

# Slovene SuperGLUE Benchmark: Translation and Evaluation

Aleš Žagar, Marko Robnik-Šikonja

University of Ljubljana, Faculty of Computer and Information Science, Večna pot 113, 1000 Ljubljana,  
ales.zagar@fri.uni-lj.si, marko.robnik@fri.uni-lj.si

## Abstract

We present SuperGLUE benchmark adapted and translated into Slovene using a combination of human and machine translation. We describe the translation process and problems arising due to differences in morphology and grammar. We evaluate the translated datasets in several modes: monolingual, cross-lingual, and multilingual, taking into account differences between machine and human translated training sets. The results show that the monolingual Slovene SloBERTa model is superior to massively multilingual and trilingual BERT models, but these also show a good cross-lingual performance on certain tasks. The performance of Slovene models still lags behind the best English models.

**Keywords:** benchmarks, evaluation, transformer models, natural language understanding, cross-lingual models, multilingual models

## 1. Introduction

To measure the progress across the fast evolving area of natural language processing (NLP), several prominent benchmarking suites have been proposed, such as SentEval (Conneau et al., 2018), GLUE (Wang et al., 2019b), and SuperGLUE (Wang et al., 2019a). SentEval aims to evaluate sentence encoders, while GLUE and its more challenging successor, SuperGLUE, evaluate general language understanding. These two benchmark suites contain public training sets and private test sets that can be assessed through an evaluation server giving separate scores for each of the tasks and an overall single-number score. The tasks in SuperGLUE are diverse and comprised of question answering (QA), natural language inference (NLI), coreference resolution, and word sense disambiguation (WSD). Non-expert humans evaluated all the tasks to provide a human performance baseline.

Despite recent criticism of the way how the evaluation scores in these benchmarks are computed (e.g., the arithmetic mean of separate metrics is used for all tasks of different complexity and different sizes of testing sets) (Shavrina and Malykh, 2021; Kasai et al., 2021), there is little doubt that the datasets contained in the benchmarks significantly contribute to the progress in the NLP area. Unfortunately, most of the research is focused on English, which often limits the generality of the approaches and does not address the full scale of language complexity. Still, we can observe significant efforts to make relevant datasets cross-lingual or adapted to more languages. Broad multilingual benchmarking suites, such as XTREME (Hu et al., 2020) and XGLUE (Liang et al., 2020), provide machine-translated datasets for several relevant tasks and languages. Complementary to that, we present a combined machine-human translation of SuperGLUE benchmarking suite for less-resourced Slovene language. We describe the difficulties arising in the translation and provide the first evaluation

of the datasets using large pretrained BERT-like language models. Our evaluation encompasses monolingual, cross-lingual and multilingual approaches, comparison of human and machine-translated parts, and comparison of two state-of-the-art machine translation systems.

The paper is structured into further five sections. In Section 2, we present related work on non-English benchmarking suites. In Section 3, we describe the translation process and the final SuperGLUE datasets in Slovene. We describe three evaluation settings (monolingual, cross-lingual, and multilingual) in Section 4, and present the results in Section 5. Section 6 presents conclusions, limitations of our work, and ideas for further improvements. To make the paper self-contained and ease understanding of the different datasets in the SuperGLUE benchmark, we follow Wang et al. (2019a) and present an example for each of the tasks in Appendix.

## 2. Related work

In the area of NLP, there are several prominent benchmarking suites, the most prominent being GLUE (Wang et al., 2019b) and SuperGLUE (Wang et al., 2019a). While most benchmarking datasets are available in English, there are a few notable attempts to extend them to other languages.

XTREME (Hu et al., 2020) is a multi-task benchmark for evaluating the cross-lingual representations across 40 languages and 9 tasks. It focuses on the zero-shot cross-lingual transfer and provides large English training sets, much smaller testing sets in target languages (from 5 to 40 human translated or checked instances, the rest are machine translated without human interventions). The suite covers tasks such as natural language inference (NLI), paraphrasing, part-of-speech (POS) tagging, named entity recognition (NER), question answering, and parallel sentence extraction and detection. XGLUE (Liang et al., 2020) is intended to evaluate

cross-lingual pretrained models on 11 tasks: NER, POS-tagging, question answering, paraphrasing, ad relevance, web page relevance, question-answer matching, question generation, and news title generation. Similarly to XTREME, the suite provides training data in English and tests the cross-lingual models on testing data in several target languages (from 3 to 18).

