

# Enriching Word Usage Graphs with Cluster Definitions

Mariia Fedorova<sup>1</sup>, Andrey Kutuzov<sup>1</sup>, Nikolay Arefyev<sup>1</sup>, Dominik Schlechtweg<sup>2</sup>

<sup>1</sup>University of Oslo, <sup>2</sup>University of Stuttgart  
{mariiaf, andreku, nikolare}@ifi.uio.no  
dominik.schlechtweg@ims.uni-stuttgart.de

## Abstract

We present a dataset of word usage graphs (WUGs), where the existing WUGs for multiple languages are enriched with cluster labels functioning as sense definitions. They are generated from scratch by fine-tuned encoder-decoder language models. The conducted human evaluation has shown that these definitions match the existing clusters in WUGs better than the definitions chosen from WordNet by two baseline systems. At the same time, the method is straightforward to use and easy to extend to new languages. The resulting enriched datasets can be extremely helpful for moving on to explainable semantic change modeling.

**Keywords:** word usage graphs, semantic change detection, definition generation

## 1. Introduction and related work

In natural languages most words are polysemous, i.e. the same word in different contexts can have different meanings. This gave rise to NLP tasks like word sense induction and word sense disambiguation. To solve them, the NLP community came up with many valuable resources related to word senses. One specific use case for such datasets is lexical semantic change modeling, where diachronic changes of word meaning are traced in an automated way.

The field of semantic change modeling makes a heavy use of the so called ‘Word Usage Graphs’ (WUGs) (Schlechtweg et al., 2021). Each WUG is associated with a particular target word, and is a weighted, undirected graph, with nodes corresponding to target word usages in a corpus, while edges are weighted with the semantic proximity of particular pairs of usages. WUGs are human-annotated, with annotators yielding contextualized graded judgments about the said semantic proximity. After the annotation is complete, graph nodes (target word occurrences) are automatically clustered into groups roughly corresponding to word senses. When a WUG contains target word occurrences from different time periods, it is called a *Diachronic Word Usage Graph* (DWUG).

DWUGs are available for several languages and are often used for evaluating semantic change detection and discovery systems (Schlechtweg et al., 2020; Kutuzov et al., 2022b). However, one of the features they lack is human interpretability of senses: clusters are labeled only with their numerical identifiers. To analyze a cluster from a DWUG, one has to browse through the actual target word usages, which is cumbersome and time-consuming. Giulianielli et al. (2023) addressed this issue and suggested a method to automatically

generate human-readable sense definitions (cluster labels) using an encoder-decoder language model fine-tuned on the task of contextualized definition generation (Mickus et al., 2022). However, their work was limited to the English DWUGs and they did not release the labels they produced.

In this paper, we apply this method to actually generate sense definitions for the available diachronic and synchronic WUGs in several languages (English, German, Norwegian, and Russian). We publicly release these enriched WUGs (mappings from clusters to definitions) with definitions both in English and in the WUG-specific languages.<sup>1</sup> In addition, we compare the performance of the LLM-based definition generation system with other approaches which choose a definition from an existing ontology like Wordnet or Wiktionary. Thus, our contribution is twofold: first, we make the existing word usage graphs more useful for linguists and lexicographers; second, we evaluate definition generation and aggregation methods on multiple languages and release the best fine-tuned models.

## 2. Data Description

### 2.1. Word usage graphs

The WUG repository (Dominik Schlechtweg and Sabine Schulte im Walde, 2023) features word usage graphs for eight languages. We chose English, German, Norwegian and Russian to experiment with. We ignore Chinese, Latin, Spanish and Swedish for the time being: mostly because of the lack of available datasets to fine-tune definition generators for these languages.

<sup>1</sup>[https://github.com/ltagoslo/wugs\\_with\\_definitions](https://github.com/ltagoslo/wugs_with_definitions)

Language	Target words	Clusters	Clusters annotated	Diachronic?
English (Schlechtweg et al., 2020)	46	819	120	True
German (Schlechtweg et al., 2020)	50	488	95	True
Norwegian-1 (Kutuzov et al., 2022a)	40	99	23	True
Norwegian-2 (Kutuzov et al., 2022a)	40	78	17	True
Russian (Aksenova et al., 2022)	24	90	39	False

Table 1: Main statistics of the enriched WUGs. We labeled and annotated only clusters which feature at least three usage examples.

Table 1 provides the main statistics of the word usage graphs we employ. Note that there are *two* Norwegian DWUGs with different target word sets and time periods (Norwegian-1 and Norwegian-2). Also note that for Russian we use the ‘*RuDSI*’ WUGs (Aksenova et al., 2022) which are not diachronic (target word usages may come from different time period, but it is not specially accounted for). There exist two semantic change detection datasets for Russian (‘*RuSemShift*’ and ‘*RuShiftEval*’), but due to their annotation procedure, they do not feature any meaningful clusters to label.

