the system are judgments of semantic proximity between word uses (Blank, 1997; Erk et al., 2013) from multiple human or computational annotators. The tool facilitates the annotation task by providing a data inspection interface and automatic data validation for researchers, an intuitive task interface for annotators, guidelines in multiple languages for annotator training as well as tutorial data for annotator testing. Computational annotators are provided by optimized Word-in-Context models (Pilehvar and Camacho-Collados, 2019; Arefyev et al., 2021) trained on human semantic proximity judgments (i.a. Schlechtweg et al., 2021; Kutuzov and Pivovarova, 2021b; Zamora-Reina et al., 2022). Semantic proximity judgments are represented in a graph (McCarthy et al., 2016), clustered with an automatic **graph clustering** technique (Schlechtweg et al., 2020, 2021) and visualized for analysis (Theuer Linke, 2023). This allows to measure word senses from **simple** and **intuitive** semantic proximity judgments between use

<sup>1</sup><https://www.ims.uni-stuttgart.de/data/durel-tool>

pairs. The tool offers functionalities to compare the agreement between annotators to guarantee the **inter-subjectivity** of the obtained judgments. It provides further functionalities to calculate summary statistics over the annotated data giving insights into sense frequency distributions, semantic variation or changes of senses over time. The computational annotator component allows to generate word sense clusters for large sets of words and word uses, making it possible, for instance, to analyze large amounts of data or to search unlabelled data systematically for new senses.

## 2 Related Work

We now compare DUREl to existing text annotation tools and related tools from electronic lexicography. There is a number of general-purpose text annotation tools such as CATMA (Gius et al., 2022), INCEpTION (Klie et al., 2018), MTURK<sup>2</sup>, PhiTag<sup>3</sup>, POTATO (Pei et al., 2022) or Toloka<sup>4</sup>. Many of these allow to define a wide range of custom tasks and can in principle cover use pair annotation (cf. e.g. Giulianelli et al., 2020; Kutuzov and Pivovarova, 2021a). However, this often requires preparation efforts including the writing of small programs as well as the formulation of guidelines and data for annotator training. The aim of the DUREl tool is to minimize such additional efforts around organizing a use pair annotation study. DUREl achieves this by focusing on this particular task, implementing standardized procedures which have proven to work well in previous studies (i.a. Schlechtweg et al., 2018; Hätyy et al., 2019; Schlechtweg et al., 2020; Baldissin et al., 2022). Further unique features of the tool are the task-specific data analysis (see Section 4.2.5) and computational annotators (see Section 4.2.2). Other tools, while offering annotation for a range of tasks, cannot offer these specific possibilities.

With its focus on a semantic annotation task (semantic proximity) and the sense inference functionality (see Section 4.2.4), the DUREl tool is relevant to lexicography. The most widely used lexicographic tool is Sketch Engine.<sup>5</sup> The DUREl tool differs from Sketch Engine by focusing primarily on crowdsourcing human annotations and offering computational annotation models, whereas

<sup>2</sup><https://www.mturk.com/>

<sup>3</sup><https://www.ims.uni-stuttgart.de/data/phetag>

<sup>4</sup><https://toloka.ai/en/docs/>

<sup>5</sup><https://www.sketchengine.eu>




Table 1: The DUREl relatedness scale (Schlechtweg et al., 2018) on the left and its interpretation from Schlechtweg (2023, p. 33) on the right.

Sketch Engine offers frequency-based corpus analysis and manual dictionary making. The two tools could be integrated to provide improved analysis of word meaning.

## 3 Background

The DUREl tool implements the word sense annotation scheme developed in Schlechtweg et al. (2018, 2020, 2021) and described in detail in Schlechtweg (2023, pp. 31ff.). The scheme builds on use pair proximity annotation on a relatedness scale combined with a graph clustering procedure. Annotators are asked to judge the semantic relatedness of use pairs, such as the two uses of *arm* in (1) and (2), on the scale in Table 1.

- (1) [...] she opened a vein in her little **arm**, and dipping a feather in the blood [...]
- (2) [...] he saw her within reach of his **arm**, yet the light of her eyes seemed as far off [...]

The annotated data of a word is then represented in a graph, which we call Word Usage Graph (WUG), where vertices represent word uses, and weights on edges represent the (median) semantic relatedness judgment of a use pair. The final WUGs are clustered with Correlation Clustering (Bansal et al., 2004; Schlechtweg et al., 2020).

