

# Multilingual Substitution-based Word Sense Induction

Denis Kokosinski<sup>1,2</sup>, Nikolay Arefyev<sup>3</sup>

<sup>1</sup>Lomonosov Moscow State University, Russia

<sup>2</sup>SaluteDevices, Russia

<sup>3</sup>University of Oslo, Norway

dvl.kokosinskiy@sberbank.ru, nikolare@uio.no

## Abstract

Word Sense Induction (WSI) is the task of discovering senses of an ambiguous word by grouping usages of this word into clusters corresponding to these senses. Many approaches were proposed to solve WSI in English and a few other languages, but these approaches are not easily adaptable to new languages. We present multilingual substitution-based WSI methods that support any of 100 languages covered by the underlying multilingual language model with minimal to no adaptation required. Despite the multilingual capabilities, our methods perform on par with the existing monolingual approaches on popular English WSI datasets. At the same time, they will be most useful for lower-resourced languages which miss lexical resources available for English, thus, have higher demand for unsupervised methods like WSI.

**Keywords:** Word Sense Induction, Multilinguality, Lexical Substitution

## 1. Introduction

The task of word sense induction (WSI) is to cluster occurrences (*instances*) of an ambiguous word into clusters corresponding to the senses of this word. E.g., for (1) *He sat on the **bank** of the river*, (2) *He cashed a check at the **bank***, and (3) *That **bank** holds the mortgage on my home*, the desired clusters are {1} and {2, 3}.

The core assumption in WSI is the lack of a pre-defined sense inventory. If a word sense inventory is available, the clustering task can be reformulated as a classification task over senses in this inventory. It is usually referred to as Word Sense Disambiguation (WSD) in this case. It is worth noticing that WSI is more relevant than WSD for lower-resourced languages where linguistic resources are scarce.

Although many approaches exist for WSI, the SOTA or near-SOTA ones are language-specific and are not easily adaptable to new languages. Our main goal is to create a WSI system that works seamlessly in many languages. The contributions of this paper are the following.

1. Multilingual WSI methods are proposed that generalize to any language supported by the underlying multilingual language model with minimal<sup>1</sup> to no adaptation while performing on par with the state-of-the-art WSI approaches for English<sup>2</sup>.

2. We perform rigorous evaluation and analysis of various configurations of our system on datasets

in 11 languages.

3. We discover remarkable abilities of a multilingual masked language model after monolingual unsupervised finetuning to generate cross-lingual lexical substitutes performing on par with the monolingual ones for WSI.

## 2. Related Work

WSI can be formulated as either a hard clustering problem, meaning each instance must be put to a single cluster, or a soft clustering problem where each instance must receive some weights or probabilities of belonging to each cluster. SemEval 2010 task 14 (Manandhar et al., 2010) and SemEval 2013 task 13 (Jurgens and Klapaftis, 2013) are the standard WSI benchmarks in English requiring hard and soft clustering, respectively. RUSSE'2018 (Panchenko et al., 2018), and RuDSI (Aksenova et al., 2022) are publicly available WSI benchmarks in Russian requiring hard clustering.

To the best of our knowledge, no prior work has attempted to design a multilingual WSI model. However, recently in (Pasini et al., 2021) a benchmark and several methods for multilingual WSD have been introduced. Their multilingual methods rely on XLM-R (Conneau et al., 2020), a Transformer-based masked language model trained on a corpus of 100 languages. For WSD methods labeled training data for each word sense and/or sense definitions are required. This limits the applicability of the proposed methods for lower-resourced languages.

As for monolingual WSI methods, one popular approach implemented in (Lau et al., 2013; Wang et al., 2015; Komninos and Manandhar, 2016; Amplayo et al., 2019) relies on the Latent Dirichlet

<sup>1</sup>For 78 languages supported by both the Stanza (Qi et al., 2020) lemmatizer and the XLM-R multilingual LM (Conneau et al., 2020) only one Hearst-like pattern should be translated.

<sup>2</sup>Code is available at <https://github.com/deniskokoski/mwsi.git>.

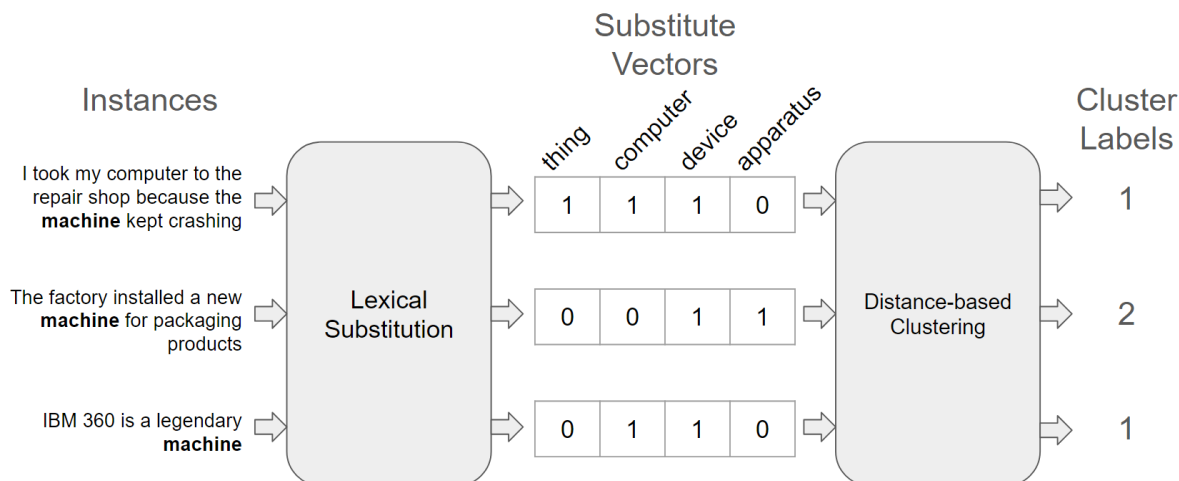


Figure 1: Substitution-based approach to Word Sense Induction.

Allocation method of topic modeling. AutoSense (Amplayo et al., 2019), the latest and the best-performing method in this group, incorporates topic and sense latent variables, as well as a switch variable to determine whether to use the local context of the target word or the whole instance.

The performance of WSI methods improved significantly with the adoption of neural masked language models. Recently, PolyLM (Ansell et al., 2021) introduced a transformer-based architecture for learning sense embeddings from an unlabeled corpus. PolyLM allocates eight embeddings for each word to represent up to eight different senses per word. The model is trained to make these embeddings distinct and helpful for the standard masked language modeling task. An important detail is that PolyLM relies on a word-level vocabulary. However, a multilingual word-level vocabulary will be prohibitively large and it is not clear how this model can be adapted to work with a subword-level vocabulary.