We did not produce labels for clusters with less than three example usages (this excluded a large amount of singleton clusters). The reason is that many of such ‘small’ groups are clustering errors, and even if they are not, it is extremely difficult to come with a good definition based on one or two examples; for the DefGen method (subsection 3.3), generating a definition for a cluster with two examples is ill-defined. We leave labeling this long tail of small clusters to future work.<sup>2</sup>

## 2.2. Definition datasets

To be able to generate contextualized dictionary-like definitions, a pre-trained encoder-decoder language model has to be fine-tuned on a dataset of definitions coupled to example target word usages. In addition, one has to choose a prompt which serves as an instruction for the language model to do the correct task. We follow the logic of Giulianelli et al. (2023) in choosing both the prompt (‘*What is the definition of TARGET\_WORD?*’ and its translations into corresponding languages) and the fine-tuning datasets for English: **WordNet** (Ishiwatari et al., 2019), **Oxford** (Gadetsky et al., 2018) and **CoDWoE** (Mickus et al., 2022) (all CoDWoE datasets originally come from Wiktionary). For other languages we used the following resources:

- Russian: the corresponding part of **CoDWoE**;

<sup>2</sup>Also, in some original DWUGs, there are clusters labeled with ‘-1’, which means annotators were unsure about semantic proximity for nodes within these clusters most of the time; we ignored these clusters as well.

- Norwegian: **Bokmålsordboka** (The Norwegian Language Council and the University of Bergen, 2023).

We are not aware of any readily available definition dataset for German; however, we still decided to assign cluster labels to German DWUGs as an experiment in zero-shot cross-lingual definition generation (also, we had available resources for human evaluation of German cluster labels). See 3 for more details on definition generation for different languages.

Table 2 shows the statistics of the definition datasets for fine-tuning.

## 3. Definition generators

We present three methods that can enrich DWUG clusters with definitions. Two of them are our baselines: they select a definition for each cluster from a human-curated lexical ontology, the English WordNet in our case. The third one generates definitions from scratch and is our main method.

### 3.1. Lesk

The simplest method uses the Lesk algorithm (Lesk, 1986) originally developed to solve the Word Sense Disambiguation (WSD) task. For a target word in a given context, it browses all the definitions listed for this word in a glossary or some WordNet-like resource and picks up a definition with the highest lexical overlap with the given context. To construct a context for our task we concatenate all usages from the same cluster, thus selecting the most suitable definition for this particular cluster. The NLTK (Bird et al., 2009) implementation of the Lesk algorithm was employed, together with WordNet as the source of senses and their definitions. We used the pre-tokenized versions of WUG examples, while WordNet definitions were split into tokens by whitespaces. The part-of-speech tags of the target words (where available) were taken into account when selecting synsets. This method is doomed to fail when the target word occurs in a sense that is not listed among senses of this target word in WordNet, or if the target word

Dataset	Entries	Lemmas	Ratio	Usage Length	Definition Length
English	175 332	49 238	3.56	31.64 $\pm$ 23.82	17.56 $\pm$ 11.13
Norwegian	70 711	30 551	2.31	8.69 $\pm$ 3.40	12.25 $\pm$ 7.35
Russian	72 872	34 495	2.11	52.03 $\pm$ 29.37	24.65 $\pm$ 15.84

Table 2: Definition datasets. Average usage and definition lengths are given in `mT0` sub-words.

is simply absent in WordNet. Also, it does not work with languages other than English unless using lexical resources for those languages. In our experiments with the English DWUGs, we observed very low accuracy compared to other methods even for English which has well-developed lexical resources (see Section 4). Thus, we decided not to adapt Lesk to other languages.

### 3.2. GlossReader

GlossReader (Rachinskiy and Arefyev, 2021) is a system originally developed to produce the contextualized embeddings that were shown to outperform standard LLM embeddings in lexical semantic tasks such as Multilingual and Cross-lingual Word-in-Context (Martelli et al., 2021) and Lexical Semantic Change Discovery (Zamora-Reina et al., 2022). The system is based on the English WSD model by (Blevins and Zettlemoyer, 2020), which consists of a definition (gloss) encoder and a context encoder, both initialized with the English BERT weights and jointly fine-tuned on a WSD dataset to maximize the dot product between a context embedding and a definition embedding of the corresponding word sense. In GlossReader, the English BERT backbone was replaced by the multilingual `XLM-R` (Conneau et al., 2020) in both encoders, and it was shown that fine-tuning `XLM-R` as part of such WSD system significantly improves the performance of its contextualized embeddings for several lexical semantic tasks in various languages even when fine-tuning on the same English WSD dataset only (Rachinskiy and Arefyev, 2021, 2022). In this work, we first build gloss embeddings for all WordNet definitions (about 117K) with the gloss encoder. Then, for each word usage we build the contextualized embedding of the target word and retrieve  $k$  most similar definition embeddings<sup>3</sup> as measured by the dot product. Finally, for each cluster we select the definition retrieved for the largest

<sup>3</sup>The pairwise human annotations from DWUGs were employed to select the optimal value of  $k$  among 1,3,10. For each DWUG and each  $k$  we built a definition for each cluster, then assigned cluster definitions to all of its usages and estimated the probability that two usages with entirely different senses (the annotation of 1) obtain the same definition. We selected  $k = 3$  for English DWUG and  $k = 10$  for other DWUGs to minimize this probability.

number of usages in this cluster.