SuperGLUE is likely the most prominent benchmarking suite for English. It follows the design of the GLUE suite, consisting of a public leaderboard for eight tasks, which are evaluated individually and jointly. The suite provides public training and development datasets, while testing data is hidden and only used to evaluate predictions submitted to the leaderboard. The benchmark includes QA, NLI, coreference resolution, and WSD tasks, for which non-expert human baselines are provided.

Following the English example, Russian SuperGLUE (Shavrina et al., 2020) was developed independently or manually translated from English SuperGLUE. It contains nine tasks: linguistic diagnostic, WSD, QA, NLI, and coreference resolution.

Slovene language possesses significantly fewer resources than English or Russian and is not included in the XTREME and XGLUE massively cross-lingual benchmarks. We aimed to provide a modern natural language understanding suite for this less-resourced language using modest resources at our disposal. For that purpose, we followed the English SuperGLUE design and translated the entire datasets (except WiC) to Slovene, mainly using machine translation and in small part human translation (120,000 words). This allows testing of monolingual, cross-lingual and multilingual approaches. Keeping the format of the original SuperGLUE, our datasets can be evaluated with the original SuperGLUE leaderboard and can be compared to English baselines and state-of-the-art.

### 3. Slovene SuperGLUE translation

The SuperGLUE benchmark consists of 8 tasks. Boolean Questions (BoolQ) is a QA task where each instance contains a short text and a question with a true/false answer related to it. CommitmentBank (CB) is an NLI task in which each example contains a short passage and hypothesis, for which a solver has to decide the degree of agreement with the truth of the hypothesis. The Choice Of Plausible Alternatives (COPA) is a causal reasoning task where a system must determine which of the proposed two sentences is the cause for the premise. Multi-Sentence Reading Comprehension (MultiRC) is a QA task, consisting of a paragraph, questions, and a list of true/false answers. Reading Comprehension with Commonsense Reasoning Dataset (ReCoRD) is a multiple-choice QA task where one must choose the correct entities from a paragraph to fill in a masked word in a given query. Recognizing Textual Entailment (RTE) is an NLI task in which a solver must determine whether the hypothesis

entails a given text or not. Word-in-Context (WiC) is a word sense disambiguation task where two contexts with an annotated polysemous word are presented to a system that must determine whether the word senses match or not. Winograd Schema Challenge (WSC) is a coreference resolution task in which a solver determines whether two annotated words refer to the same entity or not. Refer to Table 7 for an example of each task.

To evaluate cross-lingual transfer and test specifics of morphologically rich languages, we translated the SuperGLUE datasets to Slovene. Due to limited funds, we partially used human translation (HT) and partially machine translation (MT). Altogether, approximately 120,000 words were human translated. Some datasets are too large (BoolQ, MultiRC, ReCoRD, RTE) to be fully human translated. We thus provide ratios between the human translated and the original English sizes in Table 1. For MT from English to Slovene, we used the GoogleMT Cloud service. In our evaluation, we use six of the original eight tasks. As explained below, we excluded ReCoRD and WiC, which need extensive manual editing (ReCoRD) or creation from scratch (WiC).

Table 1: Number of instances in the original English and translated Slovene SuperGLUE tasks. HT stands for human translation and MT for machine translation. The ‘‘ratio’’ indicates the ratio between the number of human translated instances and all instances. The ‘/’ sign indicates that an indicated version is unavailable due to reasons explained in the text.

Dataset	split	English	HT	ratio	MT
<b>BoolQ</b>	train	9427	92	0.0098	yes
	val	3270	18	0.0055	yes
	test	3245	30	0.0092	yes
<b>CB</b>	train	250	250	1.0000	yes
	val	56	56	1.0000	yes
	test	250	250	1.0000	yes
<b>COPA</b>	train	400	400	1.0000	yes
	val	100	100	1.0000	yes
	test	500	500	1.0000	yes
<b>MultiRC</b>	train	5100	15	0.0029	yes
	val	953	3	0.0031	yes
	test	1800	30	0.0167	yes
<b>ReCoRD</b>	train	101000	60	0.0006	/
	val	10000	6	0.0006	/
	test	10000	30	0.0030	/
<b>RTE</b>	train	2500	232	0.0928	yes
	val	278	29	0.1043	yes
	test	300	29	0.0967	yes
<b>WiC</b>	train	6000	/	/	/
	val	638	/	/	/
	test	1400	/	/	/
<b>WSC</b>	train	554	554	1.0000	/
	val	104	104	1.0000	/
	test	146	146	1.0000	/

We decided to use only the HT translated part of test

sets in our evaluations to avoid noise due to translations. This makes some test sets much smaller compared to English test sets. We did not include ReCoRD in the Slovene benchmark due to the low quality of the obtained test set, consisting of confusing and ambiguous examples. Further, there are differences between English and Slovene ReCoRD tasks due to the morphological richness of Slovene. Namely, in Slovene, the correct declension of a query is often not present in the text, making it impossible to provide the right answer just by copying the contents from the provided paragraph. Finally, similarly to WSC (discussed below), ReCoRD is also affected by the problem of translating HTML tags with GoogleMT. We leave these challenges for further work.