Another class of methods for WSI, including AIKU (Başkaya et al., 2013), LSDP ELMo (Amrami and Goldberg, 2018), LSDP BERT (Amrami and Goldberg, 2019), BERT NoDP (Amrami and Goldberg, 2019; Eyal et al., 2022) and +embs (Arefyev et al., 2020), employ lexical substitution. A *lexical substitute* is a word that can replace some target word in a given sentence without making the sentence ungrammatical, and also have the same or similar/related meaning. These methods use a two-step process (see Figure 1): first, generate a vector of lexical substitutes for each instance using a substitution model and then cluster the resulting substitute vectors. An important advantage of lexical substitution methods for WSI is their interpretability. Each cluster can be labeled with a few substitutes that describe this particular sense. It is worth noticing that lexical substitution

is a long-standing NLP task by itself. There are several human-annotated benchmarks for lexical substitution: COInCo (Kremer et al., 2014) and SemEval2007 task 10 (McCarthy and Navigli, 2007) in English, GermEval 2015 (Miller et al., 2015) in German, and SemDisFr 2014 (Fabre et al., 2014) in French. We employ these datasets for intrinsic comparison of our lexical substitution methods and studying the relations between target words and lexical substitutes generated by these methods. Next we describe the process of generating substitutes in two of the substitution-based WSI methods, BERT LSDP (Amrami and Goldberg, 2019) and +embs (Arefyev et al., 2020). These methods show SOTA results on English WSI benchmarks.

BERT LSDP (Language model Substitutions with Dynamic Patterns) follows the substitution-based approach to WSI and uses BERT to produce lexical substitutes. In theory, BERT can estimate probabilities of lexical substitutes by simply replacing the target word with the <mask> token. However, Amrami and Goldberg (2019) show that such substitutes are too general and often unrelated to the target word. BERT LSDP overcomes this issue and "injects" the information about the target word into substitutes. Specifically, it averages two vectors of unnormalized probabilities (logits) of potential substitutes. The first vector of logits is taken from the BERT MLM head directly at the first token of the target word in a given instance without masking it. The second vector of logits is taken from the <mask> token after replacing the target word with the *dynamic pattern* "<T> (or even <mask>)". For example, for the sentence "This cat is cute" the input to BERT after applying the dynamic pattern is "This cat (or even <mask>) is cute." Arefyev and Bykov, 2021 have shown that combining substitutes from two *symmetric dynamic patterns*, e.g. "<T> (or even <mask>)" and "<mask> (or even <T>)" consistently

improves WSI performance compared to using one of these patterns alone. This was demonstrated for several patterns on a WSI dataset in Russian.

The substitution-based method "+embs" (Arefyev et al., 2020) relies on non-contextualized embeddings to inject the target word. The "general" substitutes are first produced by a language model with the target word masked. These substitutes are then reranked based on the a priori similarity to the target word. More specifically, the probabilities of substitutes to appear in the given context as estimated by a masked language model are multiplied by the exponentiated and re-normalized cosine similarities between their non-contextualized embeddings and the non-contextualized embedding of the target word. Several underlying language models were tested with +embs. XLNET (Yang et al., 2019) has shown superior performance, while the results with BERT are comparable to BERT LSDP.

### 3. Substitution-based Methods for Multilingual WSI

#### 3.1. From BERT to XLM-R

Incorporating a multilingual language model into BERT-based methods for WSI seems straightforward. XLM-R (Conneau et al., 2020) is a publicly available Masked Language Model (MLM) trained on a corpus of 100 languages. Theoretically, it can replace BERT (Devlin et al., 2018) in the monolingual methods and make them multilingual. However, we found the results after this replacement unsatisfactorily low and had to analyze the underlying reasons to improve the quality.

The main reason is the difference in tokenization between BERT and XLM-R. Despite having a significantly larger token vocabulary (250k in XLM-R versus 30k in English BERT), the average number of tokens per word is significantly higher on average for XLM-R due to its multilinguality. Figure 2 shows the distribution of the number of XLM-R tokens per word for human-written lexical substitutes in English, French, and German substitution datasets. While only approximately 15% of English substitutes consist of multiple tokens in BERT, for XLM-R this number increases to 20% and reaches 40% for French substitutes and 55% for German substitutes. The original BERT LSDP method only produces single-subword substitutes, thus limiting the diversity of substitutes yielded by XLM-R. Arefyev and Bykov (2021) have shown that two-subword XLM-R substitutes achieve significantly better WSI results for Russian. In general we want all reasonable substitutes to be generated independently of their length. Thus, we introduce methods that produce substitutes of variable length.

Employing the multilingual and cross-lingual ca-




Figure 2: Cumulative distribution of the number of XLM-R and BERT tokens per substitute in human-annotated lexical substitution datasets. English substitutes are from ColnCo (Kremer et al., 2014), German substitutes are from GermEval 2015 (Miller et al., 2015), and French substitutes are from SemDis 2014 (Fabre et al., 2014). Multi-word substitutes are not taken into account.

pabilities of XLM-R, we experiment with generation of substitutes in the language of the input word occurrence, and also cross-lingual substitutes in English or Russian independently of the input language.

#### 3.1.1. Concat

In the *Concat* substitute generator substitutes are generated by a masked language model that receives an instance with a target word replaced by 1, 2, or 3 consecutive <mask> tokens.<sup>3</sup> For example, an instance "cats are cute" turns into three input sequences:

"<mask> are cute",  
 "<mask><mask> are cute",  
 "<mask><mask><mask> are cute".

For each input we select  $k$  most probable tokens for the first mask and decode tokens for the following masks using greedy decoding<sup>4</sup>, resulting in  $k$  generated sequences of subwords. We then post-process these sequences by merging subwords into words and keeping only the first word from each sequence<sup>5</sup>. The resulting probability of a sub-

<sup>3</sup>We do not generate substitutes consisting of more than 3 subwords because such substitutes are very rarely suggested by humans as figure 2 shows. Experimental results on the importance of the used number of masks for WSI are presented in Appendix C.

<sup>4</sup>We tried using beam search instead, but did not observe significant improvements.

<sup>5</sup>We do not use the whole multi-word expressions generated because they increase the sparsity of the substitute vectors without providing intuitive benefits.

stitute is the product of conditional probabilities of its subwords. In total,  $3k$  substitutes are generated and  $k$  most probable are selected among them.

### 3.1.2. WCM

The Word Continuation Masking (WCM) approach requires self-supervised finetuning of the original masked language model, unlike concat. The idea of WCM is to train the masked language model to produce tokens of a lexical substitute one at a time, similarly to autoregressive language models, while leveraging both left and right contexts during generation. This kind of finetuning can be done on an unlabeled corpus of texts.

The autoregressive generation of substitutes cannot be adequately implemented for masked language models out of the box. For instance, in case when only one mask is present (e.g., "cats and <mask> are cute") only single-token substitutes will likely be generated by an MLM. The first tokens of multi-token words will have low probabilities because the model "sees" only one mask and "understands" that there is no space for a multi-token word.

WCM is similar to the standard masked language modeling objective with an important difference. Similarly to BERT 15% of tokens in an input text are masked and the objective is to restore the masked tokens. However, in WCM the boundaries of the words fully or partially replaced by a <mask> are then identified. All tokens of such words following the <mask> token are removed. The tokens that precede the <mask> token and the tokens of other words are left intact. Thus, the model is trained to predict the first token of a masked word or to continue a word without knowing the exact length of this word. During inference, masks get inserted into the context one by one to generate substitutes of arbitrary length. For example, the following decoding path is possible: "cats and <mask> are cute", "cats and copy<mask> are cute", "cats and copybara<mask> are cute", resulting in a three-token substitute "copybaras".

