

RuDSI: graph-based word sense induction dataset for Russian

Anna Aksenova **Ekaterina Gavrishina, Elisey Rykov** **Andrey Kutuzov**
National Research University Higher School of Economics National Research University Higher School of Economics University of Oslo
Russia Russia Norway
a.aksenova@hse.ru eigavrishina@edu.hse.ru andreku@ifi.uio.no
esrykov@edu.hse.ru

Abstract

We present RuDSI, a new benchmark for word sense induction (WSI) in Russian. The dataset was created using manual annotation and semi-automatic clustering of Word Usage Graphs (WUGs). Unlike prior WSI datasets for Russian, RuDSI is completely data-driven (based on texts from Russian National Corpus), with no external word senses imposed on annotators. Depending on the parameters of graph clustering, different derivative datasets can be produced from raw annotation. We report the performance that several baseline WSI methods obtain on RuDSI and discuss possibilities for improving these scores.

1 Introduction

Word sense induction (WSI) is among the most challenging problems in computational linguistics. The difficulty lies not only in the character of the task itself but also in the lack of datasets properly designed for it. We have developed such a dataset for the Russian language by means of manual annotation and clustering of the obtained senses. We dub it *Russian Data-driven Sense Induction* dataset (RuDSI)¹. Its annotation was based on so-called Word Usage Graphs (WUGs), where word usages in context are nodes connected by edges with weights corresponding to semantic proximity (Schlechtweg et al., 2020). This workflow has been already used to create diachronic semantic change datasets for Russian (Rodina and Kutuzov, 2020; Kutuzov and Pivovarova, 2021), but it is the first time it is employed for designing synchronic WSI benchmarks.

Graphs representing semantic relations between word usages were crucial for the creation of RuDSI. Communities or clusters induced from these graphs correspond to lexical senses; the number and composition of clusters for each word depends not only

¹<https://github.com/kategavrishina/RuDSI>

on human annotation, but also on the particular clustering procedure. Since we provide raw annotators' judgments, other researchers can apply their preferred graph processing techniques and obtain slightly different sense assignments.

The rest of the paper is organized as follows. In Section 2, we talk about the WSI datasets created earlier and their limitations. In Section 3, we present and analyze RuDSI and describe our annotation workflow. In Section 4, we show how graph clustering parameters affect the dataset. Section 5 reports the performance of several baseline WSI methods. In Section 6, we describe to whom and how RuDSI will be useful.²

2 Related work

In this section, we give a brief overview of word sense induction datasets for English developed as a part of SemEval competition, take a look at RUSSE'18 dataset and discuss the approaches towards WSI dataset creation.

2.1 SemEval datasets

Existing sense-annotated corpora like SemCor (Miller et al., 1993) allow for building competitive word sense disambiguation (WSD) models since they provide sufficient amount of training data. However, the major problem of such sources is the fact that word sense inventories vary depending on text domain and time period. Thus, WSD models are never universal. To solve this issue, word sense induction task was created. WSI systems aim to infer word senses from the given corpus.

In 2010, a WSI dataset was introduced during the SemEval competition (Manandhar et al., 2010). Compared to SemEval 2007 (Agirre and Soroa, 2007), it was more balanced in terms of nouns and

²This research was supported in part through computational resources of HPC facilities at HSE University (Kostenetskiy et al., 2021).

verbs distribution (50 verbs and 50 nouns in English). The main difference was in the evaluation procedure. The authors assumed that although WSI task is unsupervised, evaluating the methods on unseen test set of contexts would be more realistic. Different metrics for the clustering quality evaluation were inspected (V-measure, paired F-score) and all of them turned to be biased by number of senses predicted by WSI algorithms.

In 2013, another task setting was suggested by [Jurgens and Klapaftis \(2013\)](#). They claimed that there are contexts where multiple sense tags might be used. Therefore, the setup required predicting the weighted distribution of word senses for each context, i.e., perform *graded* word sense induction. To evaluate this task, two novel measures were introduced: fuzzy B-Cubed and fuzzy normalized mutual information. We should emphasize that our RuDSI dataset is aimed to test systems for *non-graded* word sense induction, although it could be transformed into graded setup (see Section 3.3).

2.2 Russian WSI datasets

Despite the fact that word sense induction task was well-developed for English, there were no manually annotated data for Russian until recently. In the last years, the interest to WSI and WSD tasks in Russian has increased due to the appearance of the first Russian WSI dataset. It was created as a part of RUSSE-18 shared task ([Panchenko et al., 2018](#)) and contains three subsets:

1. **wiki-wiki** (automatically extracted examples and senses from Wikipedia articles, mainly homonyms and homographs)
2. **bts-rnc** (examples from the Russian National Corpus (RNC), labeled with senses from the ‘Big Explanatory dictionary’)
3. **active-dict** (examples and senses from the ‘Active dictionary of the Russian language’ by Yuri Apresjan ([Apresjan, 2014](#)))

The training sections contained a total of about 17 thousand contexts. The key metric for the competition was Adjusted Rand Index (ARI) score ([Hubert and Arabie, 1985](#)). The Rand Index calculates the similarity between two clusterings by counting object pairs that were assigned the same or different clusters in golden labeling and in predictions. ARI adds adjustment for chance and gives score close

to 0 for random labeling and 1 for identical clusterings. When the clustering is worse than random, ARI is negative.