Unlike Lesk, this method selects from *all* definitions in WordNet, thus, it can produce reasonable definitions even for senses not listed for the target word in WordNet. Moreover, due to zero-shot cross-lingual transferability of the `XLM-R` backbone, GlossReader can produce reasonable embeddings for definitions and word usages in various languages, and retrieve definitions in one language for word usages in another language. In this work, we use GlossReader to retrieve definitions in English for word usages in several other languages (cross-lingual setup), but we leave experiments with retrieving definitions in other languages (multi-lingual setup) for the future work.

### 3.3. DefGen

We dub ‘DefGen’ our main system, which takes as an input a text string containing a target word usage, and generates a human-readable definition of the target word in this particular context from scratch. We use the method proposed by Giulianelli et al. (2023): an encoder-decoder language model is fine-tuned on a dataset of target word usages and the corresponding definitions (see the subsection 2.2). Then, definitions are conditionally generated for every example in the test set (in our case, WUGs). Giulianelli et al. (2023) used `Flan-T5` (Chung et al., 2022) as the underlying language model. However, it was trained predominantly on English and lacks the capability to properly encode or generate texts in languages with significantly different writing systems (especially true for Russian, but some German and Norwegian characters are also not processed by the `Flan-T5` tokenizer). Because of that, in this work we are using `mT0` (Muennighoff et al., 2023) which is essentially a multilingual version of `Flan-T5` (also fine-tuned on many datasets cast as natural language instructions). For all the experiments, we employ the `mT0-xl` version of the model<sup>4</sup>, 3.7B parameters in size.

Fine-tuning was done in a standard text-to-text setup, for every language (English, Norwegian, Russian) separately, so that in the end we had three language-specific models. However, the

<sup>4</sup><https://huggingface.co/bigscience/mT0-xl>

underlying model is multilingual, so e.g., the English DefGen can still take sentences in other languages as input, and produce target word definitions in English with the standard English prompt. We used it to generate English definitions for the German DWUGs. Here, the input to the model looked like *‘Ist eine Prüfung erforderlich, so erfolgt eine Entscheidung über den Antrag durch die zuständige Behörde. What is the definition of Entscheidung?’* (*‘If an examination is necessary, a decision on the application will be made by the responsible authority. What is the definition of Decision?’*)<sup>5</sup>. Note that the models were not specifically fine-tuned on producing English definitions for German examples: this is a zero-shot multilingual ability probably stemming from the large-scale instruction fine-tuning of the base model on other tasks.

Table 3 shows the raw performance of the fine-tuned models with English and language-specific (‘native’) prompts as ROUGE-L (Lin, 2004) on the validation set. Note that here we report only the performance for definitions generated in the same language as the original examples (we don’t have any datasets with examples in one language and definitions in another to evaluate against). ROUGE-L measures only surface form overlap, so it should not be considered as the definitive metric. Still, the performance of our definition generators for English and Norwegian is on par with the English results reported in Giulianelli et al. (2023). The lower score on the Russian CodWoE dataset is most probably related to the abundance of dictionary-specific abbreviations, to morphological richness of Russian and to the fact that the reference Russian definitions are longer on average than those for other languages, as well as usage examples (Table 2). This is why the lengths of the longest common sub-sequence present in the reference definition and the generated definition are lower on average. Still, the definitions generated by the Russian model are mostly relevant and meaningful, as confirmed with visual examination by native Russian speakers and with the human evaluation experiments in Section 4. It is also worth noting that the performance when using English or language-specific prompts is not significantly different. In all the subsequent experiments, we use native prompts when generating language-specific definitions.

Once the definitions for all the examples of a specific WUG cluster in the desired language are generated, we again follow Giulianelli et al. (2023) in choosing *the most prototypical definition* as the cluster label. All the generated definitions are en-

<sup>5</sup>The mT0 answer in this case was *‘The act of making up your mind about something; a decision.’*

Language	English prompt	Native prompt
English	39.14	39.14
Norwegian	28.16	27.76
Russian	17.26	17.25

Table 3: Performance of mT0-based definition generators (ROUGE-L \* 100) on the validation sets.

coded by a multilingual SBERT model<sup>6</sup> (Reimers and Gurevych, 2019) and the most prototypical definition is trivially the one with the vector representation closest (by cosine similarity) to the average of all the definition embeddings in the cluster. No search for hyperparameters was conducted, we used the default values from Giulianelli et al. (2023) at all the DefGen stages.