Consider the example in Figure 1 to understand how senses and semantic change can be inferred with the DUREl annotation procedure: Assume that WUG  $G$  represents the semantic proximity structure annotated for the set of word uses  $U$  of the English word *arm* displayed in Table 2. The uses  $U_1 = \{A, B, C\}$  and  $U_2 = \{D, E, F\}$  were sampled from the two time periods 1820–1860 and 1950–1990 respectively ( $t_1, t_2$ ). We derive sense clusters on  $G$  by building three clusters of uses with high semantic proximity and low semantic proximity to other clusters:  $C_1 = \{A, C, F\}$  (blue),




Figure 1: Clustered WUG  $G$  of *arm* (left), subgraph for 1st time period  $G_1$  (middle) and subgraph for 2nd time period  $G_2$  (right). **black/gray** lines for **high** ( $\geq 2.5$ )/**low** ( $< 2.5$ ) edge weights. Spatial proximity of nodes loosely corresponds to their semantic proximity annotation. Visualization taken from [Schlechtweg \(2023, p. 40\)](#).

$C_2 = \{D, E\}$  (orange),  $C_3 = \{B\}$  (green). We then build the time-specific subgraphs  $G_1$  and  $G_2$  and are now able to compare the clusters between time periods. For instance,  $C_3$  only exists in the first time period while  $C_2$  only exists in the second time period.

## 4 Tool description

DURel is a web application supporting user interaction by browser (Mozilla Firefox is recommended).<sup>6</sup> It was created on an architecture incorporating Java Spring for backend, PostgreSQL for the database, HTML/CSS/JavaScript plus Thymeleaf for frontend, and a CSV format for transferring annotations. The source code of the tool is publicly available under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.<sup>7</sup>

### 4.1 Data format

Sets of a word uses as in Table 2 are the basic data type that DURel works with. These can be sampled from any corpus by the user and then uploaded to the tool (after registration) as one CSV per word through an interface. The files should contain at least one use context per line along with the target word and target sentence character indices. Additionally, they can contain meta-information such as the target word’s POS, the date of the use or a user-specified grouping tag for uses, which can be used later for statistics and visualization.

<sup>6</sup>Find a demo video at <https://www2.ims.uni-stuttgart.de/video/durel-tool/231121-durel-tool-demo.mp4>.

<sup>7</sup>[https://github.com/ChangeIsKey/durel\\_tool](https://github.com/ChangeIsKey/durel_tool)

### 4.2 Features

We now describe the central features of the DURel tool.

#### 4.2.1 Project management

By clicking on the tab "Upload Project" the user can create an annotation project. He has to specify the project language and choose a set of CSV files containing word uses to upload. He has then two options options to create annotation instances (use pairs): (i) Let the system generate random sequences per annotator of all possible combinations of use pairs per word. (ii) Upload a user-defined sequence of annotation pairs, which will be presented to each annotator in randomized order. In both cases, the sequence within use pairs is swapped with a 0.5 probability. Alternatively, the user can upload his own gold annotations as annotation project.

Under the tab "My Projects", users can manage, download or delete their projects, as well as assign specific users to projects or make a project entirely public (see Figure 4). A user that has been granted access can annotate for that project.

#### 4.2.2 Use pair annotation

Use pairs are judged on the semantic proximity scale shown in Table 1. An annotator has the option of assigning a label between 1 (unrelated) and 4 (identical), or to assign no label (Cannot decide).

**Human annotation** When an annotator registers to the DURel tool, they first need to successfully complete an annotation tutorial. Before starting the tutorial, the annotator is asked to read the guidelines page, which explains how to make semantic proximity judgments. In the tutorial, the annotator annotates only a few use pairs of different words.

1824	and taking a knife from her pocket, she opened a vein in her little <b>arm</b> ,
1842	And those who remained at home had been heavily taxed to pay for the <b>arms</b> , ammunition;
1860	and though he saw her within reach of his <b>arm</b> , yet the light of her eyes seemed as far off
...	
1953	overlooking an <b>arm</b> of the sea which, at low tide, was a black and stinking mud-flat
1975	twelve miles of coastline lies in the southwest on the Gulf of Aqaba, an <b>arm</b> of the Red Sea.
1985	when the disembodied <b>arm</b> of the Statue of Liberty jets spectacularly out of the

Table 2: Sample of diachronic corpus taken from [Schlechtweg \(2023, p. 41\)](#).

These judgments are then checked against a hidden gold standard. The annotator only passes the tutorial if he reaches a certain level of agreement.

After passing the tutorial, annotators can annotate for projects to which they have been assigned by clicking on the "Annotate" tab (see Figure 6 in Appendix). Each project is divided into several words. Annotators can decide which word they annotate and the annotation can be paused at any time.