2.3 Limitations of previous datasets

Unfortunately, RUSSE-18 shared task data has a number of significant limitations. Linguistically, it includes homonyms, polysemous words and homographs, which does not correspond to the original WSI task setting: inducing senses of lexemes with the same set of word forms. In addition, some of the contexts in RUSSE-18 are noisy: there are cases where the target word is actually a root of a composite or a derivation (e.g., ‘луковица’ *bulb* is suggested as one of the words in context set for target word ‘лук’ *onion/bow*). The key issue is that word sense cannot be induced in these cases since derivations are mostly non-compositional and do not necessarily maintain the ambiguity relations of parent word. Finally, none of the target words of RUSSE-18 are monosemous, hence the dataset does not test WSI systems for polysemy detection, which is a critical issue in terms of developing a universal algorithm.

All the datasets for both Russian and English SemEval discussed above were automatically or manually tagged with dictionary-based sense inventories. We believe that it might be more realistic to derive word sense inventories for WSI pipelines evaluation not from linguistic sources, but directly from corpora, since the sets of senses vary in different corpora and domains ([Kilgarriff, 1997](#)).

2.4 Graph-based WSI datasets

A possible solution comes from combining word-in-context disambiguation and graph clustering. Conceptualization of semantic relationships as graphs empowered the approaches that represent the ambiguous lexeme as a central node in graph where nodes are the words and edge weights represent the measure of association between those words. [Hope and Keller \(2013\)](#) suggests calculating edge weights as a frequency measure for word co-occurrence similarity: the more similar are the contexts of the node lexemes, the higher will be the edge weight bridging them. Such co-occurrence graphs are calculated automatically. The similarity networks are afterwards clustered to induce word senses ([Hope and Keller, 2013](#); [Sherstuk, 2020](#)).

[McCarthy et al. \(2016\)](#) highlighted the problem of using fixed sets of senses for word sense inventory representation. Graphs used in the paper

were derived from word in context disambiguation annotation. They suggested treating annotators' judgements as graph edges and investigated different clusterability measures of such graphs.

Graph clustering has been successfully employed in annotating datasets for semantic change detection task (Schlechtweg et al., 2020, 2021). The annotation process is essentially word-in-context disambiguation: the annotators have to decide whether a pair of sentences represent the same target word sense or not. The annotation forms a *word usage graph* combining the uses from each pair of word contexts, where the nodes are the contexts themselves (sentences), and edges are weighted with the medians of annotators' judgements for a particular pair. Then, using correlation clustering, the graph is separated into clusters (communities of nodes) that correspond to the senses. The method is simple yet quite efficient as the annotators do not assign sense labels directly and the resulting clusters represent a set of data-driven senses³. Such a method does not only represent the relations between word usages, but also allows for choosing the granularity of the final word sense inventory. Moreover, the resulting senses are derived from data and not biased by lexicographic information; also, the number of clusters is determined automatically (Schlechtweg et al., 2021).

3 RuDSI dataset

3.1 Target words selection

To create RuDSI, it was first necessary to select a limited number of target words for further manual annotation. As we aimed at having words of different degree of polysemy presented in the final dataset, we extracted the total number of senses for each word in three distinct resources: Russian National Corpus (RNC)⁴, representative collection of texts in Russian with linguistic annotation; Wiktionary⁵, web-based free dictionary; and RuWordNet(Loukachevitch et al., 2016), a thesaurus of the Russian language created in the format of English WordNet (Miller, 1995). All non-noun words were discarded from this set.

³As opposed to dictionary-based senses, since the obtained senses are not taken from any resources, they are the result of automatic clustering.

⁴<https://ruscorpora.ru>; in particular, we used the RNC semantic markup (Rahilina et al., 2009) which includes parts of speech and semantic classes for a large number of lexemes (for example, *fruit/food* for the word 'apricot').

⁵<https://www.wiktionary.org>

Since the purpose of the annotation was to create a dataset with a balanced number of mono- and polysemous lexemes, we selected eight most frequent words (according to the dictionary by Lya-shevskaya and Sharov (2009)) in each of three groups: words with one sense, words with 2-4 senses (moderately polysemous), words with five or more senses (highly polysemous). The value of eight was chosen because of our limitations on the volume of annotation. The final number of senses was calculated as the average⁶ between RNC, Wiktionary and RuWordNet for each target word. Note that we did not consider these values as any sort of a gold standard, and they did not affect our human judgements in any way: annotators were not aware about the polysemy groups which the target words belonged to.

Thus, 24 target nouns were prepared for the annotation. For each word from the resulting set, 35 sentences containing this word were randomly sampled from the RNC. Next, annotators were given pairs of these sentences to estimate the relatedness of target word senses between each element in the pair.

3.2 Annotation

The annotation was performed using the DUREl web service⁷, which allows to annotate pairs of contexts for each word from the loaded sample. At each step of the annotation, a human is offered a pair of sentences to judge. For each pair, the columns 'Sentence 1' and 'Sentence 2' are presented with contexts containing the target word, which is highlighted in bold. The task is to assess how close in meaning the occurrences of the target word are in the two presented sentences on the scale from 1 (Unrelated) to 4 (Identical). The scores of 2 (Distantly Related) and 3 (Closely Related) are more subjective. In general, the 2 rating is for the uses that have different senses, but are somewhat related, and the 3 rating is for the cases when two uses have the same sense with some variation. So, a score of 1 is implied in the following example with the target word 'сторона' which is presented in the Figure 2, indicating that there is no connection between the senses (direct and figurative meaning of the lexeme):

⁶The average was preferable to minimum and maximum, since they would give more weight to one of the resources: in Wiktionary, words usually have few senses (1-2), but in RNC, same words can have a lot more senses (6 on average).