We trained **WCM**, **WCM-en**, and **WCM-ru**, three models based on XLM-R and finetuned with the WCM objective. WCM is trained on texts in all 100 languages supported by XLM-R, preserving the original proportions of languages, while WCM-en is trained only on texts in English and WCM-ru only on texts in Russian<sup>6</sup>.

<sup>6</sup>We trained each model for 2 epochs on  $\approx 6$ GB of texts randomly sampled from CommonCrawl100 (Wenzek et al., 2020).

### 3.1.3. Monolingual fine-tuning results in cross-lingual substitutes

We empirically observe a remarkable property of monolingual WCM finetuning. WCM-en and WCM-ru always yield substitutes in the finetuning language, not in the language of the input text. Surprisingly, these substitutes often constitute a reasonable description of the meaning of the target word. For example, for a sentence *à la <mask> comme à la guerre!*<sup>7</sup> the top-5 substitutes by WCM-en are words in English that are semantically similar to the masked target word in French: *war, peace, liberty, faith, freedom*. The same behavior holds for more complex examples, such as the ones found in WSI datasets.

The possibility of generating sensible English substitutes for a word occurrence in any language allows using English lemmatizer instead of searching for a lemmatizer for each input language. This significantly simplifies adaptation of our WSI methods to new languages. Cross-lingual substitutes in English may also enable direct utilization of rich English knowledge resources for other languages, but we leave this for future research. It also provides interpretable sense labels in English for an ambiguous word in any language. We provide examples of such labels in section 5.

## 3.2. Target Injection

A straightforward approach to generate substitutes with a masked language model is to replace the target word with the <mask> token. This way the language model produces substitutes that fit a given context. However, these substitutes may not be semantically related to the original target word. Target injection helps to condition the language model substitutes on the target word and get substitutes better describing its meaning. We adapt two target injection methods described in Section 2 to multilingual applications.

**SDP (Symmetric Dynamic Patterns)**. Due to the intuitive nature of dynamic patterns for target injection, they are straightforward to adapt to multilingual applications. We select the "<T> (or even <mask>)" pattern with its symmetric counterpart "<mask> (or even <T>)" for our experiments following (Amrami and Goldberg, 2019). We have experimented with several other patterns and observed no significant differences in WSI performance (refer to Appendix A for details). We use Google Translate to translate the patterns to the target language. It is the only action required to adapt our methods to any of the 78 languages supported by both XLM-R and the Stanza lemmatizer employed to lemmatize substitutes.

<sup>7</sup>A French proverb *à la guerre comme à la guerre!*, translates as *at war as at war*.

We also try simplifying the adaptation even further and experiment with applying non-translated versions of the dynamic patterns in English or Russian along with the cross-lingual WCM-en and WCM-ru models, resulting in code-switched inputs to these models. This gives a WSI method that supports all 100 languages of XLM-R without any adaptation.

In our implementation of the target injection with SDP, the probabilities of substitutes for two symmetric patterns are multiplied with smoothing<sup>8</sup>. Thus, the substitutes in the top  $k$  for both dynamic patterns will rise to the top of the ranking.

**+embs.** In the original +embs method, the substitutes can only consist of a single BERT token, which is a problem for multilingual WSI as discussed previously. The original method uses the embedding layer of a language model as the source of non-contextualized embeddings, however, this does not generalize to multi-token substitute generation.

We use FastText (Bojanowski et al., 2016) embeddings instead of relying on the embedding layer of the MLM. Fasttext provides out-of-the-box non-contextualized embeddings for all languages supported by XLM-R. We compute the distances between the FastText embeddings of the target word and each of  $3k$  most probable substitutes yielded by a multi-token substitute generator. As in the original method, all  $3k$  non-contextualized similarities pass through a temperature softmax. It helps align the peaky distribution from the substitute generator with a more flat distribution of FastText similarities. The final ranking of substitutes is the product of the non-contextualized (FastText) and the contextualized (MLM) distributions.

### 3.3. Multilingual WSI Methods

Our methods for multilingual WSI follow the general substitution-based approach to WSI. We generate  $k$  substitutes of 1 to 3 XLM-R tokens with one of the multi-token substitute generators described in Section 3.1. We use one of the methods from Section 3.2 for target injection. We see no theoretical preference in any of the given options and rely on experiments to determine the best configuration.

Intuitively, a reasonable number of lexical substitutes  $k$  for a word in a sentence should not exceed 5 or 10. It is often difficult to find many synonyms for a target word in a given context. However, (Amrami and Goldberg, 2019) have experimentally found the optimal  $k = 200$  for BERT LSDP on WSI benchmarks in English. We confirm their findings on a WSI dataset in Russian for our SDP method. The WSI performance for SDP reaches a plateau only

<sup>8</sup>If a substitute is not present among top  $k$  substitutes of one of the patterns, it is assigned the probability of  $1e-5$ .

at  $k \approx 100$  (see Appendix B). In the case of +embs, however, we found that  $k = 20$  is optimal with performance significantly declining for larger values of  $k$ . Thus, we have chosen  $k = 150$  for SDP and  $k = 20$  for +embs. We further discuss possible reasons for this distinction in section 5.

The next step in the substitution-based methods is to build a numeric vector for each instance based on the generated substitutes. We lemmatize the substitutes with the Stanza lemmatizer<sup>9</sup>. It supports 78 of 100 XLM-R languages out of the box. We experimentally compared Stanza to other multilingual and monolingual lemmatizers and observed no significant impact on WSI results<sup>10</sup>. We then build a TF-IDF vector from the lemmatized substitutes for each instance.

Finally, TF-IDF vectors are clustered. This step is language-independent. While plenty of clustering methods exist, we choose simple agglomerative clustering with cosine distance and average linkage, following previous work on WSI (Amrami and Goldberg, 2019; Arefyev et al., 2020). For each target word in the dataset, we perform clustering with the number of clusters ranging from 2 to 9 and select the clustering with the highest Calinski-Harabasz score (Caliński and Harabasz, 1974), an unsupervised score of clustering quality. In the soft clustering benchmarks we use a trivial hard-to-soft clustering conversion, where the hard cluster is assigned a probability of 1. While more complicated clustering approaches may increase the overall WSI performance, our primary focus is on the substitute generation, as we consider it to be the most language-dependent step in the lexical-substitution WSI pipeline.

## 4. WSI Evaluation

### 4.1. Datasets

We use the existing WSI datasets as well as datasets created for other lexical semantic tasks for our multilingual WSI experiments.

The most popular WSI benchmarks for English are **SE10** (SemEval-2010 task 1, Manandhar et al., 2010) and **SE13** (SemEval-2013 task 13, Jurgens and Klapaftis, 2013). They were annotated manually. SE10 provides gold hard clustering of instances, while SE13 provides gold soft clustering.