For each annotation instance, the tool records the judgment label, an optional comment, the annotator name and the timestamp of the judgment.

**Computational annotation** We use optimized, multi-lingual Word-in-Context (WiC) models as computational annotators. These are treated analogous to human annotators in the DUREl system. They appear as users/annotators in the project management and can be assigned to an annotation project or an individual word by the user through clicking on the "Tasks" tab. Creating a task with a computational annotator will trigger an automatic annotation pipeline in the DUREl backend retrieving annotation labels with the respective model and storing them in the DUREl database. This allows large-scale data labeling for uploaded projects. The currently available computational annotators are:

- **Random** samples a random integer between 1 and 4 with uniform probability (as baseline).
- **XLMR+MLP+Binary**: XLMR ([Conneau et al., 2020](#)) vectorizer with multi-layer perception and binary classification head on concatenated vectors; trained on WiC dataset; predicts either value 1 or 4.
- **XL-Lexeme**: bi-encoder that vectorizes the input sequences using a XLMR-based Siamese Network ([Cassotti et al., 2023](#)); trained to minimize the contrastive loss with cosine distance on several WiC datasets; predicts either value 1 or 4 based on thresholding cosine similarity between vectors at 0.5.

### 4.2.3 Use Analysis

By clicking on the "Data" tab the user can (i) inspect uses or (ii) annotator judgments for each word in their projects (see Figure 5 in Appendix). Uses are displayed in concordance tables showing the aligned target words with additional context. The table offers sorting functions according to multiple criteria. Annotator judgments are displayed in a table with additional information such as the contexts of both uses, the data IDs, annotator name and annotator comment.

### 4.2.4 Annotation statistics

On the "Statistics" tab the user can calculate a range of summary statistics over the annotated data for analysis. These include (i) various measures of annotator agreement, (ii) label averages for words, groupings and annotators, (iii) comparisons between human and computational annotation. Further, the system allows to infer various meta-measures from the basic semantic proximity annotations. These include (i) word sense clustering, (ii) semantic variation measures, (iii) semantic change measures.

### 4.2.5 Annotation visualization

The visualization takes the form of a clustered network graph, calculated with Python in the backend, which the Pyvis visualization library, an interface for Vis.js, connects to a HTML, JavaScript and CSS frontend.<sup>8</sup> Nodes in the graph represent word uses and edge weights represent the annotations of use pairs (see Figure 7 in Appendix). Each edge's proximity score is calculated by taking the median between annotator judgments. When all of the annotators have given a score of "Cannot decide", the edge is marked with "NaN". If a node has at least half of its judgments as "Cannot decide", it is excluded from the clustering process. The plotted graph can be filtered according to annotation group, date, edge weight, annotator and noise

<sup>8</sup><https://www.ims.uni-stuttgart.de/data/wugs>



nodes or edges. The visualization also includes additional information, such as sense frequency and probability distributions, meta-information on uses, the clustering method and the node positioning method. Users can explore detailed information in additional "Stats" dropdowns. The main goal of the graph visualization is to make it easier to find information on word uses. In contrast to a large text document or table with hundreds of uses, in a clustered graph different meanings can be found quickly.

### 4.3 Frontend

The user interface of DUREl is designed with HTML, CSS, JavaScript and the Thymeleaf template engine. Thymeleaf is a natural choice for an application with Java Spring Boot as it provides full Spring Framework integration.

### 4.4 Backend

The DUREl backend is built with the Spring framework and a PostgreSQL database, both of which are open-source and widely used in industry. It is responsible for user and project management, transferring project data (upload, download), handling the annotation process, and data analysis (use analysis, statistics). The backend runs the WUG visualization pipeline as a subprocess (see Section 4.2.5).

The computational annotation pipeline is implemented as a separate component, built in Python with PyTorch<sup>9</sup> and Hugging Face<sup>10</sup>: It retrieves annotation tasks users create on the DUREl website, automatically generates annotations and sends them back to the DUREl backend. We deploy the two components on different servers and let them interact with each other by sending REST API requests. The separation gives us the possibility to use the computational annotation pipeline independently from DUREl with other annotation tools. It also allows us to deploy the computational annotation pipeline on any server, depending on computational workload, and to run multiple instances of the pipeline to spread the workload to multiple servers.

## 5 Case Studies and Evaluation

In this section, we describe two case studies to show the usefulness of the tool and to evaluate

<sup>9</sup><https://pytorch.org/>

<sup>10</sup><https://huggingface.co>

the computational annotators: (i) The *arm* example is a small scale study on a chosen test word for which we selected word uses from a corpus. (ii) The lexicographer case study conducted on a set of 18 Swedish words is related to an ongoing lexicographic project.