⁷<https://durel.ims.uni-stuttgart.de>

- (1) a. ‘При этом важны не только масштабы производства, но и его качественная сторона, то есть эффективное управление активами...’ (the meaning of ‘*component, element*’)
*At the same time, not only the scale of production is important, but also its qualitative **side**, that is, effective asset management...*
- b. ‘Так, донеся государю императору Александру о занятии Реймса, полки разошлись на пространстве от города вёрст до тридцати на квартиры в разные стороны.’ (the meaning of ‘*space, direction*’)
*So, having informed the Emperor Alexander about the occupation of Reims, the regiments dispersed in the space from the city to thirty versts to apartments in different **directions**.*

The next example presents the case of two uses with identical senses for the word ‘день’ (*day*) requiring the score of 4:

- (2) a. ‘Вещи не были еще расставлены, рамы были частью без стекол, частью с остатками расколотых, и (был дождливый день) с потолка текло.’
*Things were not yet arranged, the window frames were partly without glass, partly with the remains of splintered ones, and (it was a rainy **day**) the ceiling was flowing.*
- b. ‘Каждый день с раннего утра до обеда и с обеда до вечера я занят был работою или в доме, или в саду, или в огороде’
*Every **day**, from early morning to lunch and from lunch to evening, I was busy working either in the house, or in the garden, or in the vegetable garden.*

For each of the 24 words, as mentioned earlier, 35 sentences were sampled from the RNC. The DUREl platform automatically generated random sentence pairs, and at the first stage of our workflow, 180 pairs were annotated for each target word⁸. As a result, 24 separate word usage graphs with 35 nodes each were obtained.

⁸Annotation was performed by a subset of the authors of the article as native Russian speakers.

3.3 Aggregation of senses via graph clustering

Clustering of the sentences obtained as a result of the annotation for each lexeme was performed using the pipeline from (Schlechtweg et al., 2021) based on the variation of correlation clustering (Bansal et al., 2004; Schlechtweg et al., 2020). The DUREl relatedness scale from 1 to 4 was derived from continuum of semantic proximity (Blank, 1997): Homonymy - Proximity - Context Variance - Identity. Based on the continuum, the authors rescaled the annotators’ judgements for clusterization to represent the idea of usage pairs with 1 and 2 scores belonging to different senses, and with 3 and 4 scores — to the same sense. For this purpose they created the *threshold* parameter which was used to calculate the resulting edge weight: $W'(e) = W(e) - \text{threshold}$, and equated it to 2.5 (e.g., a score of 1 became -1.5). The division into clusters is based on the similarity between the target word senses within the sentences in a pair: clustering algorithm minimizes the sum of positive edge weights (3 and 4 scores in the original) across clusters and the sum of negative edge weights (1 and 2 scores) within clusters. Correlation clustering yields only one cluster label for a node (sentence), but by replacing it with a fuzzy graph clustering algorithm like the one in (Peng et al., 2021), it is possible to come up with a graded variation of RuDSI.

As a result of clustering, sense clusters were obtained, which contain examples for each target word, labeled with sense number and connected by edges (the edge weight depends on the number and values of annotators’ judgements).

After the first round of annotation, we analyzed the number of *uncompared clusters* — those clusters whose sentences have never been compared in the process of annotation. The existence of uncompared clusters indicates that the graph is not connected enough. We decided that for the five words with the number of uncompared clusters exceeding the average (2.75) additional annotation is required. After the second annotation round (60 extra pairs of sentences for each of five words) there were still four words left for which the number of uncompared clusters has remained almost the same and still exceeded the average number. For these words, sentences from the corresponding clusters were manually selected, organized into pairs and annotated following our regular workflow. After all the annotation rounds, the number of uncompared

clusters is not more than two for any target word, and the average number of annotated sentence pairs per word is 215.

Initially, we got a large number of singleton clusters (1.13 on average across words). These are clusters containing only one node (context, usage example). They may appear when the target word is used in a specific context, for example, in an idiomatic expression. Singleton clusters are problematic, since in these cases it is difficult to tell legitimate exotic senses from clustering errors. We planned to filter them out in one of the following ways: not to consider examples from singleton clusters or to attach singleton examples to the largest cluster of a particular word, reducing the total number of senses. However, after reviewing the clusters manually, we noticed that in some lexemes singleton clusters can be aggregated with a larger one, but not with the largest one, and in other lexemes singleton clusters, on the contrary, express a very specific idiomatic expression that can neither be attached to another cluster nor removed from the sample without loss of representative power. So, we decided to leave the singleton clusters untouched and did not filter them out.

Figure 1 shows the distribution of the number of senses for the target words yielded by the annotation procedure (per-word numbers can be found in the Appendix). As can be seen, most words tend to end up having 3-5 senses.

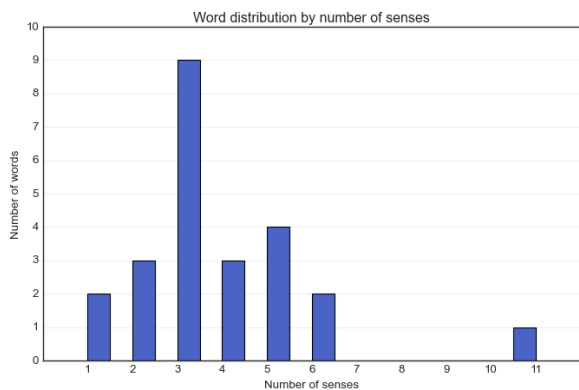


Figure 1: Word distribution by the number of senses obtained in RuDSI.