RUSSE'2018 (Panchenko et al., 2018) is a shared task that presented three manually annotated WSI datasets in Russian: wiki-wiki, bts-rnc

<sup>9</sup><https://stanfordnlp.github.io/stanza/lemma.html>

<sup>10</sup>We experimented with Spacy (<https://spacy.io/api/lemmatizer>), WordNet lemmatizer (nlkt.stem.wordnet) and PyMorphy2 (Korobov, 2015a). The differences in WSI metrics are within 5% and are not consistent between datasets and languages.

<i>Dataset</i>	<i>Instances</i>	<i>words</i>	<i>Instances per word</i>	<i>Senses per word</i>	<i>Avg instance length (words)</i>
SE10	8915	100	89.2	3.9	64.9
SE13	4664	48	97.2	-	30.3
bts-rnc-ru-test-private	4335	34	127.5	3.0	25.0
DeWUG-de-Sense	826	24	34.4	2.8	33.4
XL-WSD-bg	3259	380	8.6	2.4	3.3
XL-WSD-de	117	14	8.4	2.1	28.3
XL-WSD-en	3126	396	7.9	2.8	26.3
XL-WSD-es	213	36	5.9	2.2	31.4
XL-WSD-fr	181	33	5.5	2.2	34.9
XL-WSD-gl	972	184	5.3	2.8	4.9
XL-WSD-it	319	54	5.9	2.3	31.7
XL-WSD-ko	865	164	5.3	2.5	6.0*
XL-WSD-sl	1711	70	24.4	3.2	22.2
XL-WSD-zh	6774	696	9.7	3.5	35.0*

Table 1: Basic statistics for the used datasets. ISO 639 is used to encode language names.  
\* Number of hieroglyphs

and active-dict. We employ the private test subset of **bts-rnc-ru** in our experiments<sup>11</sup>. Bts-rnc-ru also has a predefined train subset, which we use for hyperparameter selection and ablation study.

We use the largest existing Word Sense Disambiguation dataset **XL-WSD** (Pasini et al., 2021) for WSI to extend the number of languages for evaluation. The test subset of XL-WSD provides sense annotations for over 50 thousand instances in 18 languages. We replace the provided sense definitions with cluster labels. We only apply minimal filtering to the test subset of XL-WSD, keeping all the target words with at least two senses and at least three instances. We include only 10 out of 18 languages of XL-WSD because datasets in the other eight languages frequently contain only one instance per sense. We also use train subsets of XL-WSD in English and Slovenian with a stricter filter (at least eight instances per word) as our dev sets.

We employed the dataset in German **DWUG de Sense** (Schlechtweg et al., 2021, Schlechtweg, 2022) which, unlike XL-WSD, is human-annotated. Originally, DWUG de Sense is designed to study the change in word usage in particular senses over time. The annotators were given a task to map target words to one of the predefined sense definitions in instances from "old" (XIX century) and "new" (XX century) corpora. We concatenated the "old" and the "new" parts of the dataset with no further processing<sup>12</sup>. Like in XL-WSD, we replaced sense definitions with cluster labels.

To summarize, we use SE10, SE13, the test sub-

<sup>11</sup>Wiki-wiki dataset contains easy examples and has been solved with a perfect score in the competition and active-dict has a significantly lower instances per target word ratio.

<sup>12</sup>We only removed three instances that are longer than 512 tokens.

set of bts-rnc-ru, the XL-WSD in 10 languages, and DeWUG de Sense for WSI evaluation. Basic statistics for the used datasets are provided in table 1. Notably, the converted datasets (DWUG de Sense and XL-WSD) have a much lower instances per word ratio in comparison to SE10, SE13, and bts-rnc-ru. It makes the task of WSI more challenging, as in many cases clustering must be performed in a very sparse vector space. Also, instances for Bulgarian and Galician in XL-WSD are short. They do not contain full sentences but rather short phrases.

## 4.2. Results

We include a total of 6 configurations in our experiments. We employ four multi-token substitute generators (concat, WCM, WCM-en, and WCM-ru) and combine them with two target injection methods (SDP and +embs). As mentioned above, WCM-en and WCM-ru always yield substitutes in English and Russian, respectively, for context in any language. We exclude WCM-en +embs and WCM-ru +embs from experiments. These configurations require aligned cross-lingual embeddings, which we consider a source of additional error.

### 4.2.1. Comparison to Monolingual Methods

We compare our methods to existing monolingual ones on established WSI benchmarks: bts-rnc-ru, SE10, and SE13. The results are provided in Table 2.

Our SDP-based methods perform on par with SOTA English monolingual approaches in hard clustering on SemEval 2010 and do not fall far behind on SemEval 2013 with trivial hard-to-soft clustering conversion. For +embs the results are noticeably worse. We discuss possible explanations for this fact in Section 5. It is worth mentioning that BERT LSDP (Amrami and Goldberg, 2019), PolyLM Large

(Ansell et al., 2021), XLNET +embs (Arefyev et al., 2020), and our methods have roughly the same number of parameters (350 million) in the underlying neural language models.

In `bts-rnc-ru`, our systems fall behind only of (Arefyev and Bykov, 2021) that uses a fixed number of clusters for all words. This number is fitted on the train set of `bts-rnc`, and is unlikely to transfer to other languages or datasets.

Notably, WCM-en is the second best-performing configuration on the Russian benchmark `bts-rnc-ru` while producing only English substitutes. We confirm that substitutes are indeed in English while the dataset is in Russian by building a simple character-based classifier. Russian and English have different alphabets, therefore we consider a substitute to be in English if all its characters are Latin. Over 99% of the generated substitutes get classified as English. We also performed a visual inspection, confirming this result. WCM-ru is close to WCM-en on English benchmarks with over 99% of substitutes in Russian.

#### 4.2.2. Multilingual WSI evaluation

We also evaluate the proposed methods on our custom WSI benchmarks XL-WSD and DWUG de Sense. We employ a trivial adaptation of the substitute generator from the original BERT LSDP approach as a baseline. In the baseline, XLM-R replaces BERT as an underlying model, but multitoken substitutes are not allowed. Like in the original method, the target injection relies on both the unmasked instance and the non-symmetric dynamic pattern. However, we replace the original clustering approach with ours to allow for a better comparison of the substitute generators.

The results on all the test datasets we use for WSI are presented in Table 3. All of our methods surpass the baseline on each dataset, proving the effectiveness of our approach in a multilingual setup. No configuration is an overall winner.

We include two more configurations in this experiment: **WCM-en SDP-en** and **WCM-ru SDP-ru**. These configurations are identical to WCM-en SDP and WCM-ru SDP, respectively, but do not involve the translation of the dynamic patterns. WCM-en SDP-en uses the pattern in English for all languages, and WCM-ru SDP-ru always uses the pattern in Russian. Due to the properties of WCM-en and WCM-ru they also do not require the use of a multilingual lemmatizer. We can see from Table 3 that **WCM-en SDP-en shows strong WSI results and requires no adaptation to new languages..** We thus consider WCM-en SDP-en to be a truly multilingual WSI system.