The WiC task cannot be translated and would have to be conceived anew because it is impossible to transfer the same set of meanings of a given word from English to a target language. Taking the WiC example in Table 7, we note that the word *board* in two different contexts translates to two completely different words in Slovene (Context 1: Bivanje in *hrana*. Context 2: Čez okna je nabil *deske*.), which makes the task inappropriate. We intend to prepare Slovene WiC task in the future.

For WSC dataset, we only provide HT as its MT would require human assistance and verification. First, GoogleMT translations cannot handle the correct placement of HTML tags indicating coreferences. The second reason is that in Slovene coreferences can be expressed with verbs, while coreferences in English are mainly nouns, proper names and pronouns. This makes the task different in Slovene compared to English. On the one hand, the task is more difficult in Slovene because solutions cover more types of words; on the other hand, the Slovene verbs might reveal the coreference information for some instances. We included the HT translated WSC in the Slovene SuperGLUE benchmark but intend to systematically analyze the differences in further work.

## 4. Evaluation settings

SuperGLUE benchmark is extensively used to compare large pretrained models in English<sup>1</sup>. In contrast to that, we concentrate on the Slovene translation of the SuperGLUE tasks and new cross-lingual and multi-lingual opportunities which arise from that.

We analyze the performance of four BERT models capable to handle Slovene: monolingual Slovene SloBERTa<sup>2</sup> (Ulčar and Robnik-Šikonja, 2021), trilingual (Croatian-Slovene-English) CroSloEngual BERT<sup>3</sup> (Ulčar and Robnik-Šikonja, 2020), massively multilingual mBERT (Devlin et al., 2019) (bert-base-multilingual-cased<sup>4</sup>), and XLM-R (Conneau et al.,

2019) (xlm-roberta-base<sup>5</sup>).

The single-number overall average score (e.g., Avg in the second column of Table 2) comprises equally weighted tasks. For English, there are eight tasks; for Slovene, there are all six translated tasks: BoolQ, CB, COPA, MultiRC, RTE, and WSC. In tasks with multiple metrics, we averaged those metrics to get a single task score. For details on how the score is calculated for each task, see (Wang et al., 2019a).

We test models in three settings. The monolingual setting uses the same language (Slovene or English) for training and testing. In the cross-lingual setting, we tested the cross-lingual models (CroSloEngual, mBERT, and XLM-R) and transfer between English and Slovene datasets in both directions. In the multilingual setting, the models were trained on the combined full size English and Slovene data. The results for the three settings as well as a comparison of two MT systems (GoogleMT and DeepL) are reported in Section 5.

We fine-tuned the available Slovene BERT models on SuperGLUE tasks using the Jiant tool (Phang et al., 2020). We used a single-task learning setting for each task and fine-tuned models for 100 epochs, with the initial learning rate of  $10^{-5}$ . Each model was fine-tuned using the dataset corresponding to one of the three settings.

## 5. Results

We first report the results for each of the three settings (monolingual, cross-lingual, and multilingual) separately. We end the section with the comparison between human and machine-translated data and the comparison of machine translation systems.

### 5.1. Monolingual results

Our monolingual analysis compares different Slovene prediction models on the complete datasets, composed of existing human translated instances while the remaining instances are machine-translated. Table 2 shows the results together with several baselines trained on the original English datasets. Some comparisons to English baselines are questionable as the Slovene models are trained on only a small fraction of better quality HT data (BoolQ, MultiRC) and tested on smaller set of HT test data. In the case of the BERT++ model, the English model was additionally pretrained with transfer learning tasks similar to a target one (CB, RTE, BoolQ, COPA). In terms of datasets, the only fair comparison is possible with the CB, COPA, and WSC. Considering the Avg scores in Table 2, the monolingual SloBERTa is the best performing Slovene model. On average, all Slovene BERT models perform better than the Most Frequent baseline. Concerning individual tasks, none of the Slovene models exceeds the Most Frequent baseline in the MultiRC task. SloBERTa was significantly better than the rest of the models in