### 5.1 The *arm* example

For various senses of the word *arm*, Schlechtweg (2023, p. 41) selected the sentences (i.e. word uses) displayed in Table 2 from the online interface of COHA corpus (Davies, 2012). We uploaded these uses to the DUREl tool, which combined each sentence with every other sentence into use pairs. These pairs were then annotated with the XL-Lexeme annotator (see Section 4.2.2) and clustered using the correlation clustering algorithm. The resulting cluster structures were then visualized via the system's visualization function (see Figure 8 in Appendix). The generated cluster structures are very similar to the manual annotations: both clusterings distinguish the metaphorical sense 'arm of the sea' (recall Section 3). However, the computational annotator merges the 'body part' and the 'weapon' sense. Hence, while the computational annotator is not perfect, it works reasonably well for this example. We hope that this annotator will prove useful to discover meaning structures in various meaning-related study areas. One such application is given in the next subsection. (We also provide a video demonstration of this example.<sup>11</sup>)

### 5.2 The lexicographer case study

To test DUREl in a practical application, we used it in a small study on revealing semantic variation in Swedish. The study is closely related to the work of revising the comprehensive Swedish dictionary published by the Swedish Academy ('The Contemporary Dictionary of the Swedish Academy'; in short SO<sup>12</sup>), a general language definition dictionary with about 65,000 headwords. An important part of the revision work within the dictionary project is to examine whether the meanings of the headwords in SO have developed in some way since the last edition was published. However, the lexicographic team of SO currently do not, in a systematic way, use any formal computational methods for discovering semantic change

<sup>11</sup><https://www2.ims.uni-stuttgart.de/video/durel-tool/230623-durel-tool-demo.mp4>.

<sup>12</sup><https://svenska.se/>

Word	ARI	Word	ARI
ofantlig	1.0	klimat	0.083
enkelspårig	1.0	vansinnig	0.0
baksida	0.912	lirka	0.0
bagage	0.785	kapitulera	0.0
fasad	0.652	hemmaplan	0.0
vissen	0.645	hagla	0.0
skör	0.507	fotavtryck	0.0
rutten	0.333	tvärnita	-0.019
ventilera	0.303	kriga	-0.025
<b>Average</b>		<b>0.343</b>	

Table 3: Cluster evaluation based on ARI.

on the lexical level. Hence, the lexicographic challenge is finding relatively new meanings not already recorded. The question is whether the DUREl tool can point out meaning variation (as an indicator of change) by clustering different meanings of a word. For that purpose, as a first round of experiments, we selected a set of 18 established Swedish words that are already in SO. These words were thoughtfully selected based on their semantic characteristics: all of them have a main sense and one or more subsenses. For each word, we automatically extracted a collection of 50 random sentences from the SVT (Swedish Television) corpus available through Korp (Borin et al., 2012). These sentences were then uploaded into the tool, automatically paired and annotated with XL-Lexeme, and finally clustered using the correlation algorithm. The resulting clusters were evaluated against lexicographer judgments for sense clusters (gold data).

Table 3 displays the evaluation results using the Adjusted Rand Index (ARI, Hubert and Arabie, 1985) to compare automatically generated clusters with manually curated gold clusters. Each row represents a specific word used to form the clusters, with the corresponding ARI value indicating the similarity between the automatic and gold clusters generated by DUREl and the lexicographer respectively. ARI generally ranges from -0.5 to 1, where higher values closer to 1 signify better agreement between the gold clusters and the derived clusters. A value of 0 suggests a random clustering. As can be noted words like *ofantlig*, *enkelspårig*, *baksida*, and *bagage* demonstrate relatively high ARI values, indicating stronger alignment between their respective clusters and the gold clusters. In contrast *tvärnita*, *hemmaplan* and *vansinnig* show ARI values of 0 or slightly negative. The ‘Average’ row




Figure 2: Clusters for *bagage*.

provides an overall assessment by presenting the average ARI across all words. The moderately positive average value suggests that the automatically derived clusters generally encode meaningful semantic information which could be useful for lexicographic work.

