

How to Probe Sentence Embeddings in Low-Resource Languages: On Structural Design Choices for Probing Task Evaluation

Steffen Eger, Johannes Daxenberger, Iryna Gurevych

Computer Science Department, Technische Universität Darmstadt, Germany

eger@aiphes.tu-darmstadt.de,

{daxenberger, gurevych}@ukp.informatik.tu-darmstadt.de

Abstract

Sentence encoders map sentences to real valued vectors for use in downstream applications. To peek into these representations—e.g., to increase interpretability of their results—probing tasks have been designed which query them for linguistic knowledge. However, designing probing tasks for lesser-resourced languages is tricky, because these often lack large-scale annotated data or (high-quality) dependency parsers as a prerequisite of probing task design in English. To investigate how to probe sentence embeddings in such cases, we investigate sensitivity of probing task results to structural design choices, conducting the first such large scale study. We show that design choices like size of the annotated probing dataset and type of classifier used for evaluation do (sometimes substantially) influence probing outcomes. We then probe embeddings in a multilingual setup with design choices that lie in a ‘stable region’, as we identify for English, and find that results on English do not transfer to other languages. Fairer and more comprehensive sentence-level probing evaluation should thus be carried out on multiple languages in the future.

1 Introduction

Sentence embeddings (a.k.a. sentence encoders) have become ubiquitous in NLP (Kiros et al., 2015; Conneau et al., 2017), extending the concept of word embeddings to the sentence level. In the context of recent efforts to open the black box of deep learning models and representations (Linzen et al., 2019), it has also become fashionable to *probe* sentence embeddings for the linguistic information signals they contain (Perone et al., 2018), as this may not be clear from their performances in downstream tasks. Such probes are linguistic micro tasks—like detecting the length of a sentence or its dependency tree depth—that have to be solved by

		classifier			
		LR	MLP	NB	RF
size	High	(A,B,C)	(A,B,C)	(C,A,B)	(C,B,A)
	Mid	(A,C,B)	(C,B,A)	(A,B,C)	(C,B,A)
	Low	(A,B,C)	(B,A,C)	(B,C,A)	(A,B,C)

Table 1: Schematic illustration of our concept of stability across two dimensions (classifier and training size). Here, three encoders, dubbed A,B,C, are ranked. The region of stability is given by those settings that support the majority ranking of encoders, which is $A \succ B \succ C$.

a classifier using given representations.

The majority of approaches for probing sentence embeddings target English, but recently some works have also addressed other languages such as Polish, Russian, or Spanish in a multi- and cross-lingual setup (Krasnowska-Kieraś and Wróblewska, 2019; Ravishankar et al., 2019). Motivations for considering a multi-lingual analysis include knowing whether findings from English transfer to other languages and determining a universal set of probing tasks that suits multiple languages, e.g., with richer morphology and freer word order.

Our work is also inspired by probing sentence encoders in multiple (particularly low-resource) languages. We are especially interested in the formal structure of probing task design in this context. Namely, when designing probing tasks for low-resource languages, some questions arise naturally that are less critical in English. One of them is the size of training data for probing tasks, as this training data typically needs to be (automatically or manually) annotated, an inherent obstacle in low-resource settings.¹

Thus, at first, we ask for the training data size

¹The main issue is that high-quality dependency parsers, as required for standard probing tasks, exist only for a handful of languages. E.g., UDPipe (Straka, 2018) is available for only about 100 languages, and performance scores for some of these are considerably below those of English (Straka, 2018).

required for obtaining reliable probing task results. This question is also relevant for English: on the one hand, [Conneau et al. \(2018\)](#) claim that training data for a probing task should be plentiful, as otherwise (highly parametrized) classifiers on top of representations may be unable to extract the relevant information signals; on the other hand, [Hewitt and Liang \(2019\)](#) note that a sufficiently powerful classifier with enough training data can in principle learn any task, without this necessarily allowing to conclude that the representations adequately store the linguistic signal under scrutiny. Second, we ask how stable probing task results are across different classifiers (e.g., MLP vs. Naive Bayes). This question is closely related to the question about size, since different classifiers have different sensitivities to data size; especially deep models are claimed to require more training data.