<sup>1</sup><https://super.gluebenchmark.com/leaderboard>

<sup>2</sup><https://huggingface.co/EMBEDDIA/sloberta>

<sup>3</sup><https://huggingface.co/EMBEDDIA/crosloengual-bert>

<sup>4</sup><https://huggingface.co/bert-base-multilingual-cased>

<sup>5</sup><https://huggingface.co/xlm-roberta-base>

Table 2: The SuperGLUE benchmarks in English (upper part) and Slovene (lower part). All English results are taken from (Wang et al., 2019a). The best score for each task and language is in **bold**. The best average scores (Avg) for each language are underlined. \*MultiRC has multiple test sets released on a staggered schedule, and these results evaluate on an installation of the test set that is a subset of ours.

Task	Avg	BoolQ	CB	COPA	MultiRC	ReCoRD	RTE	WiC	WSC
Models/Metrics		Acc.	F1/Acc.	Acc.	F1 <sub>a</sub> /EM	F1/EM	Acc.	Acc.	Acc.
Most Frequent	45.7	62.3	21.7/48.4	50.0	61.1/0.3	33.4/32.5	50.3	50.0	<b>65.1</b>
CBoW	44.7	62.1	49.0/71.2	51.6	0.0/0.4	14.0/13.6	49.7	53.0	<b>65.1</b>
BERT	69.3	77.4	75.7/83.6	70.6	70.0/24.0	<b>72.0/71.3</b>	71.6	<b>69.5</b>	64.3
BERT++	<u>73.3</u>	<b>79.0</b>	<b>84.7/90.4</b>	<b>73.8</b>	<b>70.0/24.1</b>	<b>72.0/71.3</b>	<b>79.0</b>	<b>69.5</b>	64.3
Human (est.)	89.8	89.0	95.8/98.9	100.0	81.8*/51.9*	91.7/91.3	93.6	80.0	100.0
Most Frequent	49.1	63.3	21.7/48.4	50.0	<b>76.4/0.6</b>	-	58.6	-	65.8
SloBERTa	<u>63.9</u>	66.6	<b>74.0/76.8</b>	<b>61.8</b>	62.7/21.9	-	<b>62.1</b>	-	<b>73.3</b>
CroSloEngual	57.8	66.6	62.1/72.4	58.2	56.7/15.6	-	<b>62.1</b>	-	56.2
mBERT	59.1	70.0	66.6/73.6	54.2	57.4/16.3	-	<b>62.1</b>	-	61.6
XML-R	58.7	<b>76.7</b>	66.2/73.2	50.0	55.3/13.9	-	55.2	-	65.8

CB, COPA, and WSC, while XML-R was the best on BoolQ.

Compared to English models, the best Slovene model (SloBERTa) achieved better results on WSC. It seems that none of the English models learned anything from WSC (they are below the Most Frequent baseline), but the SloBERTa model achieved the score of 73.3 (the Most Frequent baseline gives 65.8). The success of SloBERTa on WSC might stem from the morphology of Slovene verbs, which include the information on the gender; this information is beneficial in coreference resolution and makes some instances easier in Slovene than in English. Nevertheless, there is still a large gap compared to human performance. All Slovene models showed good performance on CB and fell between English CBoW and BERT.

## 5.2. Cross-lingual results

In the cross-lingual scenario, we tested the three multilingual BERT models (mBERT, CroSloEngual, XLM-R) and the transfer between English and Slovene datasets (both directions). For Slovene as the source language, we used the available human translated examples. To make the comparison balanced, we only used the same examples from the English datasets. We tested both zero-shot transfer (no training data in the target language) and few-shot transfer. In the few-shot training, we used 10 additional examples from the target language for each task. We randomly sampled these 10 examples for 5 times and reported averages to achieve more statistically valid results. The fine-tuning hyperparameters are the same as in the monolingual setup.

The results are presented in Table 3. Averaged over all tasks, some models improved the Most frequent baseline. In general, they were quite unsuccessful on BoolQ, MultiRC, and WSC but showed some promising results on COPA, RTE, and especially CB. Additional training examples in the few-shot scenario brought some visible improvements. It seems that

models perform better in the English-Slovene direction than vice versa. The best performing model is XLM-R, followed by CroSloEngual BERT and mBERT.