3.4 Important statistics

Based on the results of clustering, we computed some statistics presented in this subsection. In particular, the ratios of words by the number of senses was calculated. As it turned out, RuDSI contains 8.3% of monosemous words, 62.5% of words with

2-4 senses (moderately polysemous), and 29.2% of words with five or more senses (highly polysemous). Note that these values are different from the original percentages obtained from our linguistics sources. This is expected, since our senses are fully data-driven.

It was also interesting to consider the correlation of these ‘data-driven’ sense numbers and the degree of lexical polysemy yielded by the RNC, Wiktionary and RuWordNet, on which we relied during the selection of the target words. The Table 1 shows Spearman correlation between the number of clusters in the RuDSI word usage graphs and the number of senses in the sources mentioned above. ‘Mean number of senses’ is the average between the RNC, Wiktionary and RuWordNet. All the correlations are strong and significant at $p = 0.05$: that is, the resulting clusters based on data-driven sense induction roughly correspond to sense numbers from external linguistic sources.

Source	Spearman ρ	p value
RNC	0.84	0.000
Wiktionary	0.43	0.034
RuWordNet	0.73	0.000
Mean number of senses	0.90	0.000

Table 1: Correlation of the word sense numbers between RuDSI and other resources.

In addition, we calculated the Spearman correlation between the number of senses in RuDSI and the target word frequencies from the [Lyashevskaya and Sharov \(2009\)](#) dictionary (based on the RNC). Its value is 0.53 ($p = 0.007$). Therefore, the number of word senses in RuDSI is significantly correlated with word frequencies in the RNC. This is expected, since it is known that frequent words tend to be more polysemous ([Zipf, 1945](#); [Hernández-Fernández et al., 2016](#)). It also means that in many cases it is possible to predict the number of RuDSI senses for a word by looking at its RNC frequency.

3.5 Format and technical details

As a result of the steps described above, each example sentence (usage) for each target word was assigned an index of the cluster to which it belongs. We aggregated this data in order to compile a dataset in a format similar to RUSSE-18 ([Panchenko et al., 2018](#)). The structure of the RuDSI dataset is presented in the Table 2: word, context (sentence), positions of the word in the context and the gold identifier of the cluster (sense).

word	context	positions	sense_id
'тысяча'	'...тысяча пятьсот запорожцев...'	76-82	0
'тысяча'	'...вмещает 12 тысяч зрителей.'	34-39	0
'тысяча'	'...около 5 тыс. вагонов...'	49-52	0
'тысяча'	'...ну, сотни тысяч.'	34-39	0
'тысяча'	'...на пытки тысячи ни в чем...'	28-34	1

Table 2: RuDSI dataset sample for the word ‘тысяча’ (*thousand*).

We encourage evaluating state-of-the-art WSI approaches with RuDSI, this is why it was important for the texts in the dataset to not exceed 512 tokens in length. The maximum sequence length is always added to the Transformers architecture models due to the attention layers, which are quadratically scaled with increasing sequence length. 512 tokens is the popular maximum sequence length, which was first specified in BERT. The only sentence in RuDSI (out of 840) which has been longer than this value has been truncated to 512 tokens.

4 Robustness of clustering

In order to verify the stability of clustering algorithm we experimented with changing the default hyperparameters and analyzed the resulting data in comparison with the default sense clusters presented in RuDSI. In the pipeline (Schlechtweg et al., 2021), there were two parameters that could affect the obtained clusters: the threshold used to rescale the annotators’ judgements and the number of clustering iterations. The threshold parameter was previously described in 3.3: it affects the resulting weights on the graph edges. Originally, the threshold was 2.5 causing 1 and 2 scores (‘Unrelated’ and ‘Distantly Related’) to transform to negative values, and 3 and 4 scores (‘Closely Related’ and ‘Identical’) to remain positive to represent the contrast between different senses and the same sense of the word. We reviewed two other options: the threshold equaled to 1.5 (so that a score of 1 became negative (-0.5) and contrasted with 2, 3 and 4 scores that were matched to 0.5, 1.5 and 2.5 respectively) and equaled to 3.5 (1, 2 and 3 scores were opposed to a score of 4; only the sentences marked us ‘Identical’ were considered as containing the same sense of the word).

The number of clustering iterations (‘iters’ parameter) stands for the number of passes through the same graph given that the input graph is the result of the previous iteration. Each pass performs the clustering algorithm and minimizes the loss of

the obtained clusters.

In Table 3, are presented the mean and standard deviation of ARI score among words between the default clustering and clusterings with modified hyperparameters. We can conclude that the number of iterations does not greatly affect the resulting clusters, even as a result of a single iteration (‘iters’ = 1) approximately the same clustering is obtained. However the threshold parameter strongly influences the obtained clusters as it reforms the original idea of similarity of different judgements during the annotation.

iters	threshold	Mean ARI	SD ARI
5	2.5	–	–
5	1.5	0.12	0.29
5	3.5	0.27	0.26
1	2.5	0.95	0.13
3	2.5	0.95	0.10
4	2.5	0.95	0.11
6	2.5	0.94	0.13

Table 3: Similarity (by ARI) of the default RuDSI clustering and clusterings obtained by changing hyperparameters. ‘SD’ stands for standard deviation.