We evaluate the sensitivity of probing task results to the two outlined parameters—which are *mere* machine learning design choices that do not affect the linguistic content stored in the sentence representations under scrutiny—and then determine a ‘region of stability’ for English (`en`), where outcomes are predicted to be similar for the majority of parameter choices. [Table 1](#) illustrates this. Using parameter choices within our region of stability, we turn to three lower-resource languages, *viz.*: Turkish (`tr`), Russian (`ru`), and Georgian (`ka`). `tr` is a Turkic language written in Latin script which makes exhaustive use of agglutination. `ru` is a Slavic language written in Cyrillic script characterized by strong inflection and rich morphology. `ka` is a South-Caucasian language using its own script called Mkhedruli. It makes use of both agglutination as well as inflection. For these, our main research questions are whether probing task results transfer from English to the other languages.

Overall, our research questions are:

- **(i)** How reliable are probing task results across machine learning design choices?
- **(ii)** Will encoder performances correlate across languages, even though the languages and their linguistic properties may differ?
- **(iii)** Will probing task performances correlate across languages?
- **(iv)** Will the correlation between probing and downstream tasks be the same across languages?

These questions are important because they indicate whether or not probing tasks (and their relation

to downstream tasks) have to be re-evaluated in languages other than `en`.²

2 Related work

Our goal is to probe for sentence-level linguistic knowledge encoded in sentence embeddings ([Perone et al., 2018](#)) in a multilingual setup which marginalizes out the effects of probing task design choices when comparing sentence representations.

Sentence embeddings have become central for representing texts beyond the word level, e.g., in small data scenarios, where it is difficult to induce good higher-level text representations from word embeddings ([Subramanian et al., 2018](#)) or for clustering or text retrieval applications ([Reimers and Gurevych, 2019](#)). To standardize the comparison of sentence embeddings, [Conneau and Kiela \(2018\)](#) proposed the SentEval framework for evaluating the quality of sentence embeddings on a range of downstream and 10 probing tasks.

Probing tasks are used to introspect embeddings for linguistic knowledge, by taking “probes” as dedicated syntactic or semantic micro tasks ([Köhn, 2016](#)). As opposed to an evaluation in downstream applications or benchmarks like GLUE ([Wang et al., 2018](#)), probing tasks target very specific linguistic knowledge which may otherwise be confounded in downstream applications. Since they are artificial tasks, they can also be better controlled for to avoid dataset biases and artifacts. Probing is typically either executed on type/token (word) ([Tenney et al., 2019](#)) or sentence level ([Adi et al., 2017](#)). For sentence level evaluation, SentEval thus far only includes `en` data. Each probing task in SentEval is balanced and has 100k train, 10k dev, and 10k test instances. The effects of these design choices are unclear, which is why our work addresses their influence systematically.

In the **multilingual** setting, [Sahin et al. \(2019\)](#) propose 15 token and type level probing tasks. Their probing task data is sourced from UniMorph 2.0 ([Kirov et al., 2018](#)), Universal Dependency treebanks ([McCarthy et al., 2018](#)) and Wikipedia word frequency lists. To deal with lower-resourced languages, they only use 10K samples per probing task/language pair (7K/2K/1K for train/dev/test) and exclude task/language pairs for which this amount cannot be generated. Their final experi-

²Code and data are available from <https://github.com/UKPLab/conll2020-multilingual-sentence-probing>.

ments are carried out on five languages (Finnish, German, Spanish, `ru`, `tr`), for which enough training data is available. They find that for morphologically rich (agglutinative) languages, several probing tasks positively correlate with downstream applications. Our work also investigates correlation between probing and downstream performance, but we do so on sentence level.