The low overall performance can be explained by a low number of training examples in the source language. If we take a closer look at specific models, we can observe that XLM-R shows very good results on CB in both directions. CroSloEngual BERT achieved a similar outstanding result on COPA. It is the only model that learned something on this dataset as well.

We can conclude that for the difficult SuperGLUE benchmark, the cross-lingual transfer is challenging but not impossible. In the future, we plan to expand the current set of experiments in several directions. First, we will train English models on the complete SuperGLUE datasets and transfer them to Slovene human and machine-translated datasets. Second, we will train Slovene models on the combined machine and human translated datasets and transfer them to complete English datasets. Finally, we will also combine training for several tasks and test transfer learning scenarios.

## 5.3. Multilingual results

In the multilingual setting, the three multilingual BERT models (CroSloEngual BERT, mBERT, and XLM-R) were trained on the combined full-size English and Slovene data. Slovene data is comprised of human-translated and machine-translated data. For Slovene, the models are tested on only HT data. The results are reported in Table 4.

Interestingly, all the best scores for all tasks were achieved when tested on English testing sets (the training sets were identical and comprised of both languages). This might be due to lower amounts of data used in pretraining of Slovene BERT models and the quality of the Slovene MT.

The overall best model with the highest Avg score is mBERT (best in BoolQ, MultiRC and RTE), followed by XLM-R (best in CB and WSC) and CroSloEngual BERT (best in COPA). In the MultiRC and WSC, all

Table 3: Cross-lingual results on human translated SuperGLUE test sets. The best results for zero-shot and few-shot scenarios are in bold.

Evaluation	Model	source	target	Avg	BoolQ acc.	CB F1/acc.	COPA Acc.	MultiRC F1 <sub>a</sub> /EM	RTE Acc.	WSC Acc.
Zero-shot	CroSloEngual	english	slovene	49.8	56.7	43.7/60.0	54.6	48.0/6.6	58.6	50.7
		slovene	english	52.6	60.0	53.8/70	<b>59.6</b>	<b>56.7/9.6</b>	48.3	58.2
	mBERT	english	slovene	47.4	56.7	36.2/57.2	50.2	47.3/8.7	55.2	64.4
		slovene	english	48.3	60.0	44.6/50.4	49.8	56.2/8.7	51.7	57.5
	XLM-R	english	slovene	<b>53.8</b>	<b>63.3</b>	<b>62.9/68.4</b>	53.6	48.5/0.3	<b>62.1</b>	56.2
		slovene	english	51.7	<b>63.3</b>	59.1/67.2	47.2	52.9/12.9	51.7	<b>65.8</b>
Few-shot	CroSloEngual	english	slovene	54.4	60.0	52.4/68.6	55.0	52.8/9.72	<b>65.5</b>	54.1
		slovene	english	53.0	60.0	53.8/70.0	<b>59.5</b>	56.0/12.1	49.7	58.2
	mBERT	english	slovene	50.9	60.1	53.1/66.2	50.4	50.8/9.8	53.8	64.4
		slovene	english	51.3	60.7	51.8/58.2	50.3	57.2/11.1	56.5	56.8
	XLM-R	english	slovene	<b>57.0</b>	<b>63.3</b>	<b>65.8/69.8</b>	53.3	<b>76.4/0.6</b>	62.1	57.4
		slovene	english	53.0	<b>63.3</b>	63.0/69.6	48.3	51.4/10.6	55.8	<b>65.8</b>
	Most frequent			52.4	63.3	23.0/52.7	50.0	77.3/0.3	58.6	<b>65.8</b>

Table 4: Multilingual results on English and Slovene (human translated) test sets. Models were trained on combined full size English and Slovene data. Slovene data is comprised of human translated and machine translated data. The best results for each task is in bold.

Evaluation on	Model	Avg	BoolQ Acc.	CB F1/acc.	COPA Acc	MultiRC F1 <sub>a</sub> /EM	RTE Acc.	WSC Acc.
Slovene	CroSloEngual	59.8	70.0	67.7/74.7	59.4	58.4/15.6	51.7	58.2
	mBERT	60.2	73.0	66.5/71.9	51.6	57.5/17.0	62.1	58.9
	XML-R	59.9	63.3	69.9/74.7	52.8	58.8/18.8	58.6	61.0
English	CroSloEngual	59.9	63.3	67.3/75.5	<b>62.4</b>	59.5/16.7	55.2	57.5
	mBERT	<b>64.2</b>	<b>76.7</b>	69.9/74.7	58.6	<b>60.4/21.5</b>	<b>65.5</b>	63.0
	XML-R	61.4	70.0	<b>74.1/79.9</b>	51.8	60.1/19.4	48.3	<b>65.8</b>
	Most frequent	52.4	63.3	23.0/52.7	50	77.3/0.3	58.6	65.8

models lag behind the Most Frequent baseline.