System	F-S	V-M	AVG
<b>Concat SDP</b>	<b>67.7</b>	<b>41.7</b>	<b>53.1</b>
<b>Concat +embs</b>	62.5	32.0	44.7
<b>WCM SDP</b>	65.1	39.4	50.7
<b>WCM +embs</b>	64.1	26.9	41.5
<b>WCM-en SDP</b>	66.3	39.0	50.8
<b>WCM-ru SDP</b>	65.5	37.2	49.4
BERT LSDP	<u>71.3</u> $\pm 0.1$	40.4 $\pm 1.8$	53.6 $\pm 1.2$
XLNET +embs			54.2
PolyLM Large*	67.5	<b>43.6</b>	<b>54.3</b>
AutoSense	61.7	9.8	24.5

System	FBC	FNMI	AVG
<b>Concat SDP</b>	<b>65.1</b>	<b>19.6</b>	<b>35.7</b>
<b>Concat +embs</b>	59.6	14.0	28.8
<b>WCM SDP</b>	62.5	16.4	32.0
<b>WCM +embs</b>	61.7	13.7	29.1
<b>WCM-en SDP</b>	64.0	18.8	34.6
<b>WCM-ru SDP</b>	62.2	17.4	32.9
BERT LSDP	64.0 $\pm 0.5$	21.4 $\pm 0.5$	37.0 $\pm 0.5$
XLNET +embs			37.3
PolyLM Large*	<b>66.7</b>	<b>23.7</b>	<b>39.7</b>
ELMo LSDP	57.5	11.3	25.4
AutoSense	61.7	7.96	22.2

System	ARI
<b>Concat SDP</b>	<b>51.6</b>
<b>Concat +embs</b>	49.1
<b>WCM SDP</b>	48.9
<b>WCM +embs</b>	44.2
<b>WCM-en SDP</b>	50.4
<b>WCM-ru SDP</b>	50.2
Arefyev and Bykov (2021)	<b>57.3</b>
Arefyev et al. (2019)	45.1
RUSSE'2018 competition best [1]	33.8
RUSSE'2018 competition baseline [1]	21.3

Table 2: Evaluation results in English on SE10 (top), SE13 (middle) and in Russian on `bts-rnc-ru`-test-private (bottom). The best result in each group is in **bold**, the overall best result is also **underlined**.

\* PolyLM Large was not presented in the original paper (Ansell et al., 2021), the results are from the authors' official repository.

## 5. Semantic Relations between Substitutes and Target Words

To get more insights on substitute generation, we also perform a direct evaluation of our substitute generators on existing lexical substitution datasets: SemEval2007 task 10 (McCarthy and Navigli, 2007) in English, GermEval 2015 (Miller et al., 2015) in German and SemDisFr 2014 (Fabre et al., 2014) in French. The results are presented in Figure 3. The performance of target injection methods on these datasets is opposite to WSI (in line with the results from Arefyev et al., 2020): +embs performs better than SDP "or even" or other dynamic patterns.

To find an explanation for the divergence of results on lexical substitution and WSI, we classify relations between target words and substitutes based on WordNet (Miller, 1995) relations for different

	SE10 SE13			bts-r ru	DWUG de	XL-WSD									
	en	en				bg	de	en	es	fr	gl	it	ko	sl	zh
Baseline	24.2	23.5	35.4	40.4	5.0	1.1	18.4	20.5	20.6	7.6	20.7	11.6	3.9	20.4	
Concat SDP	<b>38.1</b>	<b>30.4</b>	47.7	<b>41.9</b>	<b>11.9</b>	18.0	32.9	<b>34.4</b>	<b>27.8</b>	<b>11.6</b>	<b>23.5</b>	<b>24.9</b>	<b>18.5</b>	<b>30.3</b>	
WCM SDP	34.9	25.4	<b>48.9</b>	39.5	10.0	<b>23.4</b>	<b>33.0</b>	29.0	27.3	10.8	22.0	17.9	13.5	26.8	
WCM-en SDP	35.8	28.8	<b>50.4</b>	41.2	<b>18.1</b>	22.7	31.9	37.0	24.7	9.3	<b>29.9</b>	<b>25.6</b>	<b>17.8</b>	<b>27.8</b>	
WCM-en SDP-en	35.8	28.8	47.5	<b>50.8</b>	16.8	<b>25.1</b>	31.9	<b>38.0</b>	<b>33.1</b>	<b>10.6</b>	27.2	22.5	15.7	26.8	
WCM-ru SDP	<b>33.3</b>	<b>28.1</b>	50.2	46.3	<b>14.8</b>	<b>20.5</b>	<b>31.4</b>	<b>39.3</b>	<b>33.7</b>	<b>10.5</b>	<b>30.6</b>	<b>25.6</b>	<b>14.3</b>	<b>27.3</b>	
WCM-ru SDP-ru	28.4	23.1	50.2	<b>46.4</b>	13.5	19.7	30.0	33.5	21.2	10.1	24.4	23.4	12.3	25.2	

Table 3: WSI evaluation results. The metric is Adjusted Rand Index (ARI, Hubert and Arabie, 1985).

target injection methods. WordNet provides "synonym" relations of the synsets as well as hyponymy relations. We classify the relations between the target word and the substitute into the following classes:

*Direct hyponym, direct hypernym, or synonym* if the two words have synsets in WordNet with corresponding relations. For example "onion" is a direct hyponym of "vegetable", "vegetable" is a direct hypernym of "onion", and "middle" is a synonym of "center".

*Transitive hyponym or transitive hypernym* if there is a path of only hyponym or hypernym relations from the target word to the substitute. "Onion" is a transitive hyponym of "food" (onion, vegetable, food), and "food" is a transitive hypernym of "onion".

*Co-hyponym* if two words share a direct hypernym, *co-hyponym-3* if they share a hypernym of a hypernym. "Onion" and "garlic" are co-hyponyms, and "onion" and "apple" are co-hyponyms-3.

Figure 4 represents the distribution of relation classes for "or even" and +embs target injection methods on all instances from SE10 where target words are nouns. SDP "or even" yields substitutes with "unknown relation" more frequently. However, SDP substitutes still allow for better clustering according to WSI evaluation.

We review the most *discriminative* substitutes for target words in SE10 as a qualitative analysis of the generated substitutes. A substitute is discriminative if it appears frequently in instances for a given sense and rarely in other senses. More formally, the substitutes for each sense get ranked by dividing the number of occurrences of the substitute in this sense by the number of occurrences of that substitute in other senses, according to gold sense clustering.

The top 3 discriminative substitutes for several senses of three target words from SE10 are presented in Table 4. Firstly, we confirm that WCM-ru yields cross-lingual substitutes for instances in the WSI datasets, and these substitutes fit the context quite well. Secondly, we find examples of substitutes that seem useful for WSI but which are not lexical substitutes in a traditional sense. E.g., "blood" and "animal" indeed help distinguish a "blood cell"




Figure 3: Evaluation of substitute generators on the lexical substitution datasets. Blue is the concat substitute generator and orange is the WCM.