In this way, we generated English cluster/sense labels for all four languages. In addition, language-specific labels for Norwegian and Russian were generated. In the next section, we evaluate them using human judgments.

## 4. Evaluation and results

### 4.1. Evaluation setup

Our ultimate aim is to assign each WUG cluster a human-readable label which is distinct enough to be useful for a linguist or lexicographer. In other words, the label should be helpful in distinguishing one cluster from another<sup>7</sup>. To evaluate the labeling methods, we came up with the *‘guess the cluster by definition’* task for human evaluation, described below.

For every cluster of a target word, the annotators were shown the label generated for this cluster (without knowing what system it came from) and *two* clusters, represented with five randomly sampled example sentences each (or less, if the cluster has less examples in the WUG). One of those clusters is the one for which the system generated the label and another one is a randomly sampled filler cluster (the filler pairings were generated before the evaluation started and were used throughout the whole process). The humans were asked which cluster is the best fit for the shown label (definition). They also had the options ‘none of the clusters fits’ and ‘the definition fits both clusters’. Only one choice was allowed. Importantly, the annotators were instructed that the label fits the cluster if it fits the majority of the examples in it. Otherwise, fluency or factual correctness of

<sup>6</sup><https://huggingface.co/sentence-transformers/distiluse-base-multilingual-cased-v1>

<sup>7</sup>The clusters themselves can be incorrect, but their fixing is out of scope for this paper.

the definitions were not taken into account: they only had to be distinctive enough to tell one cluster from another. Our final evaluation metrics is accuracy for all the clusters, where the system label is considered to be ‘correct’ if and only if the human chose as the best fit the same cluster that the system generated the label for. Our additional metrics are the percentages of ‘fits both’ and ‘fits none’ judgments (the less the better).

German, Russian and Norwegian predictions were evaluated by native speakers of the corresponding languages (the German speaker was familiar with the corresponding WUG data beforehand). English predictions were evaluated by the paper authors who are fluent English speakers. English definitions for the RuDSI dataset were evaluated by three independent annotators (using majority voting as the aggregation method). The resulting Krippendorff’s  $\alpha$  inter-rater agreement is 0.314, which is considered a fair agreement (Artstein and Poesio, 2008).

## 4.2. Experimental results

Table 4 shows the evaluation results of the systems for the WUGs we experiment with. Note that the task we are solving (generating labels or definitions for sets of sentences) is novel in itself, so there are no prior results to compare against. DefGen outperforms both Lesk and GlossReader in accuracy for all the datasets under comparison (English, German and Russian with English definitions). Interestingly, DefGen is slightly more prone to producing too general definitions which fit both clusters, but much more rarely produces irrelevant definitions which fit none of the clusters (for Lesk, more than half of the definitions are like this). This is probably the reason for the higher DefGen accuracy score.

For the rest of the WUGs (Norwegian with Norwegian and English definitions and RuDSI with Russian definitions), we evaluated only DefGen, to make sure its performance does not drop for some reason. The results for Norwegian are on the same level or higher, but the accuracy of Russian definitions for RuDSI are much lower than the English definitions generated with the same method. This partially stems from an increased number of ‘fits none’ judgments (the definition is irrelevant for both clusters). We observed low ROUGE-L performance of the Russian DefGen before (subsection 3.3), so it does not come as a surprise.

However, this degradation can be caused by some properties of a particular dataset, as is suggested by the difference between English and Norwegian definitions for the Norwegian clusters. As mentioned before, there are two sets of Norwegian DWUGs. For Norwegian-2, the English definitions are more accurate than the Norwegian ones (88.24

System	Accuracy	Fits both	Fits none
English DWUG, English definitions			
Lesk	21.67	5.00%	53.33%
GlossReader	50.00	9.17%	37.50%
DefGen	<b>69.17</b>	10.83%	11.67%
German DWUG, English definitions			
GlossReader	53.68	13.68%	27.37%
DefGen	<b>57.89</b>	16.84%	12.63%
RuDSI, English definitions			
GlossReader	64.10	10.26%	17.95%
DefGen	<b>71.79</b>	15.38%	2.56%
Norwegian-1 DWUG, English definitions			
DefGen	60.87	13.04	21.74
Norwegian-2 DWUG, English definitions			
DefGen	88.24	5.88%	5.88%
Norwegian-1 DWUG, Norwegian definitions			
DefGen	73.91	4.35%	21.74%
Norwegian-2 DWUG, Norwegian definitions			
DefGen	76.47	11.76%	11.76%
RuDSI, Russian definitions			
DefGen	48.72	7.69%	15.38%