#### 5.4. Comparing human and machine translation

To test the difference in human and machine translation, we repeated the monolingual experiments separately on HT data and the same MT instances (the sizes of these HT datasets are reported in Table 1). Each model was fine-tuned using either MT or HT datasets of the same size. Only the translated content varies between both translation types; otherwise, they contain exactly the same examples. The splits of instances into train, validation and test sets is the same as in the English variant (but mostly considerably smaller, see Table 1). WSC is excluded from this evaluation as it can only be human translated, so the average score (Avg) is computed from the five remaining tasks. Table 5 shows the results. The comparison with results in Table 2 is inadequate because the models there used significantly more training examples.

Considering the Avg scores in Table 5, the monolingual SloBERTa is again the best performing Slovene model and all BERT models, regardless of translation type, perform better than the Most Frequent baselines. From the translation type perspective, the models trained on HT datasets perform better than those trained on MT datasets by 1.5 points. The only task where MT is bet-

ter than HT is MultiRC, but looking at single scores, we can observe that none of the models learned anything in this task as there is a large gap between the Most Frequent baseline and the rest of the models.

Analysis of specific tasks shows that for the BoolQ dataset, all models predict the most frequent class (the testing set might be too small for reliable conclusions in BoolQ). We can safely assume that training sample sizes are too small in BoolQ and MultiRC (none of the models learned anything) and must be increased (we have only 92 HT examples in BoolQ and 15 HT examples in MultiRC). The same is also true for RTE.

#### 5.5. Comparing machine translation systems

To check the effect of MT systems on the performance of models, we perform a separate study in the monolingual setting using two currently best MT systems for Slovene: GoogleMT<sup>6</sup> and DeepL<sup>7</sup>.

We compare all four Slovene BERT models on two datasets where all models performed better than the Most Frequent baseline, CB and COPA (WSC needs manual translation as discussed in Section 3). Results in Table 6 show surprisingly large differences between the two MT systems. DeepL gives favourable results

<sup>6</sup><https://translate.google.com/>

<sup>7</sup><https://www.deepl.com/translator>

Table 5: Comparing HT and MT on the Slovene SuperGLUE benchmark. Note that for this experiment, the average score Avg is computed as the average of the five listed tasks and not six as in Table 2. The best score for each task is in **bold**.

Task Models/Metrics	Avg	BoolQ	CB	COPA	MultiRC	RTE
		Acc.	F1/Acc.	Acc.	F1 <sub>a</sub> /EM	Acc.
Most Frequent	49.1	63.3	21.7/48.4	50.0	<b>76.4/0.6</b>	<b>58.6</b>
HT-mBERT	54.3	63.3	66.6/73.6	54.2	45.1/8.1	57.2
MT-mBERT	55.2	63.3	65.1/68.8	54.4	55.4/11.7	57.9
HT-CroSloEngual	55.6	63.3	62.1/72.4	58.2	53.0/8.4	<b>58.6</b>
MT-CroSloEngual	53.4	63.3	59.8/68.4	55.0	51.2/10.5	53.8
HT-SloBERTa	<b>57.2</b>	63.3	<b>74.0/76.8</b>	<b>61.8</b>	53.0/10.8	53.8
MT-SloBERTa	55.8	63.3	68.6/74.8	58.2	57.1/12.0	49.6
HT-XLM-R	53.5	63.3	66.2/73.2	50.0	53.3/0.9	57.2
MT-XLM-R	50.1	63.3	62.0/68.4	51.4	55.3/0.6	42.8
HT-Avg	<b>55.1</b>	63.3	<b>70.6</b>	<b>56.0</b>	29.1	<b>56.7</b>
MT-Avg	53.6	63.3	67.0	54.8	<b>31.7</b>	51.0

Table 6: Comparison of GoogleMT and DeepL MT systems on the CB and COPA datasets. The best score for each task is in bold.

Model	GoogleMT		DeepL	
	CB	COPA	CB	COPA
SloBERTa	68.6/74.8	<b>58.2</b>	<b>73.6/79.1</b>	54.2
CroSloEngual	59.8/68.4	55.0	66.1/75.1	53.8
mBERT	65.1/68.8	54.4	60.4/70.7	50.8
XML-R	62.0/68.4	51.4	44.3/63.5	53.6

on the CB dataset, while the inverse is true on the COPA dataset. We cannot explain the differences and leave a more detailed analysis of the issue for further work.