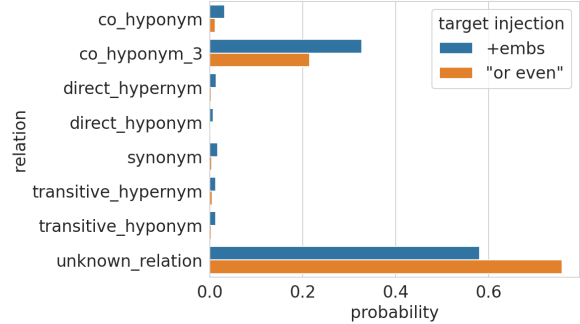


Figure 4: Distribution of Wordnet hyponymy relations for different target injection methods. Only top-20 substitutes are used for each instance.

from a "prison cell" but it is hard to describe the exact relation between the target word "cell" and these substitutes.

In general, we conclude that the superior performance of SDPs on WSI can be explained by the presence of words that are not lexical substitutes by conventional definition (they cannot *replace* the target word in the sentence) but tend to co-occur with the target word. It may also explain the optimal WSI  $k$  in +embs being lower than in SDP. While +embs is better at producing closely related substitutes, SDP substitutes are more diverse. Words usually do not have many closely related substitutes and +embs quickly "runs out" of adequate substitutes.

## 6. Conclusion

By harnessing the capabilities of the multilingual masked language model XLM-R and carefully



Target word	Context Example	Concat +embs	Concat SDP	WCM-ru SDP
body (33 instances)	... the highest legislative <b>body</b> of China ...	institution govern organization	organization committee office	чиновник(official) учреждение(establishment) государство(state)
body (70 instances)	... the <b>body's</b> immune response...	shape, bodily, muscle	personality, muscle, health	мышца(muscle), взгляд (sight), красота(beauty)
body (78 instances)	... his <b>body</b> bore several wounds...	dead, victim, remain	victim, debris, grave	имущество(property), флаг (flag), контейнер(container)
cell (11 instances)	... to its prison <b>cell</b> shortage ...	jail, room, house	apartment, cellar, detention	коридор(corridor), гараж (garage), кабинет(cabinet)
cell (58 instances)	... of white blood <b>cells</b> involved in the body's immune ...	tissue tumor stem	blood virus animal	вирус(virus), организм(organism) новообразование(tumor)
display (11 instances)	... on a TV <b>display</b> ...	lcd, monitor, system	LCD, digital, component	аналог(analogue), дисплей (display) монитор (monitor)
display (19 instances)	... diamonds on <b>display</b> ...	sale, showcase, view	sell, presentation, public	продажа (sale), купить (to buy), продать (to sell)

Table 4: Top-3 discriminative substitutes for Concat +embs, Concat SDP, and WCM-ru SDP for three nouns from the SE10 dataset.

adapting existing monolingual WSI methods, we were able to build a multilingual WSI system that does not require any adaptation to be used in a new language. We performed extensive experiments across several languages and datasets to prove that our approach is robust and performs on par with the state-of-the-art monolingual methods. We also discovered that by simply finetuning XLM-R with the WCM objective in a single language, we obtain a model capable of producing cross-lingual lexical substitutes, and these substitutes allow for adequate WSI performance.

## 7. Limitations

In terms of languages, our system has only one minimal requirement. The language of an instance must be one of 100 languages supported by XLM-R. However, we have only tested our systems on limited academic benchmarks with sense annotations. These benchmarks may not reflect all the nuances of word usage in real corpora.

Our methods tackle WSI as a hard clustering task, which suggests distinct borders between the senses of an ambiguous word. The existence of such borders is hardly the case for many words in the natural languages (Jurgens and Klapaftis, 2013). More than that, our WSI systems always consider the target word to be ambiguous, i.e., having at least two senses. We do not tackle the problem of ambiguity detection in our work. Overall, we believe that deeper research into clustering methods for WSI is required to use current WSI systems "in the wild".

## 8. Acknowledgements

Nikolay Arefyev has received funding from the European Union's Horizon Europe research and innovation program under Grant agreement No 101070350 (HPLT).

## 9. Bibliographical References

### References

- Reinald Kim Amplayo, Seung-won Hwang, and Min Song. 2019. Autosense model for word sense induction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6212–6219.
- Asaf Amrami and Yoav Goldberg. 2018. Word sense induction with neural bilm and symmetric patterns. *arXiv preprint arXiv:1808.08518*.
- Asaf Amrami and Yoav Goldberg. 2019. Towards better substitution-based word sense induction. *arXiv preprint arXiv:1905.12598*.
- Alan Ansell, Felipe Bravo-Marquez, and Bernhard Pfahringer. 2021. PolyIm: Learning about polysemy through language modeling. *arXiv preprint arXiv:2101.10448*.
- Nikolay Arefyev, Boris Sheludko, and Alexander Panchenko. 2019. Combining lexical substitutes in neural word sense induction. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 62–70.
- Nikolay Arefyev, Boris Sheludko, Alexander Podolskiy, and Alexander Panchenko. 2020. Always

- keep your target in mind: Studying semantics and improving performance of neural lexical substitution. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1242–1255.
- Nikolay V Arefyev and Dmitry A. Bykov. 2021. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the International Conference on Computational Linguistics and Intellectual Technologies (Dialogue)*.
- Osman Başkaya, Enis Sert, Volkan Cirik, and Deniz Yuret. 2013. Ai-ku: Using substitute vectors and co-occurrence modeling for word sense induction and disambiguation. In *Second Joint Conference on Lexical and Computational Semantics (\*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 300–306.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc."
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching word vectors with subword information. *arXiv preprint arXiv:1607.04606*.
- T. Caliński and J Harabasz. 1974. [A dendrite method for cluster analysis](#). *Communications in Statistics*, 3(1):1–27.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Matan Eyal, Shoval Sadde, Hillel Taub-Tabib, and Yoav Goldberg. 2022. Large scale substitution-based word sense induction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4738–4752.
- Lawrence Hubert and Phipps Arabie. 1985. [Comparing partitions](#). *Journal of Classification*, 2(1):193–218.
- Alexandros Komninos and Suresh Manandhar. 2016. Structured generative models of continuous features for word sense induction. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3577–3587.
- Mikhail Korobov. 2015a. [Morphological analyzer and generator for russian and ukrainian languages](#). In Mikhail Yu. Khachay, Natalia Konstantinova, Alexander Panchenko, Dmitry I. Ignatov, and Valeri G. Labunets, editors, *Analysis of Images, Social Networks and Texts*, volume 542 of *Communications in Computer and Information Science*, pages 320–332. Springer International Publishing.
- Mikhail Korobov. 2015b. Morphological analyzer and generator for russian and ukrainian languages. In *Analysis of Images, Social Networks and Texts: 4th International Conference, AIST 2015, Yekaterinburg, Russia, April 9–11, 2015, Revised Selected Papers 4*, pages 320–332. Springer.
- Jey Han Lau, Paul Cook, and Timothy Baldwin. 2013. unimelb: Topic modelling-based word sense induction for web snippet clustering. In *Second Joint Conference on Lexical and Computational Semantics (\*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 217–221.
- Matthew E. 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, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A Python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Dominik Schlechtweg. 2022. *Human and Computational Measurement of Lexical Semantic Change*. Ph.D. thesis, University of Stuttgart.
- Jing Wang, Mohit Bansal, Kevin Gimpel, Brian D Ziebart, and Clement T Yu. 2015. A sense-topic model for word sense induction with unsupervised data enrichment. *Transactions of the Association for Computational Linguistics*, 3:59–71.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pre-training for language understanding. *Advances in neural information processing systems*, 32.