On sentence level, [Ravishankar et al. \(2019\)](#) train an InferSent-like encoder ([Conneau et al., 2017](#)) on `en` and map this encoder to four languages (`ru`, French, German, Spanish) using parallel data. Subsequently, they probe the encoders on the probing tasks proposed by [Conneau et al. \(2018\)](#) using Wikipedia data for each language, with the same size of probing task data as in SentEval, i.e., 100k/10k/10k for train/dev/test. Their interest is in whether probing tasks results are higher/lower compared to `en` scores. They find particularly the `ru` probing scores to be low, which they speculate to be an artifact of cross-lingual word embedding induction and the language distance of `ru` to `en`. In contrast to us, their focus is particularly on the effect of transferring sentence representations from `en` to other languages. The problem of such an analysis is that results may be affected by the nature of the cross-lingual mapping techniques.

[Krasnowska-Kieraś and Wróblewska \(2019\)](#) probe sentence encoders in `en` and Polish (`pl`). They use tasks defined in [Conneau et al. \(2018\)](#) but slightly modify them (e.g., replacing dependency with constituency trees), reject some tasks (Bigram-Shift, as word order may play a minor role in `pl`), and add two new tasks (Voice and Sentence Type). Since `pl` data is less abundant, they shrink the size of the `pl` datasets to 75k/7.5k/7.5k for train/dev/test and, for consistency, do the same for `en`. They extract probing datasets from an `en-pl` parallel corpus using COMBO for dependency parsing ([Rybak and Wróblewska, 2018](#)). They find that `en` and `pl` probing results mostly agree, i.e., encoders store the same linguistic information across the two languages.

3 Approach

In the absence of ground truth, our main interest is in a ‘stable’ structural setup for probing task design—with the end goal of applying this design to multilingual probing analyses (keeping their restrictions, e.g., small data sizes, in mind). To this end, we consider a 2d space \mathcal{X} comprising probing

data size and classifier choice for probing tasks.³ For a selected set of points p_0, p_1, \dots in \mathcal{X} , we evaluate all our encoders on p_i , and determine the ‘outcomes’ O_i (e.g., ranking) of the encoders at p_i . We consider a setup p_i as stable if outcome O_i is shared by a majority of other settings p_j . This can be considered a region of agreement, similarly to inter-annotator agreement ([Artstein and Poesio, 2008](#)). In other words, we identify ‘ideal’ test conditions by minimizing the influence of parameters p_i on the outcome O_i . Below, we will approximate these intuitions using correlation.

3.1 Embeddings

We consider two types of sentence encoders, **non-parametric methods** which combine word embeddings in elementary ways, without training; and **parametric methods**, which tune parameters on top of word embeddings. As non-parametric methods, we consider: (i) average word embeddings as a popular baseline, (ii) the concatenation of average, min and max pooling (*pmeans*) ([Rücklé et al., 2018](#)); and Random LSTMs ([Conneau et al., 2017](#); [Wieting and Kiela, 2019](#)), which feed word embeddings to randomly initialized LSTMs, then apply a pooling operation across time-steps. As parametric methods, we consider: InferSent ([Conneau et al., 2017](#)), which induces a sentence representation by learning a semantic entailment relationship between two sentences; QuickThought ([Logeswaran and Lee, 2018](#)) which reframes the popular SkipThought model ([Kiros et al., 2015](#)) in a classification context; LASER ([Artetxe and Schwenk, 2019](#)) derived from massively multilingual machine translation models, and BERT base ([Devlin et al., 2019](#)), where we average token embeddings of the last layer for a sentence representation. Dimensionalities of encoders are listed in the appendix.

3.2 Probing Tasks

Following [Conneau et al. \(2018\)](#), we consider the following probing tasks: **BigramShift** (`en`, `tr`, `ru`, `ka`), **TreeDepth** (`en`), **Length** (`en`, `tr`, `ru`, `ka`), **Subject Number** (`en`, `tr`, `ru`), **WordContent** (`en`, `tr`, `ru`, `ka`), and **TopConstituents** (`en`).

³We also looked at further parameters, in particular, the class (im)balances of training datasets. Details and results can be found in the appendix. Since, however, their influence seemed to be less critical and an increased search space would blow up computational costs, we decided to limit our investigation to the described dimensions.