In addition to automatic evaluation using ARI, the clusters were analyzed qualitatively by a lexicographer, which provided useful insights. For example consider Figure 2 which shows the clusters produced for the word *bagage* (Eng: ‘luggage’). As can be seen, the uses were clustered into three main clusters (colored blue, orange, and green). The blue and orange colored represent the literal and figurative usage of the word respectively while one of the uses was wrongly put into the third, green cluster. A discrepancy between the number of clusters identified by DUREl and the number of senses recorded in SO dictionary for the candidate word can be used as an indicator that the description in the dictionary is outdated. Such words can be prioritized for manual inspection in order to update the dictionary entry or otherwise to improve the computational model. An examination of the figurative examples (the orange cluster) revealed that some of the uses can be classified as different variants of the same Swedish idiomatic expression (e.g. *ha något i bagaget*, ENG: ‘to have something in the luggage’). This particular idiom, and similar others, is not treated in SO 2021 and will be included in the next edition. Further, the text examples from the corpora form also a good basis to include more language examples in the SO dictionary. In summary, the DUREl tool was evaluated by the lexicographer to be helpful in revising the SO dictionary, who otherwise has relied completely on manual cumbersome methodology.

## 6 Conclusion

We presented the DUREl, an online, open-source annotation tool for annotating semantic proximity between word uses. The tool supports human as well as computational annotation, building on recent advances with WiC models. Annotator judgments are automatically clustered and visualized for analysis. This allows to capture word senses with simple and intuitive micro-task judgments between use pairs, requiring minimal preparation efforts. Additionally, DUREl can compare the agreement between annotators to guarantee the intersubjectivity of the obtained judgments and to calculate summary statistics over the annotated data giving insights into sense frequency distributions, semantic variation or changes of senses over time. The computational annotator component allows to generate word sense clusters for large sets of words and word uses, making it possible, for instance, to analyze large amounts of data or to search systematically for new senses.

A number of different research groups around the world have used the system to annotate data (e.g. Zamora-Reina et al., 2022; Kutuzov et al., 2022; Aksenova et al., 2022; Chen et al., 2023) and continue to do so: Currently, the tool is being used to annotate an Italian dataset.

### 6.1 Limitations

DUREl provides a range of functionalities feasible only because it focuses on a single task but with the disadvantage of a more narrow application range and low customization. This also binds users to the four-level annotation scale hard-coded into the system. While this scale is motivated by theory (Schlechtweg et al., 2018), there exist valid alternatives (e.g. Erk et al., 2013; Brown, 2008). For users who want to use such an alternative scale, we recommend to conduct the annotation study within a more customizable annotation system like PhiTag. The annotated data can then be uploaded to DUREl with the "Upload Judgments" functionality, and clustering and analysis can be applied within DUREl.

The tool currently allows to specify a continuous substring of any length as target string in each uploaded word use. Discontinuous target strings, as needed e.g. for discontinuous multi-word expressions such as particle verbs in German, are currently not allowed.

The computational annotators are currently re-

stricted to predict binary labels. In the future, we will provide WiC models optimized for ordinal predictions on all levels of the annotation scale (Zhang, 2023).

### Acknowledgments

The development of the tool has been funded by the project 'Towards Computational Lexical Semantic Change Detection' supported by the Swedish Research Council (2019–2022; contract 2018-01184), the research program 'Change is Key!' supported by Riksbankens Jubileumsfond (under reference number M21-0021) and the CRETA center funded by the German Ministry for Education and Research (BMBF). We thank Anne Reuter, Annalena Streichert, Enrique Medina Waldo Castaneda, Sinan Kurtyigit, Serge Kotchourko, Pedro G. Bascoy, Francesco Periti and Jing Chen for contributions to this project.

### References

- Anna Aksenova, Ekaterina Gavrishina, Elisey Rykov, and Andrey Kutuzov. 2022. [Rudsi: graph-based word sense induction dataset for russian](#).
- Nikolay Arefyev, Maksim Fedoseev, Vitaly Protasov, Daniil Homskiy, Adis Davletov, and Alexander Panchenko. 2021. Deepmistake: Which senses are hard to distinguish for a word-in-context model. volume 2021-June, pages 16–30.
- Carlos Santos Armendariz, Matthew Purver, Matej Ulčar, Senja Pollak, Nikola Ljubešić, and Mark Granroth-Wilding. 2020. [CoSimLex: A resource for evaluating graded word similarity in context](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5878–5886, Marseille, France. European Language Resources Association.
- Gioia Baldissin, Dominik Schlechtweg, and Sabine Schulte im Walde. 2022. [DiaWUG: A Dataset for Diatopic Lexical Semantic Variation in Spanish](#). In *Proceedings of the 13th Language Resources and Evaluation Conference*, Marseille, France. European Language Resources Association.
- Nikhil Bansal, Avrim Blum, and Shuchi Chawla. 2004. [Correlation clustering](#). *Machine Learning*, 56(1-3):89–113.
- Andreas Blank. 1997. *Prinzipien des lexikalischen Bedeutungswandels am Beispiel der romanischen Sprachen*. Niemeyer, Tübingen.
- Lars Borin, Markus Forsberg, and Johan Roxendal. 2012. Korp – the corpus infrastructure of språkbanken. In *Proceedings of LREC 2012. Istanbul: ELRA*, volume Accepted, page 474–478.