## 10. Language Resource References

Aksenova, Anna and Gavrishina, Ekaterina and Rykov, Elisei and Kutuzov, Andrey. 2022. *RuDSI: Graph-based Word Sense Induction Dataset for Russian*. PID <https://github.com/kategavrishina/RuDSI>.

Fabre, Cécile and Hathout, Nabil and Ho-Dac, Lydia-Mai and Morlane-Hondère, François and Muller, Philippe and Sajous, Franck and Tanguy, Ludovic and Van de Cruys, Tim. 2014. *Présentation de l'atelier SemDis 2014 : sémantique distributionnelle pour la substitution lexicale et l'exploration de corpus spécialisés*. Laboratoire Cognition, Langues, Langage, Ergonomie (CLLE) UMR 5263. PID <http://redac.univ-tlse2.fr/datasets/semdis-gold/lexicalsubstitution/>.

Jurgens, David and Klapaftis, Ioannis. 2013. *SemEval-2013 Task 13: Word Sense Induction for Graded and Non-Graded Senses*. Association for Computational Linguistics. PID <https://lcl.uniroma1.it/wsdeval/evaluation-data>.

Kremer, Gerhard and Erk, Katrin and Padó, Sebastian and Thater, Stefan. 2014. *What substitutes tell us-analysis of an "all-words" lexical substitution corpus*. Institute for Natural Language Processing, University of Stuttgart. PID <https://www.ims.uni-stuttgart.de/en/research/resources/corpora/coincod/>.

Manandhar, Suresh and Klapaftis, Ioannis and Dligach, Dmitriy and Pradhan, Sameer. 2010. *SemEval-2010 Task 14: Word Sense Induction & Disambiguation*. Association for Computational Linguistics. PID <https://doi.org/10.5281/zenodo.5638549>.

McCarthy, Diana and Navigli, Roberto. 2007. *SemEval-2007 Task 10: English Lexical Substitution Task*. Association for Computational Linguistics.

Miller, George A. 1995. *WordNet: a lexical database for English*. ACM New York, NY, USA. PID <https://wordnet.princeton.edu/>.

Miller, Tristan and Benikova, Darina and Abualhajja, Sallam. 2015. *GermEval 2015: LexSub-A*

*shared task for German-language lexical substitution*. (1) Humboldt University Berlin, (2) FG Language Technology, (3) Ubiquitous Knowledge Processing Lab (UKP-TUDA) Dept. of Computer Science, Technische Universität Darmstadt (4) Ubiquitous Knowledge Processing Lab (UKP-DIPF) German Institute for Educational Research and Educational Information. PID <https://www.nothingisreal.com/germeval2015/datasets.html>.

Panchenko, Alexander and Lopukhina, Anastasia and Ustalov, Dmitry and Lopukhin, Konstantin and Arefyev, Nikolay and Leontyev, Alexey and Loukachevitch, Natalia. 2018. *RUSSE'2018: A Shared Task on Word Sense Induction for the Russian Language*. RSUH. PID <https://russe.nlpub.org/2018/wsi/>.

Pasini, Tommaso and Raganato, Alessandro and Navigli, Roberto. 2021. *XL-WSD: An Extra-Large and Cross-Lingual Evaluation Framework for Word Sense Disambiguation*. Sapienza NLP group. PID <https://sapienzanlp.github.io/xl-wsd/>.

Dominik Schlechtweg, Nina Tahmasebi, Simon Hengchen, Haim Dubossarsky, and Barbara McGillivray. 2021. Dwug: A large resource of diachronic word usage graphs in four languages. *arXiv preprint arXiv:2104.08540*.

Wenzek, Guillaume and Lachaux, Marie-Anne and Conneau, Alexis and Chaudhary, Vishrav and Guzmán, Francisco and Joulin, Armand and Grave, Edouard. 2020. *CCNet: Extracting High Quality Monolingual Datasets from Web Crawl Data*. European Language Resources Association. PID <https://huggingface.co/datasets/cc100>.

### A. Comparison of multi-token substitute generators and target injection methods on the dev sets

We evaluated 18 system configurations on the dev sets (russe-bts-rnc-train, XL-WSD-en-train and XL-WSD-sl-train) with different target injection methods (4 dynamic patterns and +embs) and multi-token substitute generators (WCM-multi, WCM-en, WCM-ru, and concat). The used dynamic patterns are presented in Table 5. The baseline method is described in Section 4.2.

Our primary metric, **ARI** (Adjusted Rand Index, Hubert and Arabie, 1985), is a clusterization metric and may vary significantly with changes in clustering hyperparameters (e.g., in the number of clusters). We also measure **maxARI** to evaluate the quality of substitute vectors themselves. MaxARI

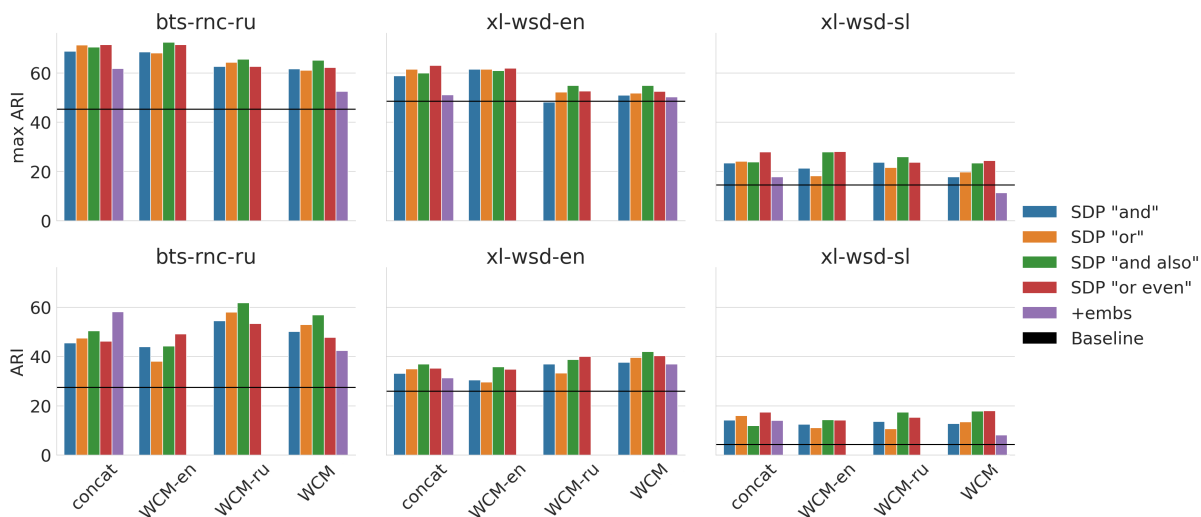


Figure 5: Evaluation of different configurations of our system on the WSI dev sets.

is the upper bound for ARI achievable given a vectorizer, a clusterer, and a grid on clustering hyperparameters. It is calculated with the clustering hyperparameters selected individually for each target word to maximize ARI. We observe maxARI to be more stable to variations in substitute generator parameters compared to regular ARI.