Table 4: Results of human evaluation in the ‘guess the cluster by definition’ task.

vs 76.47), same tendency as for Russian. But for Norwegian-1, it’s vice versa (60.87 vs 73.91). Thus, the performance of English definitions for Norwegian examples varies greatly: can be excellent for one set of WUGs but poor for another. At the same time, the performance of Norwegian definitions is fairly good for both. We conjecture that the definitions in WUG-native languages can be more robust and less influenced by peculiarities of specific target words. If this is true, it would mean that RuDSI simply happened to be a ‘lucky’ set of WUGs for English definitions, same as Norwegian-2. To test this hypothesis, another Russian dataset is needed, so this is left for future work.

## 4.3. Error analysis

We have manually classified all erroneous definitions from the English, Norwegian 1 and Russian DWUG by mistake type, the results are shown in the tables 5, 6, 7.

One of the most widespread problems, which has caused most of ‘fits both’ annotations, is too broad definition. For example, the following def-

Error type	Share	Number
too broad definition	0.41	15
wrong sense	0.32	12
similar words but not a definition	0.16	6
redundant whitespace	0.08	3
ambiguous definition, ambiguous word	0.03	1
repetitions	0.05	2
factual mistake	0.03	1
non-existing word	0.03	1
Total	1.00	37

Table 5: Erroneous English definitions by error type. Some definitions were annotated with two error types, for instance, some too broad definitions also suffer from a redundant whitespace.

inition was generated for the word ‘relationship’: ‘*The way in which two or more things are connected, or in which two or more things are related.*’. This definition did not allow the annotator to distinguish the clusters containing usages related to relationship among entities or among people. The same word causes ‘fits both’ in Norwegian: ‘forhold’ (‘relationship’) is defined as ‘*forbindelse, samsvar mellom ulike faktorer*’ (‘connection, agreement between different factors’). A similar problem was also reported by the annotator of the German data: ‘*Concreteness vs. abstractness seems very important in descriptions: Sometimes the description could fit to both clusters, depending on whether it is read concretely or abstractly (metaphorically), e.g. ‘beimischen’, ‘to mix (something) together’. Mixing could be done e.g. by mixing chemicals or ideas or thoughts. Similarly with ‘abdecken: to cover or supply’.*

Error type	Share	Number
wrong sense	0.5	3
similar words but not a definition	0.17	1
too broad definition	0.17	1
repetitions	0.17	1
wrong preprocessing	0.17	1
Total	1.00	6

Table 6: Erroneous Norwegian definitions by error type. Some definitions were annotated with two error types.

Another source of too broad definitions is joining multiple meanings of a word into one definition by semicolon (there can be also a redundant space added before this semicolon; we consider such cases to be a punctuation error. A possible reason is that the training data were pre-tokenized and joined by a space). For example, ‘plane’ is defined as ‘*A flat surface, without slope, tilts, or indentations ; a level surface.*’ which merges its everyday and mathematical meanings.

Some definitions use ambiguous words so that

the resulting sense description does not help to distinguish between clusters. For example, ‘*Something used to support or maintain an object, scene, etc.*’ for the word ‘prop’. The word ‘scene’ is also ambiguous and the annotator has selected a wrong cluster because of it. This problem arises even more clearly when the model is trained in other languages with an English prompt. The definitions are often simply direct translations of the target word into English with the same level of ambiguity. For example, the Russian word ‘сторона’ (‘side’) is defined as ‘one’s part, aspect, role’.

Some of ‘fits both’ annotations may be also explained by too granular clustering, but fixing it is beyond of the scope of this paper.

Another frequent problem are definitions that are good themselves, but describe another sense of the word, not the one used in their source clusters. Such definitions result in wrong or ‘fits none’ annotation. They are often generated by usages of the target lemma in phrasal verbs or in metaphorical meaning. For example, a cluster containing usages of ‘chef d’oeuvres’ (in the sense of ‘masterpiece’) has generated the definition ‘*A master of excellence.*’ for the target word ‘chef’.

Error type	Share	Number
wrong sense	0.5	10
repetitions	0.45	9
similar words but not a definition	0.3	6
redundant whitespace	0.2	4
redundant note	0.2	4
too broad definition	0.15	3
redundant parentheses	0.05	1
Total	1.00	20

Table 7: Erroneous Russian definitions by error type. Some definitions were annotated with more than one error type.