- Susan Windisch Brown. 2008. Choosing sense distinctions for WSD: Psycholinguistic evidence. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Short Papers*, pages 249–252, Stroudsburg, PA, USA.
- Pierluigi Cassotti, Lucia Siciliani, Marco de Gemmis, Giovanni Semeraro, and Pierpaolo Basile. 2023. Xllexeme: Wic pretrained model for cross-lingual lexical semantic change. In *Proceedings of the 61th Annual Meeting of the Association for Computational Linguistics*, Online. Association for Computational Linguistics.
- Jing Chen, Emmanuele Chersoni, Dominik Schlechtweg, Jelena Prokic, and Chu-Ren Huang. 2023. ChiWUG: A graph-based evaluation dataset for chinese lexical semantic change detection. In *Proceedings of the 4th International Workshop on Computational Approaches to Historical Language Change*, Singapore. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Mark Davies. 2012. Expanding Horizons in Historical Linguistics with the 400-Million Word Corpus of Historical American English. *Corpora*, 7(2):121–157.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Katrin Erk, Diana McCarthy, and Nicholas Gaylord. 2013. Measuring word meaning in context. *Computational Linguistics*, 39(3):511–554.
- Mario Giulianelli, Marco del Tredici, and Raquel Fernández. 2020. Analysing lexical semantic change with contextualised word representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3960–3973, Online. Association for Computational Linguistics.
- Evelyn Gius, Jan Christoph Meister, Malte Meister, Marco Petris, Christian Bruck, Janina Jacke, Mareike Schumacher, Dominik Gerstorfer, Marie Flüh, and Jan Horstmann. 2022. [Catma](#).
- Anna Häty, Dominik Schlechtweg, and Sabine Schulte im Walde. 2019. [SUREl: A gold standard for incorporating meaning shifts into term extraction](#). In *Proceedings of the 8th Joint Conference on Lexical and Computational Semantics*, pages 1–8, Minneapolis, MN, USA.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Lawrence Hubert and Phipps Arabie. 1985. Comparing partitions. *Journal of Classification*, 2:193–218.
- Adam Kilgarriff. 1997. "I don't believe in word senses". *Computers and the Humanities*, 31(2).
- Adam Kilgarriff. 2007. *Word Senses*, chapter 2. Springer.
- Jan-Christoph Klie, Michael Bugert, Beto Boulosa, Richard Eckart de Castilho, and Iryna Gurevych. 2018. [The inception platform: Machine-assisted and knowledge-oriented interactive annotation](#). In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9. Association for Computational Linguistics. Event Title: The 27th International Conference on Computational Linguistics (COLING 2018).
- Andrey Kutuzov and Lidia Pivovarova. 2021a. Rushifteval: a shared task on semantic shift detection for russian. *Komp'yuternaya Lingvistika i Intellektual'nye Tekhnologii: Dialog conference*.
- Andrey Kutuzov and Lidia Pivovarova. 2021b. [Three-part diachronic semantic change dataset for russian](#).
- Andrey Kutuzov, Samia Touileb, Petter Mæhlum, Tita Enstad, and Alexandra Wittemann. 2022. [NorDiaChange: Diachronic semantic change dataset for Norwegian](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2563–2572, Marseille, France. European Language Resources Association.
- Diana McCarthy, Marianna Apidianaki, and Katrin Erk. 2016. Word sense clustering and clusterability. *Computational Linguistics*, 42(2):245–275.
- Jiaxin Pei, Aparna Ananthasubramaniam, Xingyao Wang, Naitian Zhou, Apostolos Dedeloudis, Jackson Sargent, and David Jurgens. 2022. Potato: The portable text annotation tool. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 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*, pages 2227–2237, New Orleans, LA, USA.