We also examined the change in the number of singleton clusters depending on clustering hyperparameters. Similarly, the threshold parameter has a much stronger effect than the number of iterations. The threshold of 1.5 causes merging of most senses into one cluster (sense) and separation of the minimal number of singleton clusters (0.13 on average). In turn, the threshold of 3.5 generates division into a larger number of clusters most of which are singleton clusters (6.33 on average). Notably, the iterations parameter is inversely proportional to the number of singleton clusters: the more iterations, the more singleton clusters are attached to larger clusters (the more senses are considered the same). A summary of singleton clusters analysis can be found in Table 4.

5 Baseline WSI methods performance

In this section, we show how the existing WSI methods perform on RuDSI. We deliberately did not experiment with the state-of-the-art lexical substitution method (Amrami and Goldberg, 2019). The goal is to report the results of the baseline approaches, leaving more advanced methods for future research.

iters	threshold	# Singletons	SD
5	2.5	1.13	0.74
5	1.5	0.13	0.45
5	3.5	6.33	4.43
1	2.5	1.21	0.83
3	2.5	1.13	0.8
4	2.5	1.13	0.95
6	2.5	1.08	0.88

Table 4: Statistics for singleton clusters in the default RuDSI clustering and clusterings obtained by changing hyperparameters. ‘Singletons’ is the average number of singleton clusters among words. ‘SD’ stands for standard deviation.

5.1 Naive baselines

Two naive baselines were implemented for WSI problem solution, namely assignment of the same sense for all target words, and a random choice of two senses for each target word.

5.2 Birch

Next, we applied more advanced embedding-based approaches. One of the methods top-rated in the RUSSE-18 shared task is static embeddings clustering (Panchenko et al., 2018). After testing different clustering algorithms, we settled on Birch (Zhang et al., 1996), which provided the best results. We used the following pipeline: first, we calculated sentence embeddings as an average over word embeddings for each context, second, all embeddings within each target context were divided into two clusters. For word embedding extraction we used *ruwikiruscorpora_func_upos_skipgram_300_5_2019* Word2Vec model trained on RNC and Wikipedia from the RusVectores web service (Kutuzov and Kuzmenko, 2017).

5.3 Jamsic

Jamsic method was also included in the list of the best systems in the RUSSE-18 shared task description paper (Panchenko et al., 2018). Using the Word2Vec model specified above, the nearest neighbor for each target word is extracted. The embedding of this word represents the first sense of the target word. Then this embedding is subtracted from the embedding of the target word and the embedding of the second sense is obtained. Finally, we get an average embedding for each sentence

and determine to which sense it is closer by cosine similarity. This method works with one word sense and its nearest one, so it always distributes contexts into only two senses.

5.4 Egvi

This is a relatively new approach that has successfully proved itself in solving the WSI problem for different languages. For this method, we used Russian sense inventories pre-generated by processing ego graphs (Logacheva et al., 2020), and for each target word we received an average embedding of each sense from sense inventories. For word embeddings, we used the same *ruwikiruscorpora_func_upos_skipgram_300_5_2019* model. Then we removed the target word from RuDSI contexts, received average word embeddings and clustered them with the KMeans algorithm, passing embeddings of values from sense inventories as cluster centers. The parameter of number of clusters for each target word was equal to number of senses in sense inventories for this word.

5.5 BERT KMeans

BERT-based embeddings proved to be efficient in solving RUSSE-18 too (Slapoguzov et al., 2021). We took the *sbert_large_nlu_ru* model⁹ as a feature extractor and used token embeddings from its last layer. For calculating the representation of words split during tokenization, mean pooling was used. Word vectors were clustered by the KMeans algorithm into two senses.

We also tried to do KMeans clustering of BERT embeddings by taking the number of clusters from Egvi sense inventories.

5.6 Results

The mean and standard deviation of ARI score among words are presented in Table 5. The ARI metric takes into account randomness when clustering, so the ARI of the Random sense method is 0.0. The approaches that became the best in the RUSSE-18 shared task do not gain values higher than 0.05 on RuDSI. The simplest One sense baseline is better than BERT clustering. Arguably, the low BERT results are caused by the number of clusters parameter of the KMeans algorithm, which was equal to 2, while only two target words (out of 24) actually had two senses. Egvi algorithm proved to be the best. This method was based on the pre-generated

⁹https://huggingface.co/sberbank-ai/sbert_large_nlu_ru

sense inventories, in which the number of senses often was identical to RuDSI, so it worked better than BERT KMeans. The number of Egvi senses improved the quality of clustering of BERT embeddings, but not enough to exceed the native Egvi.

For comparison, the table shows the results of the methods on the RUSSE-18 dataset. Due to a number of limitations described earlier, a wiki-wiki dataset was taken for comparison. It is noticeable that the values of the ARI metric for the basic methods on wiki-wiki are much higher.

We also found no correlation between the density of the word graph and the values of the ARI metric, with the exception of the Jamsic method, for which the correlation results are significant at a significance level of $p = 0.05$.