The third most frequent problem are definitions that are semantically similar to what could be a good definition, but they barely describe the meaning of the cluster. Such cases are called ‘similar words but not a definition’ in our tables. An example is the definition of the word ‘gas’: ‘*Any fluid substance, especially a natural fluid, which is a mixture of air and nitrogen, usually created by the combustion of natural gasses.*’ It sounds much dictionary-like, but is meaningless and also contain a non-existing word ‘gasses’. The same problem was also reported for German, the annotator has often made his judgments based on ‘only some aspect that was not present in the other cluster’ in such cases. For Norwegian, there is a definition ‘*liten, sammenhengende fjordformasjon*’ (‘small, connected fjord formation’) generated for the word ‘rev’ (‘reef’). Not only is this definition an example of over-fitting to the Norwegian realities, but also,

as the annotator has pointed out, ‘fjordformasjon’ means ‘a formation of the actual fjord, like how the water moves and the sides of the fjord, not the ground in the water in the fjord’. Sometimes the meaninglessness of the definition is produced by repetitions in it; for example, ‘pin’ is defined as ‘*To fasten with a pins or pins.*’

Sometimes the definitions with a wrong meaning still carry the sentiment or the style of their source cluster. For example, the definition ‘*Bravery, courage, chutzpah, or brazenness.*’ containing a colloquial word ‘chutzpah’ was generated for the word ‘ball’ in its colloquial meaning. For German, the annotator could assign the definition to one of the clusters only based on the sentiment while the definition content does not fit otherwise: ‘*Abgesang: an expression of disapproval, a denial, or rejection. The negative sentiment fits very well with the "downfall" meaning of the word.*’

Factual mistakes are an important problem. For example, ‘heel’ is defined as ‘*The part of the human foot above the ankle.*’, while it is the part under the ankle. The annotation guidelines<sup>8</sup> included the instruction to ignore factual mistakes if they allow to distinguish between the clusters, so in fact this type may be not as rare as our current analysis shows. For German, the annotator reported that ‘*the meaning expressed an antonym or contrast of one of the clusters: e.g. 'Schmiere: a highly successful theatrical production' is exactly the opposite: a very bad theatrical play. I judged it still as fitting. Similarly for: Engpass: the property of a more than adequate quantity or supply; Spielball: involving active participation; Missklang: any agreeable (pleasing and harmonious) sounds.*’

Some errors are specific for the Russian training dataset and can be fixed by cleaning the dataset. These are the Wiktionary artifacts: redundant footnotes in parentheses (e.g. ‘стороны [ 1 ], участники какого-либо конфликтного процессуального действия’) and redundant dictionary notes that seem to occur randomly, e.g. ‘только полн. ф.’ (‘full form only’), which makes sense for adjectives and participles only, but is generated in the definition of a noun.

To conclude, the definitions produced by the model are not yet ready-to-use. They are in general suitable for semantic change detection or other semantic-related NLP tasks, but not for the full replacement of a human lexicographer.

More definition examples with an attempt to classify them by the error severity into good, borderline and bad ones can be found in [Appendix A](#).

---

<sup>8</sup>[https://github.com/lrgoslo/wugs\\_with\\_definitions/blob/main/human\\_evaluation/evaluation\\_guidelines.pdf](https://github.com/lrgoslo/wugs_with_definitions/blob/main/human_evaluation/evaluation_guidelines.pdf)

## 5. Conclusion

We have experimented with several methods for assigning human-readable cluster labels to word usage graphs (both synchronic and diachronic) for multiple languages. Human evaluation showed that conditional generation of such labels from scratch using a fine-tuned `mt0` language model outperforms two baselines which choose a label from an external closed set of possible definitions (senses). This ‘DefGen’ method generally follows the prior work, but we additionally show that it can generate reasonable labels even for languages on which the model was not specifically fine-tuned. Our definition generation models for English<sup>9</sup>, Norwegian<sup>10</sup> and Russian<sup>11</sup> are publicly available.

We release enriched versions of the existing WUGs for English, German, Russian and Norwegian. In these versions, clusters of usages are accompanied with labels generated by DefGen, both in English and in the native languages of the datasets. It is our hope that this makes the WUGs easier to use and analyze for linguists and lexicographers: it is now not necessary to browse through all the examples in a cluster to find out what sense of the target word does it correspond to. It is even more important for diachronic WUGs, where senses are changing over time and must be compared historically. Assigning DWUG clusters human-readable definitions is another step towards explainable semantic change modeling, thus bridging the gap between NLP and humanities.

We believe our findings support the claim of [Giulianelli et al. \(2023\)](#) that generated definitions can be used as convenient and interpretable contextualized representations for a wide range of domains. One can think of using such ‘definitions as representations’ in WSD, WSI or sentiment analysis, to name only a few NLP tasks.

As a future work, we plan to experiment with a DefGen model fine-tuned on definition datasets in several languages at once and with further fine-tuning such a model on a very small amount of manually created examples in other languages. Our preliminary experiments with German are promising (the model is able to generate definitions in German, despite the lack of proper fine-tuning dataset), but more work is needed. It is also planned to generate cluster labels for other existing WUGs not yet covered in this paper.