Task	Description	Example
Bigram Shift	Whether two words in a sentence are inverted	This is my Eve Christmas. \rightarrow True
Tree Depth	Longest path from root to leaf in constituent tree	“One hand here , one hand there , that ’s it” \rightarrow 5
Length	Number of tokens	I like cats \rightarrow 1-4 words
Subject Number	Whether the subject is in singular or plural	They work together \rightarrow Plural
Word Content	Which mid-frequency word a sentence contains	Everybody should step back \rightarrow everybody
Top Constituents	Classific. task where classes are given by 19 most common top constituent sequences in corpus	Did he buy anything from Troy \rightarrow VDP_NP_VP
Voice	Whether sent. contains a passive construct	He likes cats \rightarrow False
SV-Agree	Whether subject and verb agree	They works together \rightarrow Disagree
SV-Dist	Distance between subject and verb	The delivery was very late \rightarrow 1

Table 2: Probing tasks, their description and illustration. Upper tasks are defined as in SentEval.

We choose Length, BigramShift and WordContent because they are unsupervised tasks that require no labeled data and can thus be easily implemented across different languages—they also represent three different types of elementary probing tasks: surface, syntactic and semantic/lexical. We further include Subject Number across all our languages because number marking is extremely common across languages and it is comparatively easy to identify. We adopt **Voice** (en, tr, ru, ka) from Krasnowska-Kieraś and Wróblewska (2019). For en, we additionally evaluate on TreeDepth and TopConstituents as hard syntactic tasks. We add two tasks not present in the canon of probing tasks given in SentEval: **Subject-Verb-Agreement (SV-Agree)** (en, tr, ru, ka) and **Subject-Verb-Distance (SV-Dist)** (en, tr, ru). We probe representations for these properties because we suspect that agreement between subject and verb is a difficult task which requires inferring a relationship between pairs of words which may stand in a long-distance relationship (Gulordava et al., 2018). Moreover, we assume this task to be particularly hard in morphologically rich and word-order free languages, thus it could be a good predictor for performance in downstream tasks.

To implement the probing tasks, for en, we use the probing tasks datasets defined in Conneau and Kiela (2018) and we apply spaCy⁴ to sentences extracted from Wikipedia for the newly added probing tasks Voice and SV-Agree. For tr, ru, and ka, we do not rely on dependency parsers because of quality issues and unavailability for ka. Instead, for tr and ru, we use information from Universal Dependencies (UD) (Nivre et al., 2016). E.g., for SV-Dist, we determine the dependency distance between the main verb and the subject from UD.

⁴<https://spacy.io>

Instead of the exact distances, we predict binned classes: [1], [2,4], [5,7], [8,12], [13,∞). For ka, we use data and grammatical information from the Georgian National Corpus (GNC)⁵. We could not implement SV-Dist for ka, due to missing dependency information in the GNC. For the same reason, we omit Subject Number for ka.

For SV-Agree, we create a list of frequently occurring verbs together with their corresponding *present tense conjugations* for all involved languages including English. We check each individual candidate sentence from Wikipedia for the presence of a verb form in the list. If no word is present, we exclude the sentence from consideration. Otherwise, we randomly replace the verb form by a different conjugation in 50% of the cases.

An overview of the probing tasks, along with descriptions and examples, is given in Table 2.

3.3 Downstream Tasks

In addition to probing tasks, we test the embeddings in downstream applications. Our focus is on a diverse set of high-level sentence classification tasks. We choose **Argument Mining**, **Sentiment Analysis** and **TREC question answering**. Required training data for languages other than en has been machine translated using Google Translate⁶ for Argument Mining and TREC.⁷ Statistics for all datasets are reported in Table 6.