Method	RuDSI		RUSSE	
	Mean ARI	SD ARI	Mean ARI	SD ARI
One sense	0.08	0.28	0.00	0.00
Random sense	0.00	0.00	0.01	0.00
Birch	0.03	0.14	0.93	0.10
Jamsic	0.04	0.10	0.58	0.47
BERT KMeans	0.03	0.14	0.85	0.06
BERT KMeans + Egvi	0.08	0.31	0.64	0.31
Egvi	0.17	0.22	0.59	0.16

Table 5: Performance of WSI methods on RuDSI and RUSSE. ‘SD’ stands for standard deviation.

6 Intended RuDSI audience

Our vision is that RuDSI might be of use for three different communities.

1. Researchers analyzing NLP systems in terms their WSD and WSI abilities for Russian. It is especially important for evaluating contextualized language models trained on large-scale corpora using deep neural architectures, from RNNs to Transformers and beyond. RussianSuperGLUE benchmark (Shavrina et al., 2020) already includes the RUSSE dataset (cast as a binary classification task). We believe RuDSI might be a useful addition, representing a more difficult task related to lexical senses. As was shown in 5, it cannot be solved with trivial baselines (Iazykova et al., 2021), which makes it an interesting NLP challenge.
2. Graph theory researchers and all those interested in applications of graphs to real world tasks. Word usage graphs we are providing

are representative of contextual semantic similarity judgments by humans. These graphs can be processed and clustered in different ways, yielding different ‘views’ of Russian word sense inventories. In addition, the properties of word usage graphs themselves can bring new insights for both graph theory and Russian linguistics.

3. Finally, our work on RuDSI is a part of a larger project of implementing WSI features into the RNC web interface. RuDSI is based on RNC data, so it will be used to evaluate various WSI solutions and choose the best one. Thus, it is going to be directly or indirectly used by the large RNC audience, consisting of both linguists and general population.

7 Conclusion

We have presented RuDSI, a novel graph-based word sense induction dataset for Russian, obtained by clustering word usage graphs produced by human annotation. It includes words with different degrees of polysemy (monosemous, moderately polysemous and highly polysemous words). The sense inventories are generated in a completely data-driven way as well. Importantly, depending on what graph processing workflow is used, slightly different datasets can be produced from the same raw RuDSI human judgments.

We report the RuDSI performance for only the simplest and most basic approaches to WSI, so a possible future work would be to apply some more advanced methods to it. Also we have considered only nouns, so it would be interesting to experiment with other parts of speech as well (this will require a new round of annotation). Since most of the target words in RuDSI have 3-5 senses, the addition of highly polysemous words may become another future improvement. In addition, it would be beneficial to extend the list of contexts for each word, however extra annotation would be required.

References

- Eneko Agirre and Aitor Soroa. 2007. *SemEval-2007 task 02: Evaluating word sense induction and discrimination systems*. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 7–12, Prague, Czech Republic. Association for Computational Linguistics.

- Asaf Amrami and Yoav Goldberg. 2019. Towards better substitution-based word sense induction. *arXiv preprint arXiv:1905.12598*.
- Juri Apresjan. 2014. Active dictionary of russian. *Meaning-Text Theory*, 2014.
- Nikhil Bansal, Avrim Blum, and Shuchi Chawla. 2004. Correlation clustering. *Machine learning*, 56(1):89–113.
- Andreas Blank. 1997. *Prinzipien des lexikalischen Bedeutungswandels am Beispiel der romanischen Sprachen*. Niemeyer, Tübingen.
- Antoni Hernández-Fernández, Bernardino Casas, Ramon Ferrer-i Cancho, and Jaume Baixeries. 2016. Testing the robustness of laws of polysemy and brevity versus frequency. In *International Conference on Statistical Language and Speech Processing*, pages 19–29. Springer.
- David Hope and Bill Keller. 2013. Maxmax: a graph-based soft clustering algorithm applied to word sense induction. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 368–381. Springer.
- Lawrence Hubert and Phipps Arabie. 1985. Comparing partitions. *Journal of classification*, 2(1):193–218.
- Tatyana Iazykova, Denis Kapelyushnik, Olga Bystrova, and Andrey Kutuzov. 2021. **Unreasonable effectiveness of rule-based heuristics in solving Russian SuperGLUE tasks**. *Komp'yuternaya Lingvistika i Intellektual'nye Tekhnologii: Dialog conference*.
- David Jurgens and Ioannis Klapaftis. 2013. **SemEval-2013 task 13: Word sense induction for graded and non-graded senses**. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 290–299, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Adam Kilgarriff. 1997. I don't believe in word senses. *Computers and the Humanities*, 31(2):91–113.
- P. S. Kostenetskiy, R. A. Chulkevich, and V. I. Kozyrev. 2021. **HPC resources of the higher school of economics**. *Journal of Physics: Conference Series*, 1740(1):012050.
- Andrey Kutuzov and Elizaveta Kuzmenko. 2017. **Webvectors: A toolkit for building web interfaces for vector semantic models**. In *Analysis of Images, Social Networks and Texts: 5th International Conference, AIST 2016, Yekaterinburg, Russia, April 7-9, 2016, Revised Selected Papers*, pages 155–161, Cham. Springer International Publishing.
- Andrey Kutuzov and Lidia Pivovarova. 2021. **Three-part diachronic semantic change dataset for Russian**. In *Proceedings of the 2nd International Workshop on Computational Approaches to Historical Language Change 2021*, pages 7–13, Online. Association for Computational Linguistics.
- Varvara Logacheva, Denis Teslenko, Artem Shelmanov, Steffen Remus, Dmitry Ustalov, Andrey Kutuzov, Ekaterina Artemova, Chris Biemann, Simone Paolo Ponzetto, and Alexander Panchenko. 2020. **Word sense disambiguation for 158 languages using word embeddings only**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5943–5952, Marseille, France. European Language Resources Association.
- Natalia V Loukachevitch, German Lashevich, Anastasia A Gerasimova, Vladimir V Ivanov, and Boris V Dobrov. 2016. Creating russian wordnet by conversion. In *Computational Linguistics and Intellectual Technologies: papers from the Annual conference "Dialogue"*, pages 405–415.
- Ol'ga Nikolaevna Lyashevskaya and Sergej Aleksandrovich Sharov. 2009. *Chastotny'j slovar' sovremennogo russkogo yazy'ka: na materialax Nacional'nogo korpusa russkogo yazy'ka*. Azbukovnik.
- Suresh Manandhar, Ioannis Klapaftis, Dmitriy Dligach, and Sameer Pradhan. 2010. **SemEval-2010 task 14: Word sense induction & disambiguation**. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 63–68, Uppsala, Sweden. Association for Computational Linguistics.
- Diana McCarthy, Marianna Apidianaki, and Katrin Erk. 2016. Word sense clustering and clusterability. *Computational Linguistics*, 42(2):245–275.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- George A Miller, Claudia Leacock, Randee Teng, and Ross T Bunker. 1993. A semantic concordance. In *Human Language Technology: Proceedings of a Workshop Held at Plainsboro, New Jersey, March 21-24, 1993*.
- Alexander Panchenko, Anastasia Lopukhina, Dmitry Ustalov, Konstantin Lopukhin, Nikolay Arefyev, Alexey Leontyev, and Natalia Loukachevitch. 2018. **RUSSE'2018: A Shared Task on Word Sense Induction for the Russian Language**. In *Computational Linguistics and Intellectual Technologies: Papers from the Annual International Conference "Dialogue"*, pages 547–564, Moscow, Russia. RSUH.
- Yong Peng, Xin Zhu, Feiping Nie, Wanzeng Kong, and Yuan Ge. 2021. **Fuzzy graph clustering**. *Information Sciences*, 571:38–49.
- Ekaterina Rahilina, Galina Kustova, Ol'ga Lyashevskaya, Tatyana Reznikova, and Ol'ga Shemanayeva. 2009. Zadachi i principy semanticheskogo razmetki leksiki v nkrya. In *Russian National Corpus: 2006-2008. New results and perspectives*, pages 215–240. Nestor-Istorija.