---

<sup>9</sup><https://huggingface.co/lrgoslo/mt0-definition-en-xl>

<sup>10</sup><https://huggingface.co/lrgoslo/mt0-definition-no-xl>

<sup>11</sup><https://huggingface.co/lrgoslo/mt0-definition-ru-xl>

## 6. Acknowledgments

Andrey Kutuzov and Nikolay Arefyev have received funding from the European Union's Horizon Europe research and innovation program under Grant agreement No 101070350 (HPLT). Dominik Schlechtweg has been funded by the research program 'Change is Key!' supported by Riksbankens Jubileumsfond (under reference number M21-0021). We acknowledge the help of Petter Mæhlum in evaluating Norwegian definitions and Pavel Suvorkov in evaluating Russian definitions.

## References

- Anna Aksenova, Ekaterina Gavrishina, Elisei Rykov, and Andrey Kutuzov. 2022. [RuDSI: Graph-based word sense induction dataset for Russian](#). In *Proceedings of TextGraphs-16: Graph-based Methods for Natural Language Processing*, pages 77–88, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Ron Artstein and Massimo Poesio. 2008. [Survey article: Inter-coder agreement for computational linguistics](#). *Computational Linguistics*, 34(4):555–596.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. "O'Reilly Media, Inc."
- Terra Blevins and Luke Zettlemoyer. 2020. [Moving down the long tail of word sense disambiguation with gloss informed bi-encoders](#). In *Proceedings of the 58th Association for Computational Linguistics*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- 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.
- Artyom Gadetsky, Ilya Yakubovskiy, and Dmitry Vetrov. 2018. [Conditional generators of words definitions](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 266–271, Melbourne, Australia. Association for Computational Linguistics.
- Mario Giulianelli, Iris Luden, Raquel Fernandez, and Andrey Kutuzov. 2023. [Interpretable word sense representations via definition generation: The case of semantic change analysis](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3130–3148, Toronto, Canada. Association for Computational Linguistics.
- Shonosuke Ishiwatari, Hiroaki Hayashi, Naoki Yoshinaga, Graham Neubig, Shoetsu Sato, Masashi Toyoda, and Masaru Kitsuregawa. 2019. [Learning to describe unknown phrases with local and global contexts](#). 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 3467–3476, Minneapolis, Minnesota. Association for Computational Linguistics.
- Andrey Kutuzov, Samia Touileb, Petter Mæhlum, Tita Enstad, and Alexandra Wittemann. 2022a. [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.
- Andrey Kutuzov, Erik Velldal, and Lilja Øvrelid. 2022b. [Contextualized embeddings for semantic change detection: Lessons learned](#). In *Northern European Journal of Language Technology, Volume 8*, Copenhagen, Denmark. Northern European Association of Language Technology.
- Michael Lesk. 1986. Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. In *Proceedings of the 5th annual international conference on Systems documentation*, pages 24–26.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Federico Martelli, Najla Kalach, Gabriele Tola, and Roberto Navigli. 2021. [SemEval-2021 task 2: Multilingual and cross-lingual word-in-context disambiguation \(MCL-WiC\)](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 24–36, Online. Association for Computational Linguistics.



- Timothee Mickus, Kees Van Deemter, Mathieu Constant, and Denis Paperno. 2022. [Semeval-2022 task 1: CODWOE – comparing dictionaries and word embeddings](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1–14, Seattle, United States. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual generalization through multitask finetuning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111, Toronto, Canada. Association for Computational Linguistics.
- Maxim Rachinskiy and Nikolay Arefyev. 2021. [GlossReader at SemEval-2021 task 2: Reading definitions improves contextualized word embeddings](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 756–762, Online. Association for Computational Linguistics.
- Maxim Rachinskiy and Nikolay Arefyev. 2022. [GlossReader at LSCDiscovery: Train to select a proper gloss in English – discover lexical semantic change in Spanish](#). In *Proceedings of the 3rd Workshop on Computational Approaches to Historical Language Change*, pages 198–203, Dublin, Ireland. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- 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 Fourteenth Workshop on Semantic Evaluation*, pages 1–23, Barcelona (online). International Committee for Computational Linguistics.
- 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.
- 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 Workshop on Computational Approaches to Historical Language Change*, pages 149–164, Dublin, Ireland. Association for Computational Linguistics.
- Dominik Schlechtweg and Sabine Schulte im Walde. 2023. [Word Usage Graphs \(WUGs\)](#). University of Stuttgart.
- The Norwegian Language Council and the University of Bergen. 2023. [Bokmålsordboka](#). University of Bergen.