- Mohammad Taher Pilehvar and Jose Camacho-Collados. 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Dominik Schlechtweg. 2023. [Human and Computational Measurement of Lexical Semantic Change](#). Ph.D. thesis, University of Stuttgart, Stuttgart, Germany.
- Dominik Schlechtweg, Barbara McGillivray, Simon Hengchen, Haim Dubossarsky, and Nina Tahmasebi. 2020. [SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection](#). In *Proceedings of the 14th International Workshop on Semantic Evaluation*, Barcelona, Spain. Association for Computational Linguistics.
- Dominik Schlechtweg, Sabine Schulte im Walde, and Stefanie Eckmann. 2018. [Diachronic Usage Relatedness \(DURel\): A framework for the annotation of lexical semantic change](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 169–174, New Orleans, Louisiana.
- Dominik Schlechtweg, Nina Tahmasebi, Simon Hengchen, Haim Dubossarsky, and Barbara McGillivray. 2021. [DWUG: A large Resource of Diachronic Word Usage Graphs in Four Languages](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7079–7091, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- A. Soares da Silva. 1992. Homonímia e polissemia: Análise sémica e teoria do campoléxico. In *Actas do XIX Congreso Internacional de Lingüística e Filoloxía Románicas*, volume 2 of *Lexicología e Metalexicografía*, pages 257–287, La Coruña. Fundación Pedro Barrié de la Maza.
- Lukas Theuer Linke. 2023. Testing the effect of using crowdsourced semantic proximity judgments in the process of human lexicographical word sense clustering. Bachelor thesis, University of Stuttgart.
- Frank D. Zamora-Reina, Felipe Bravo-Marquez, and Dominik Schlechtweg. 2022. [LSCDiscovery: A shared task on semantic change discovery and detection in Spanish](#). In *Proceedings of the 3rd International Workshop on Computational Approaches to Historical Language Change*, Dublin, Ireland. Association for Computational Linguistics.
- Tuo Zhang. 2023. An ordinal formulation of the graded word-in-context task. Master thesis, University of Stuttgart.

## **A Appendix**

We have included some screenshots of the DUREl application below.

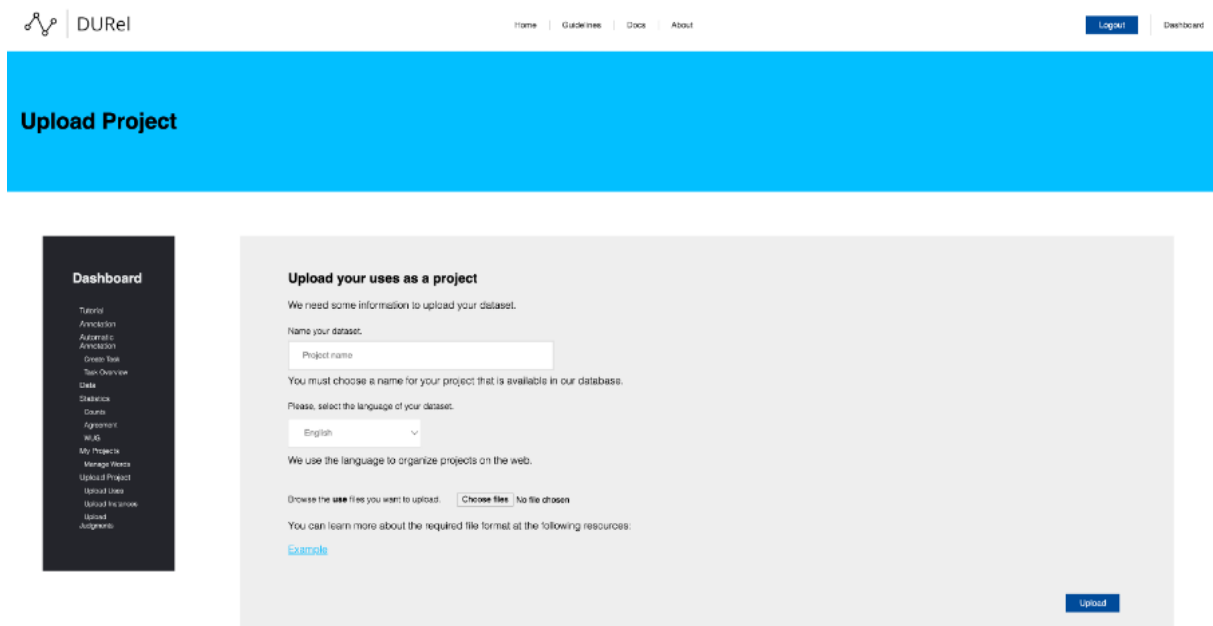


Figure 3: Upload uses tab: Upload interface for word uses.




Figure 4: My Projects tab: Assign access rights to annotators, delete or download projects.