- Julia Rodina and Andrey Kutuzov. 2020. [RuSemShift: a dataset of historical lexical semantic change in Russian](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1037–1047, Barcelona, Spain (Online). International Committee on 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.
- Tatiana Shavrina, Alena Fenogenova, Emelyanov Anton, Denis Shevelev, Ekaterina Artemova, Valentin Malykh, Vladislav Mikhailov, Maria Tikhonova, Andrey Chertok, and Andrey Evlampiev. 2020. [RussianSuperGLUE: A Russian language understanding evaluation benchmark](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4717–4726, Online. Association for Computational Linguistics.
- Leonid Sherstuk. 2020. Context-based text-graph embeddings in word-sense induction tasks. Master’s thesis, HSE.
- Aleksandr Slapoguzov, Konstantin Malyuga, and Evgenij Tsopa. 2021. Word sense induction for Russian texts using BERT. In *Proceedings of the 28th Conference of Fruct Association, Moscow, Russia*, pages 25–29.
- Tian Zhang, Raghu Ramakrishnan, and Miron Livny. 1996. [Birch: An efficient data clustering method for very large databases](#). In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data, SIGMOD ’96*, page 103–114, New York, NY, USA. Association for Computing Machinery.
- George Kingsley Zipf. 1945. The meaning-frequency relationship of words. *The Journal of general psychology*, 33(2):251–256.

A Annotation interface

B Word usage graph

C Detailed performance

Test - Please indicate the semantic relatedness of the two uses of the marked words in the sentences above.

4
Identical

3
Closely Related

2
Distantly Related

1
Unrelated

-
Cannot decide

NEXT

Annotation

Sentence 1

При этом важны не только масштабы производства, но и его качественная **сторона**, то есть эффективное управление активами компании, подготовка корпоративного менеджмента, динамика формирования акционерной прибыли, уровень капитализации.

Sentence 2

Так, донеся государю императору Александру о занятии Реймса, полки разошлись на пространстве от города верст до тридцати на квартиры в разные **стороны**.

You can use keyboard shortcuts to perform annotations. Use the numeric keypad (keys 1-4, and 0) to select the desired vote, then press Enter to submit it. You can pause the annotation process anytime by pressing the button 'Pause'. We have already saved your annotations, feel free to continue anytime.

Figure 2: Example of the word ‘сторона’ (*side, direction*) annotation in the DUREl interface.