Argument Mining (AM) AM is an emergent NLP task requiring sophisticated reasoning capabilities. We reuse the sentence-level argument (stance)

⁵<https://clarino.uib.no/gnc>

⁶<http://translate.google.com>

⁷To estimate the quality of the machine translation, we measured its performance on parallel data. Details can be found in the appendix. While the machine translation is generally of acceptable quality, we cannot exclude the possibility that it may effect some of our downstream tasks results reported below.

detection dataset by [Stab et al. \(2018\)](#), which labels sentences extracted from web pages as pro-, con-, or non-arguments for eight different topics. A sentence only qualifies as pro or con argument when it both expresses a stance towards the topic and gives a reason for that stance. The classifier input is a concatenation of the sentence embedding and the topic encoding. In total, there are about 25,000 sentences.

Sentiment Analysis As opposed to AM, sentiment analysis only determines the opinion flavor of a statement. Since sentiment analysis is a very established NLP task, we did not machine translate `en` training data, but used original data for `en`, `ru` and `tr` and created a novel dataset for `ka`. For `en`, we use the US Airline Twitter Sentiment dataset, consisting of 14,148 tweets labeled in three sentiment classes⁸. For `tr`, we took the Turkish Twitter Sentiment Dataset with 6,172 examples and three classes⁹. For `ru`, we used the Russian Twitter Corpus (RuTweetCorp), which we reduced to 30,000 examples in two classes.¹⁰ For `ka`, we followed the approach by [Choudhary et al. \(2018\)](#) and crawled sentiment flavored tweets in a distant supervision manner. Emojis were used as distant signals to indicate sentiment on preselected tweets from the Twitter API. After post-processing, we were able to collect 11,513 Georgian tweets in three sentiment classes. The dataset will be made available publicly, including more details on the creation process.

TREC Question Type Detection Question type detection is an important part of Question-Answering systems. The Text Retrieval Conference (TREC) dataset consists of a set of questions labeled with their respective question types (six labels including e.g. “description” or “location”) and is part of the SentEval benchmark ([Conneau and Kiela, 2018](#)). We used the data as provided in SentEval, yielding 5,952 instances.

4 Experiments

Experimental Setup To the SentEval toolkit ([Conneau and Kiela, 2018](#)), which addresses both probing and downstream tasks and offers Logistic Regression (LR) and MLP classifiers on top of representations, we added implementations of

⁸<https://www.kaggle.com/crowdfLOWER/twitter-airline-sentiment>

⁹<https://github.com/hilalbenzer/turkish-sentiment-analysis>

¹⁰<http://study.mokoron.com/>

Random Forest (RF) and Naive Bayes (NB) from scikit-learn as other popular but ‘simple’ classifiers. SentEval defines specific model validation techniques for each task. Following SentEval, we tune the size of the hidden layer in $\{50, 100, 200\}$, dropout in $\{0.0, 0.1, 0.2\}$ and L^2 regularization in $\{10^{-5}, 10^{-1}\}$ when training an MLP. For RF, we tune the maximum tree depth in $\{10, 50, 100, \infty\}$. For LR, we tune the L^2 regularization in $\{10^{-5}, 10^{-1}\}$. We do not tune any hyperparameters for NB. For all probing tasks and TREC, we use predefined splits. For AM and sentiment analysis, we use 10-fold inner cross validation.

4.1 Probing task design in `en`

In our design, we consider (a) four well-known and popular classifiers—LR, MLP, NB, RF—on top of sentence representations, and (b) six different training data sizes (between 2k and 100k). We perform an exhaustive grid-search for size and classifier design, considering all combinations.

Size For each classifier, we obtain results (on 10k test instances) when varying the training data size over 2k, 5k, 10k, 20k, 30k, 100k. Downsampling was implemented by selecting the same percentage of samples that appears in the full dataset for each class. We then report average Spearman/Pearson correlations ρ/p between any two training set sizes s and t over all 9 probing tasks:¹¹

$$\text{sim}_c(s, t) = \frac{1}{n} \sum_{i=1}^n \rho^*(\mathbf{c}_i(s), \mathbf{c}_i(t)) \quad (1)$$

where n is the number of probing tasks ($n = 9$ for `en`), and $\mathbf{c}_i(s)$ is the vector that holds scores for each of the 7 sentence encoders in our experiments, given training size s , for probing task i and classifier c . We set correlations to zero if the p-value > 0.2 .¹² In Table 3, we then report the minimum and average scores $\min_{(s,t)} \text{sim}_c(s, t)$ and $\frac{1}{M} \sum_{(s,t)} \text{sim}_c(s, t)$, respectively, per classifier c . We observe that the minimum values are small to moderate correlations between 0.2 (for NB) and 0.6 (for RF). The average correlations are moderate to high correlations ranging from 0.6 (for NB) to above 0.8 (for the others).