The results of the evaluation are presented in Figure 5. The first observation is that all the configurations surpass the baseline in ARI on all the dev datasets. The same can be said for maxARI in 16 out of 18 configurations. In general, SDPs perform better than +embs. We can see that more complicated patterns "and also" and "or even" perform slightly better than simpler patterns "and" and "or".

## B. Selection of the number of substitutes

We optimized the number of substitutes ( $k$ ) in our methods based on maxARI on the train subset of bts-rnc-ru for each target injection method. The results are presented in Figure 6. The WSI performance for SDP reaches a plateau only at  $k \approx 100$ .

Name	Symmetric Patterns
"and"	<T> and <mask> <mask> and <T>
"or"	<T> or <mask> <mask> or <T>
"and also"	<T> (and also <mask>) <mask> (and also <T>)
"or even"	<T> (or even <mask>) <mask> (or even <T>)

Table 5: The symmetric dynamic patterns used in our experiments.

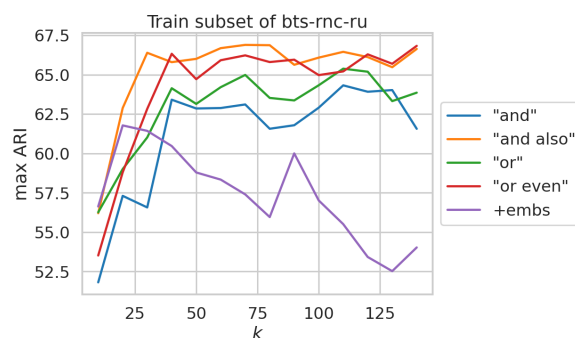


Figure 6: The relation between  $k$  and maxARI on the train subset of bts-rnc-ru for different target injection methods. Concat is used for multi-token substitute generation.

In the case of +embs, however, we found that  $k = 20$  is optimal. The performance significantly declines for larger values of  $k$ .

## C. The importance of multi-subword substitutes for WSI

Figure 7 shows the WSI performance depending on the length of substitutes generated. When comparing the more stable maxARI metric, we see small but consistent benefits from using substitutes of variable length for SE10, SE13, bts-rnc-ru, DWUG de Sense. Combining 1-, 2-, and 3-subword substitutes with the concat method usually performs best. For the same datasets generating only single-subword substitutes results in a clear drop in performance. For XL-WSD the results are mixed.




Figure 7: The relation between WSI metrics (ARI and maxARI) and the length (in subwords) of substitutes generated. The Concat method is used when combining substitutes of different lengths.

#### D. MaxARI on the test sets

Table 6 extends Table 3 and includes maxARI evaluation for the test sets. Overall, the differences in maxARI between configurations are more subtle, compared to regular ARI. The gap between ARI and maxARI is more noticeable for XL-WSD, than for other benchmarks. It means that Calinski-Harabasz score is worse at determining the optimal number of clusters in XL-WSD compared to other datasets. We think it can be explained by differences in statistics for these datasets, highlighted in Section 4.1. Specifically, they contain many words with very few examples, all belonging to different senses. Such words can be ideally clustered when swiping through the number of clusters from 2 to 9.

	SE10	SE13	bts-r	DWUG	XL-WSD									
	en	en	ru	de	bg	de	en	es	fr	gl	it	ko	sl	zh
Baseline	24.2	23.5	35.4	40.4	5.0	1.1	18.4	20.5	20.6	7.6	20.7	11.6	3.9	20.4
Concat SDP	<b>38.1</b>	<b>30.4</b>	47.7	<b>41.9</b>	<b>11.9</b>	18.0	32.9	<b>34.4</b>	<b>27.8</b>	<b>11.6</b>	<b>23.5</b>	<b>24.9</b>	<b>18.5</b>	<b>30.3</b>
WCM SDP	34.9	25.4	<b>48.9</b>	39.5	10.0	<b>23.4</b>	<b>33.0</b>	29.0	27.3	10.8	22.0	17.9	13.5	26.8
WCM-en SDP	35.8	28.8	<b>50.4</b>	41.2	<b>18.1</b>	22.7	31.9	37.0	24.7	9.3	<b>29.9</b>	<b>25.6</b>	<b>17.8</b>	<b>27.8</b>
WCM-en SDP-en	35.8	28.8	47.5	<b>50.8</b>	16.8	<b>25.1</b>	31.9	<b>38.0</b>	<b>33.1</b>	<b>10.6</b>	27.2	22.5	15.7	26.8
WCM-ru SDP	<b>33.3</b>	<b>28.1</b>	50.2	46.3	<b>14.8</b>	<b>20.5</b>	<b>31.4</b>	<b>39.3</b>	<b>33.7</b>	<b>10.5</b>	<b>30.6</b>	<b>25.6</b>	<b>14.3</b>	<b>27.3</b>
WCM-ru SDP-ru	28.4	23.1	50.2	<b>46.4</b>	13.5	19.7	30.0	33.5	21.2	10.1	24.4	23.4	12.3	25.2
Baseline	34.9	31.8	46.0	51.5	21.8	26.6	47.8	50.7	50.3	31.2	49.3	38.2	13.5	41.1
Concat SDP	<b>44.2</b>	<b>38.2</b>	<b>58.0</b>	<b>60.3</b>	<b>32.3</b>	<b>33.8</b>	<b>57.6</b>	<b>67.7</b>	51.5	<b>37.4</b>	<b>54.4</b>	<b>49.3</b>	<b>27.7</b>	<b>49.5</b>
WCM SDP	43.3	35.4	56.8	57.3	27.6	33.4	56.9	58.4	<b>53.7</b>	33.6	54.0	41.2	23.9	45.1
WCM-en SDP	44.6	37.9	<b>60.5</b>	60.7	<b>38.3</b>	35.5	57.3	<b>65.0</b>	53.8	36.7	<b>60.1</b>	<b>48.4</b>	<b>28.3</b>	46.3
WCM-en SDP-en	44.6	37.9	57.1	<b>64.8</b>	38.0	<b>36.9</b>	57.3	63.6	<b>59.6</b>	<b>37.3</b>	56.5	46.7	24.8	<b>46.6</b>
WCM-ru SDP	<b>41.7</b>	<b>36.8</b>	58.1	<b>60.6</b>	<b>34.9</b>	<b>36.2</b>	55.1	<b>66.8</b>	<b>64.8</b>	<b>36.2</b>	<b>53.9</b>	<b>50.3</b>	<b>25.0</b>	<b>46.4</b>
WCM-ru SDP-ru	40.0	33.4	58.1	58.1	32.4	36.1	<b>56.1</b>	60.1	51.5	35.1	52.2	45.5	23.0	44.1

Table 6: Extended WSI evaluation results. The metrics are ARI (top) and maxARI (bottom).