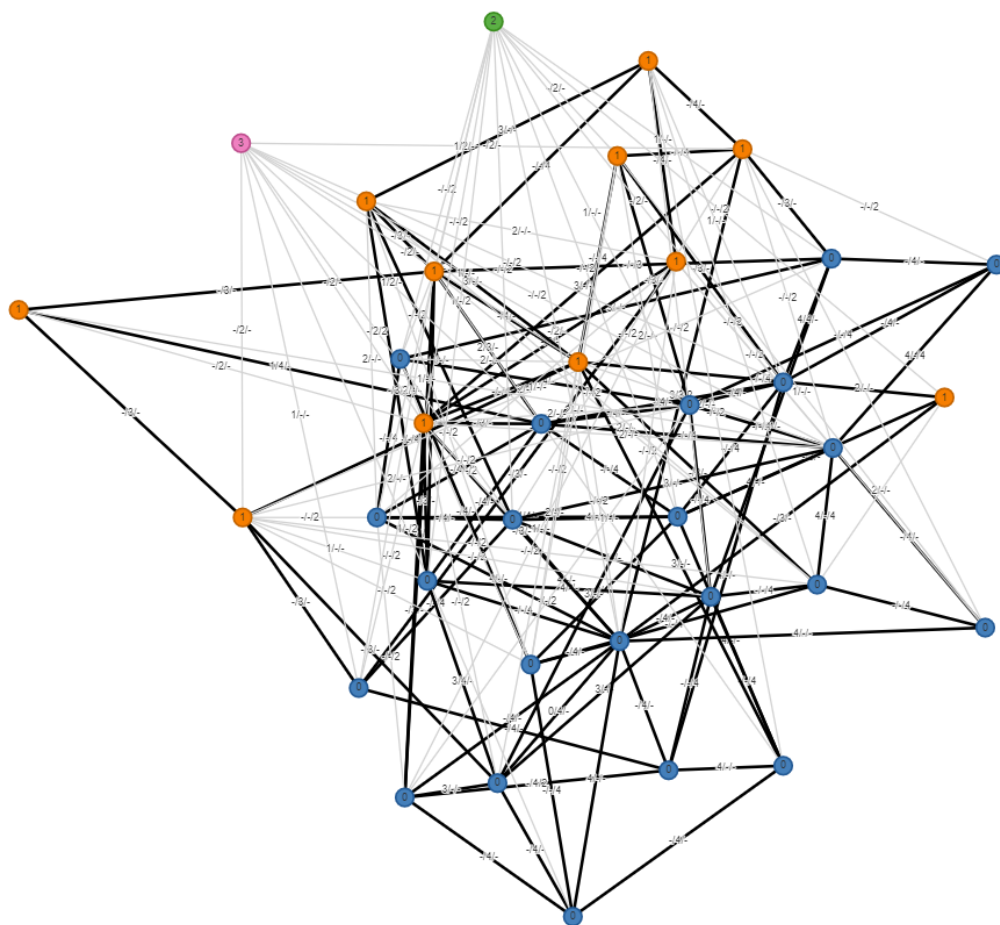


Figure 3: Word usage graph for the word ‘голова’ (*head*) as a result of clustering (four clusters, marked with node color).

Word	One sense	Random sense	Birch	Jamsic	BERT KMeans	Egvi	BERT KMeans + Egvi	№ of clusters
'Бог' (God)	0.00	-0.02	-0.04	0.13	-0.04	-0.04	0.01	3
'Время' (time)	0.00	-0.03	0.02	0.00	-0.07	-0.01	-0.04	6
'Год' (year)	0.00	0.06	-0.03	0.01	-0.02	-0.04	0.13	3
'Голова' (head)	0.00	-0.03	0.51	-0.02	0.20	0.61	-0.02	4
'Город' (city)	0.00	-0.01	-0.03	0.01	-0.01	1.00	0.00	2
'Государство' (state)	-0.06	0.02	-0.04	-0.06	-0.05	-0.04	-0.05	3
'Дело' (business)	0.00	0.03	-0.01	0.00	0.01	0.02	0.08	11
'День' (day)	0.00	0.11	-0.05	0.08	-0.02	-0.05	0.00	5
'Друг' (friend)	0.00	-0.02	0.25	-0.04	-0.09	0.12	-0.01	3
'Жена' (wife)	0.00	-0.03	0.00	-0.03	-0.04	0.00	0.0	2
'Женщина' (woman)	1.00	0.00	0.00	0.00	0.00	1.00	1.0	1
'Жизнь' (life)	0.00	0.00	0.33	-0.01	-0.13	-0.04	0.04	4
'Лицо' (face)	0.00	-0.02	0.05	0.50	0.39	0.00	0.61	3
'Место' (place)	0.00	0.05	-0.10	0.09	0.01	-0.04	-0.02	4
'Мир' (world)	0.00	-0.01	0.13	0.14	0.04	0.20	0.00	5
'Ночь' (night)	1.00	0.00	0.00	0.00	0.00	0.00	0.00	1
'Работа' (work)	0.00	-0.01	-0.04	-0.04	-0.04	0.15	0.05	5
'Результат' (result)	0.00	0.01	-0.07	0.04	0.01	0.63	-0.03	2
'Рука' (hand)	0.00	0.00	-0.08	0.06	0.38	0.16	0.05	3
'Сила' (power)	0.00	-0.01	-0.02	-0.02	0.22	-0.02	0.04	6
'Слово' (word)	0.00	-0.03	0.01	0.01	0.19	0.00	0.00	3
'Сторона' (side)	0.00	-0.01	-0.04	0.21	0.30	0.41	0.20	5
'Тысяча' (thousand)	0.00	0.00	-0.08	-0.01	-0.05	-0.04	-0.01	3
'Человек' (human)	0.00	0.00	-0.04	-0.06	-0.00	0.00	0.00	3

Table 6: Detailed performance of WSI methods.