The screenshot shows the 'Data' tab in the DUREl interface. On the left is a dark sidebar with a 'Dashboard' menu. The main content area is titled 'Uses' and contains a table with the following data:

ID	Lemma	POS	Date	Left Context	Target	Right Context
1200765	arm	N	1824	and taking a knife from her pocket, she opened a vein in her little	arm	, and dipping a feather in the blood, wrote something on a piece of white cloth, which was spread before her.
1200766	arm	N	1842	And those who remained at home had been heavily taxed to pay for the	arms	, ammunition, fortifications, and all the other endless expenses of a war.
1200767	arm	N	1890	and though he saw her within reach of his	arm	, yet the light of her eyes seemed as far off as that of a
1200768	arm	N	1903	It stood behind a high brick wall, its back windows overlooking an	arm	of the sea which, at low tide, was a black and stinking mud-flat
1200769	arm	N	1975	twelve miles of coastline lies in the southwest on the Gulf of Aquaba, an	arm	of the Red Sea. The city of Aquaba, the only port, plays
1200770	arm	N	1985	when the steamblock	arm	of the Statue of Liberty jets spectacularly out of the sandy beach.

Figure 5: Data tab: Shows concordances in a table.

The screenshot shows the 'Annotation' interface in the DUREl system. On the left is a dark sidebar with a 'Dashboard' menu. The main content area is titled 'Annotation (1/15)' and displays two sentences for annotation:

**Sentence 1**  
 And those who remained at home had been heavily taxed to pay for the **arms**, ammunition, fortifications, and all the other endless expenses of a war.

**Sentence 2**  
 and taking a knife from her pocket, she opened a vein in her little **arm**, and dipping a feather in the blood, wrote something on a piece of white cloth, which was spread before her.

Below the sentences is an 'Optional Comment' text box and a 'Keyboard shortcuts' icon. A 'Next' button is located at the bottom right of the annotation area.

Figure 6: Annotation interface: Annotation instance (use pair) presented to annotator.



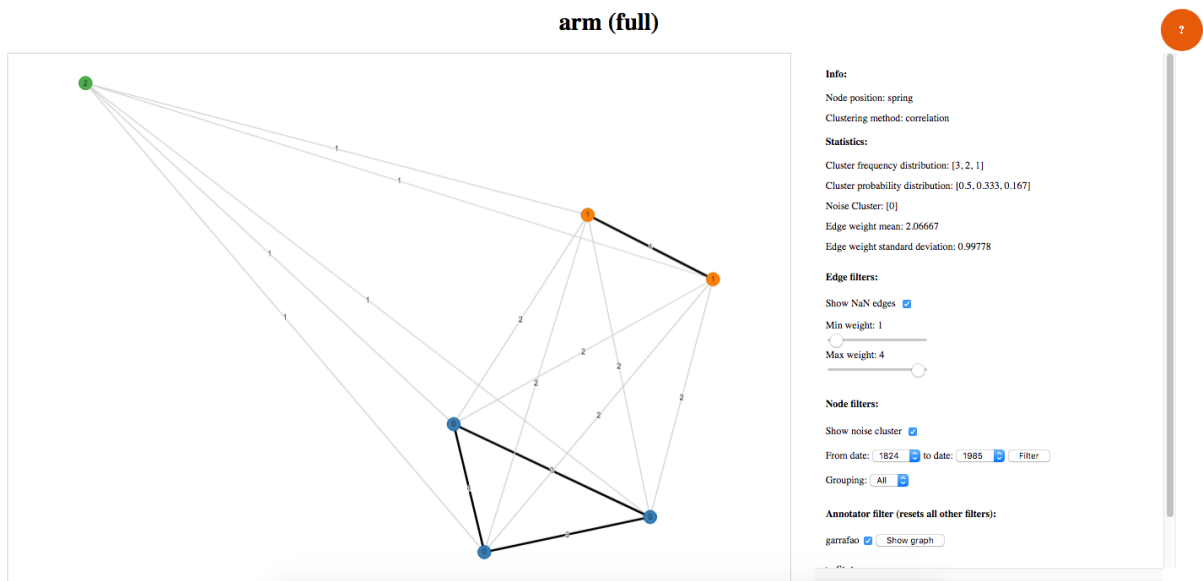


Figure 7: Visualization: Result of visualization pipeline available on "Statistics" tab.




Figure 8: Clustering obtained from human vs. computational annotations of *arm* in DUREL tool. Orange cluster corresponds to metaphorical sense 'arm of the sea' in both cases. Computational annotator merges 'body part' and 'weapon' sense.