## 7. Language Resource References

### A. Appendix

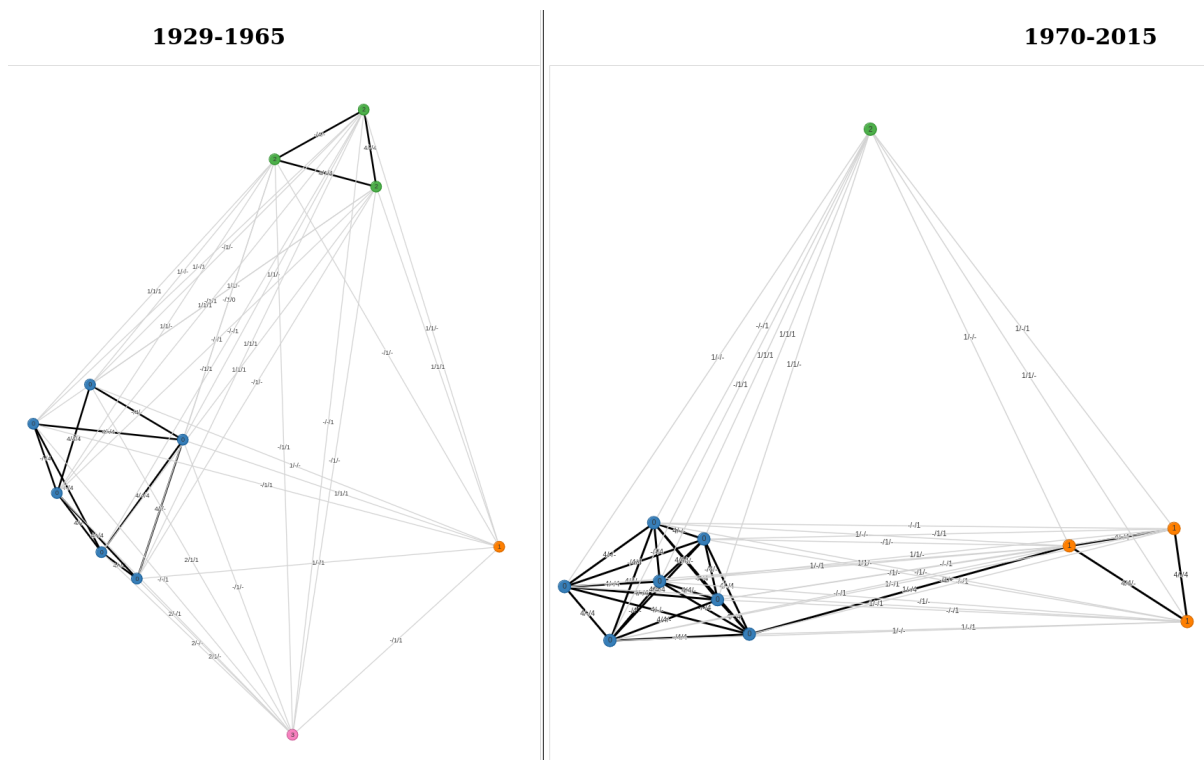


Figure 1: Diachronic word usage graph of the Norwegian word 'rev' ('fox'). The blue cluster contains the usages with the sense 'fox', the green cluster the usages with the sense 'strip of sail', the orange one the usages with the sense 'reef' and the rose cluster one usage with the sense 'card play'. It can be seen that the sense 'card play' no longer exists in the second time period, and the sense 'strip of sail' has become rare.

word	DWUG	sense	prompt	definition	translation
Entscheidung	German	decision	What is the definition of Entscheidung?	The act of making up your mind about something; a decision.	
мир	Russian	peace	Что такое мир?	полит. соглашение между сторонами о взаимных отношениях	in politics: agreement between two sides about their mutual relationships

Table 8: Good definition examples

word	DWUG	sense	prompt	definition	translation	error
rev	Norwegian-1	fox	Hva betyr rev?	rovdyr som blir brukt som inntektskilde	predator that is used as a source of income	too broad definition
rev	Norwegian-1	reef	What is the definition of rev?	A reefs, especially coral reef walls.		grammar mistake
мир	Russian	community of people	Что такое мир?	перен. совокупность людей, объединённых общностью интересов, интересов, а также их общая совокупность	metaphorical: a community of people with shared interests, interests, and also their entirety	repetitions
plane	English		What is the definition of plane?	An aircraft.		too broad definition

Table 9: Borderline definition examples

word	DWUG	sense	prompt	definition	translation	error
rev	Norwegian-1	fox	What is the definition of rev?	A rabbit.		factual mistake
rev	Norwegian-1	strip of sail	Hva betyr rev?	særlig skrevet særlig skrevet ...	-, meaningless repetition of random words	wrong preprocessing or bad generation parameters
rev	Norwegian-1	strip of sail	What is the definition of rev?	A twisting blow ; a punch.		similar words but not a definition

Table 10: Bad definition examples