¹¹We report both Spearman and Pearson for some of the results but give only Pearson for the remainder, reporting Spearman in the appendix.

¹²We choose a high p-value, because we correlate small vectors of size 7.

In Figure 1 (left), we show all the values $\text{sim}_c(s, t)$ for $c = \text{LR}, \text{NB}$. We observe that, indeed, LR has high correlations between training sizes especially starting from 10k training data points. The corresponding correlations of NB are much lower comparatively.

In Figure 2, we plot the stability of each training data size s

$$\text{sim}_c(s) = \frac{1}{N} \sum_t \text{sim}_c(s, t) \quad (2)$$

for all of our classifiers c and where N is a normalizer equal to the number of different training set sizes, $N = 6$ in our case. The higher this score for a training size s , the more similar are the probing results for another training size t , on average. Across all classifiers, 2k and 100k are least stable—100k is the default setting of SentEval. Most stable are 10k and 20k.

Classifier	Min		Avg	
	ρ	p	ρ	p
MLP	.480	.420	.810	.843
LR	.524	.502	.808	.805
RF	.529	.623	.800	.853
NB	.174	.292	.626	.671

Table 3: Stability over training sizes, in terms of minimum and average Spearman (ρ) / Pearson (p) correlation between any two sizes.

Classifier Next, we add the classifier choice as a second dimension: we examine whether correlations (Spearman/Pearson) between vectors \mathbf{c} (holding scores for each of 7 sentence encoders for a classifier c) and \mathbf{d} (holding the same scores for a classifier d) are similar in the same sense as in Eq. (1):

$$\text{sim}_{c,d}(s, t) = \frac{1}{n} \sum_{i=1}^n \rho^*(\mathbf{c}_i(s), \mathbf{d}_i(t)) \quad (3)$$

Again, we average across all probing tasks, and set correlation values to zero if the p-value exceeds 0.2. In Table 4, we give min/avg values across data set sizes in this setup. We observe that LR and MLP most strongly agree. They have acceptable average agreement with RF, but low agreement with NB, on average, and, in the worst cases, even negative correlations with NB.

In Figure 1 (right), we illustrate correlations between LR and NB, on the one side, and LR and RF, on the other side, across all possible training set

	LR	RF	NB
MLP	.481/.790	.492/.632	-.043/.236
LR		.406/.640	-.057/.197
RF			.029/.320

Table 4: Min/Avg values $\text{sim}_{c,d}(s, t)$ across (s, t) (using Pearson) between classifiers c and d .

sizes. We observe that as the training data set sizes for RF and LR become larger, these two classifiers agree more strongly with LR. RF starts to have acceptable agreement with LR from 10k training instances onwards, while NB has acceptable agreement with LR only in the case of 100k training instances.

We now operationalize our intuition of ‘region of stability’ outlined in Table 1. For each of nine probing tasks, we compute the following. Let $r_j = E_{\zeta(1)} \succ E_{\zeta(2)} \succ E_{\zeta(3)} \succ \dots$ be a specific ranking of encoders, where ζ is a fixed permutation. Let $r_{(c,s)}$ be the ranking of encoders according to the classifier, size combination (c, s) . We compute the Spearman correlation $\tau_{(c,s,j)}$ between $r_{(c,s)}$ and r_j . For each possible ranking r_j of our 7 encoders, we then determine its support as the average over all values $\tau_{(c,s,j)}$ and then find the ranking r_{\max} with most support according to this definition. Finally, we assign a score to the combination (c, s) not only when $r_{(c,s)}$ equals r_{\max} , but also when $r_{(c,s)}$ is close to r_{\max} : we again use the Spearman correlation between $r_{(c,s)}$ and r_{\max} as a measure of closeness (we require a closeness of at least 0.75). The final score for (c, s) is given by (ignoring the threshold of 0.75 in the equation):