## 6. Conclusions

We prepared the Slovene translation of natural language understanding benchmark suite SuperGLUE and released it under an open-source licence (Žagar et al., 2021)<sup>8</sup>. We described the translation process and obstacles in the transfer to a morphologically rich language. The partially machine and partially human translated datasets were used in the assessment of four BERT-based models capable of handling Slovene. The results show that the monolingual SloBERTa model is currently the best performing Slovene pretrained model. However, the performance of Slovene models is still significantly worse compared to the state-of-the-art English models, showing considerable potential for improvement of NLP approaches for less-resourced languages.

Our analyses show that the models’ performance improved with human translated datasets, and in future, we intend to increase the share of human translated data. For two of the English SuperGLUE datasets, the MT was not possible. We intend to create a Slovene version of the WiC task from scratch using word sense disambiguation tasks and manually adapt the ReCoRD task to get the full SuperGLUE suite. In the WSC tasks,

Slovene verb morphology leaks some coreference information. We intend to analyze this issue and form a more challenging Slovene WSC. Next, both Slovene and Russian are Slavic languages, and therefore more similar to each other than English. It would be interesting to combine equivalent Russian and Slovene tasks from the translated SuperGLUE tasks in further multilingual and cross-lingual experiments. Finally, as we kept the format of the original SuperGLUE benchmarks, Slovene datasets can be evaluated with the original SuperGLUE leaderboard. While this allows comparison to English models, the obtained results are not publicly available. We intend to prepare a separate Slovene leaderboard and encourage the NLP community to pay attention to less-resourced languages.

## Acknowledgements

The work was supported by the Slovenian Research Agency (ARRS) core research programme P6-0411 and the Ministry of Culture of Republic of Slovenia through the project Development of Slovene in Digital Environment (RSDO). This paper is supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 825153, project EMBEDDIA (Cross-Lingual Embeddings for Less-Represented Languages in European News Media). We thank the authors of SuperGLUE benchmark for sharing test set answers for some of the tasks.

<sup>8</sup><http://hdl.handle.net/11356/1380>

## Appendix

Table 7: Examples from the development set of SuperGLUE tasks. **Bold** texts represent parts of examples' format. Texts in *italics* are part of models' input. Underlined texts are specially marked in inputs. Texts in the monospaced font represent the expected models' outputs. The WiC task cannot be translated into other languages as different meanings of a single word would translate into different words. The ReCoRD cannot be handled with MT systems as they do not preserve tagged words. Further, in morphologically rich languages such as Slovene, the correct answer would not necessarily use the same word form as the one tagged.