$$\mu_{(c,s)} = \sum_{i=1}^n \rho^*(r_{(c,s)}^{(i)}, r_{\max}^{(i)}) \quad (4)$$

Table 5 shows classifier, size combinations with highest μ scores. LR and MLP are at the top, along with RF in the setting of 100k training data size. LR with size 10k is most stable overall, but the distance to the other top settings is small. Least stable (not shown) is NB.

classifier size	LR 10k	LR 20k	RF 100k	MLP 20k	MLP 30k	MLP 10k
$\mu_{(c,s)}$	7.6	7.3	7.2	7.0	7.0	6.9

Table 5: Most stable classifier, size combinations according to Eq. (4).

Overall, we answer our first research question—




Figure 1: Left two: Average correlations $\text{sim}_c(s, t)$ for LR (first) and NB (second), using Pearson, for any two sizes $s, t \in \{2k, 5k, 10k, 20k, 30k, 100k\}$. Right two: Average correlations $\text{sim}_{c,d}(s, t)$ for $c = \text{LR}$ and $d = \text{RF}$ (third) and $c = \text{LR}$ and $d = \text{NB}$ (fourth). Best viewed in color.




Figure 2: Stability of each training size computed using Eq. (2) for different classifiers c .

(i) **How reliable are probing task results across machine learning design choices?**—as follows (for en): probing tasks results can be little reliable and may vary greatly among machine learning parameter choices. The standard training set size of SentEval, 100k, appears to be less stable. As region of stability, we postulate especially the setting with 10k training instances for the LR classifier.

4.2 Multi-lingual results

Experimental Setup Given our results for en , we choose the LR classifier with a size of roughly 10k instances overall. Table 6 provides more details about the datasets. In line with SentEval (and partly supported by our results on dataset balance given in the appendix), we aim for as balanced label distributions as possible. Because of the small test sizes, we use inner 5-fold cross validation for all tasks except for SubjNumber, where we use pre-defined train/dev/test splits as in [Conneau et al. \(2018\)](#) to avoid leaking lexical information from train to test splits.

We obtain average and pmeans embeddings through pooling over pre-trained FastText embeddings ([Grave et al., 2018](#)). The same embeddings

are used for the random LSTM. For average BERT, we use the base-multilingual-cased model. We machine translate the AINLI corpus into tr , ru and ka , to obtain training data for Infsent.¹³ The models are then trained using default hyperparameters and with pre-trained FastText embeddings. Compared to en , we modify the WC probing task in the multilingual setting to only predict 30 mid-frequency words instead of 1000. This is more appropriate for our much smaller data sizes.

4.2.1 Probing tasks

Results are shown in Figures 3 and 4.




Figure 3: Pearson correlations across languages for different encoders.

(ii) **Will encoder performances correlate across languages?** For each encoder e , we correlate performances of e between en and the other languages on 5 (for ka) and 7 (for tr , ru) probing tasks (using 10k dataset size and LR for all involved languages, including en). In Figure 3, we observe that correlations between en and other languages are generally either zero or weakly positive. Only average embeddings have more than 1 positive correlation scores across the 3 language combinations with en . Among low-resource languages, there

¹³Using Google Translate, see appendix for details.