<b>BoolQ</b>	<p><b>Passage:</b><i>Barq's – Barq's is an American soft drink. Its brand of root beer is notable for having caffeine. Barq's, created by Edward Barq and bottled since the turn of the 20th century, is owned by the Barq family but bottled by the Coca-Cola Company. It was known as Barq's Famous Olde Tyme Root Beer until 2012.</i></p> <p><b>Question:</b>is barq's root beer a pepsi product</p> <p><b>Answer:</b>No</p>
<b>CB</b>	<p><b>Text:</b><i>B: And yet, uh, I we-, I hope to see employer based, you know, helping out. You know, child, uh, care centers at the place of employment and things like that, that will help out. A: Uh-huh. B: What do you think, do you think we are, setting a trend?</i></p> <p><b>Hypothesis:</b>they are setting a trend</p> <p><b>Entailment:</b>Unknown</p>
<b>COPA</b>	<p><b>Premise:</b><i>My body cast a shadow over the grass.</i></p> <p><b>Question:</b>What's the CAUSE for this</p> <p><b>Alternative 1:</b>The sun was rising.</p> <p><b>Alternative 2:</b>The grass was cut.</p> <p><b>Correct Alternative:</b>1</p>
<b>Multirc</b>	<p><b>Paragraph:</b><i>Susan wanted to have a birthday party. She called all of her friends. She has five friends. Her mom said that Susan can invite them all to the party. Her first friend could not go to the party because she was sick. Her second friend was going out of town. Her third friend was not so sure if her parents would let her. The fourth friend said maybe. The fifth friend could go to the party for sure. Susan was a little sad. On the day of the party, all five friends showed up. Each friend had a present for Susan. Susan was happy and sent each friend a thank you card the next week.</i></p> <p><b>Question:</b>Did Susan's sick friend recover?</p> <p><b>Candidate answers:</b> Yes, she recoverd, No (F), Yes (T), No, she didn't recover (F), Yes, she was at Susan's party (T)</p>
<b>ReCoRD</b>	<p><b>Paragraph:</b><i>(CNN) Puerto Rico on Sunday overwhelmingly voted for statehood. But Congress, the only body that can approve new states, will ultimately decide whether the status of the US commonwealth changes. Ninety-seven percent of the votes in the nonbinding referendum favored statehood, an increase over the results of a 2012 referendum, official results from the <u>State Electoral Commission</u> show. It was the fifth such vote on statehood. Today, we the people of <u>Puerto Rico</u> are sending a strong and clear message to the <u>US Congress</u> ... and to the world ... claiming our equal rights as <u>American citizens</u>, <u>Puerto Rico</u> Gov. <u>Ricardo Rossello</u> said in a news release.@highlight <u>Puerto Rico</u> voted Sunday in favor of <u>US</u> statehood</i></p> <p><b>Query:</b>For one, they can truthfully say, "Don't blame me, I didn't vote for them," when discussing the {placeholder} presidency.</p> <p><b>Correct Entities:</b> US</p>
<b>RTE</b>	<p><b>Text:</b><i>Dana Reeve, the widow of the actor Christopher Reeve, has died of lung cancer at age 44, according to the Christopher Reeve Foundation.</i></p> <p><b>Hypothesis:</b>Christopher Reeve had an accident.</p> <p><b>Entailment:</b> False</p>
<b>WiC</b>	<p><b>Context 1:</b><i>Room and board.</i></p> <p><b>Context 2:</b>He nailed <u>boards</u> across the windows.</p> <p><b>Sense match:</b> False</p>
<b>WSC</b>	<p><b>Text:</b><i>Mark told <u>Pete</u> many lies about himself, which Pete included in his book. <u>He</u> should have been more truthful.</i></p> <p><b>Coreference:</b> False</p>

## 7. Bibliographical References

- Conneau, A., Lample, G., Ranzato, M., Denoyer, L., and Jégou, H. (2018). Word translation without parallel data. In *Proceedings of International Conference on Learning Representation ICLR*.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. (2019). Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Hu, J., Ruder, S., Siddhant, A., Neubig, G., Firat, O., and Johnson, M. (2020). Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- Kasai, J., Sakaguchi, K., Bras, R. L., Dunagan, L., Morrison, J., Fabbri, A. R., Choi, Y., and Smith, N. A. (2021). Bidimensional leaderboards: Generate and evaluate language hand in hand. ArXiv preprint 2112.04139.
- Liang, Y., Duan, N., Gong, Y., Wu, N., Guo, F., Qi, W., Gong, M., Shou, L., Jiang, D., Cao, G., et al. (2020). Xglue: A new benchmark dataset for cross-lingual pre-training, understanding and generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018.
- Phang, J., Yeres, P., Swanson, J., Liu, H., Tenney, I. F., Htut, P. M., Vania, C., Wang, A., and Bowman, S. R. (2020). *jiant 2.0*: A software toolkit for research on general-purpose text understanding models. <http://jiant.info/>.
- Shavrina, T. and Malykh, V. (2021). How not to lie with a benchmark: Rearranging NLP leaderboards. ArXiv preprint 2112.01342.
- Shavrina, T., Fenogenova, A., Anton, E., Shevelev, D., Artemova, E., Malykh, V., Mikhailov, V., Tikhonova, M., Chertok, A., and Evlampiev, A. (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.
- Ulčar, M. and Robnik-Šikonja, M. (2021). SloBERTa: Slovene monolingual large pretrained masked language model. In *Proceedings of Data Mining and Data Warehousing, SiKDD*.
- Ulčar, M. and Robnik-Šikonja, M. (2020). FinEst BERT and CroSloEngual BERT: less is more in multilingual models. In *Proceedings of Text, Speech, and Dialogue, TSD 2020*, pages 104–111.
- Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. (2019a). SuperGLUE: A stickier benchmark for general-purpose language understanding systems. *Advances in Neural Information Processing Systems*, 32.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. (2019b). GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations*.

## 8. Language Resource References

- Žagar, Aleš and Robnik-Šikonja, Marko and Goli, Teja and Arhar Holdt, Špela. (2021). *Slovene translation of SuperGLUE*. Slovenian language resource repository CLARIN.SI.