Task	EN		TR		RU		KA	
	Size	Balance	Size	Balance	Size	Balance	Size	Balance
Bigram Shift	100k	1:1	10k	1:1	10k	1.1:1	10k	1.1:1
Length	100k	1:1	10k	1:1	12k	1:1	10k	1:1
Subject Number	100k	1:1	4,093	5:1	11,877	1:1	-	-
Word Content	100k	1:1	10k	1.5:1	10k	1.2:1	10k	4:1
Top Constituents	100k	1:1	-	-	-	-	-	-
Tree Depth	100k	2.2:1	-	-	-	-	-	-
Voice	100k	1:1	8,417	6:1	10k	2:1	10k	1.9:1
SV-Agree	100k	1:1	10k	1:1	10k	1:1	10k	1:1
SV-Dist	100k	1:1	1,642	1.9:1	8,231	1.1:1	-	-
Arg. Mining (macro-F1)	25,303	3:1	25,303	3:1	25,303	3:1	25,303	3:1
TREC (Accuracy)	5,952	14:1	5,952	14:1	5,952	14:1	5,952	14:1
Sentiment Analysis (macro-F1)	14,148	4.2:1	6,172	1.7:1	30k	1:1	11,513	5.5:1

Table 6: Probing and downstream tasks. We report the balance between the class with the most and the least samples. For downstream tasks, the evaluation measure is given in brackets.

are no negative correlations and fewer zero correlations. All of the low-resource languages correlate more among themselves than with `en`. This makes sense from a linguistic point of view, since `en` is clearly the outlier in our sample given its minimal inflection and fixed word order. Thus, the answer to this research question is that our results support the view that transfer is better for typologically similar languages.

(iii) Will probing task performances correlate across languages? For each probing task π , we report Pearson correlations, between all language pairs, of vectors holding scores of 7 encoders on π . Figure 4 shows the results. The pattern is overall




Figure 4: Pearson correlations across languages for different probing tasks.

similar as for (ii) in that there are many zero correlations between `en` and the other languages. `tr` also negatively correlates with `en` for SV-Agree. Only BigramShift has positive correlations throughout. Low-resource languages correlate better among themselves as with `en`. Our conclusions are the same as for question (ii).

Note that our findings contrast with [Krasnowska-](#)

[Kieraś and Wróblewska \(2019\)](#), who report that probing results for `en` and `pl` are mostly the same. Our results are more intuitively plausible: e.g., a good encoder should store linguistic information relevant for a particular language.

4.2.2 Downstream Tasks

Results are shown in Figure 5.

(iv) Will the correlation between probing and downstream tasks be the same across languages? For each of our languages, we correlate probing and downstream task performances. The results show that the answer to research question (iv) is clearly negative. In particular, `en` behaves differently to the other languages—while `ru` and `tr` behave more similarly. `ka` is the only language with negative correlations for Length, `en` the only one with positive scores. For the sentiment task, Word Content correlates positively for all languages except `ka`. The AM task correlates only in `en` and `ka`, but with different probing tasks. SV-Agree correlates positively with TREC and sentiment in all languages but `en`. This might be because determining the agreement of subject and verb is more grammatically complex in the other languages compared to English, and storing an adequate amount of grammatical information may be beneficial for certain downstream tasks. Predicting the performances of embeddings in downstream tasks via probing tasks thus appears idiosyncratic for individual languages. Opposed to [Sahin et al. \(2019\)](#), who suggest a direct relation between word level probing and downstream performance on agglutinative languages, we see little to no systematic correlation on the sentence level. Overall, SV-Agree is the best predictor across languages, with




Figure 5: Pearson correlation among probing task and downstream performance for all languages.

7 positive correlations out of 12 possible. Interestingly, this task is missing from the current canon of SentEval.

5 Concluding Remarks

We investigated formal aspects of probing task design, including probing data size and classifier choice, in order to determine structural conditions for multilingual (low-resource) probing. We showed that probing tasks results are at best partly stable even for `en` and that the rankings of encoders varies with design choices. However, we identified a partial region of stability where results are supported by a majority of settings—even though this may not be mistaken for a ‘region of truth’. This region was identified in `en`, which has most resources available. Our further findings then showed that probing and downstream results do not transfer well from English to our other languages, which in turn challenges our identified region of stability.

Overall, our results have partly negative implications for current practices of probing task design as they indicate that probing tasks are to some degree unreliable tools for introspecting linguistic information contained in sentence encoders. The relation of probing to downstream tasks is also unclear, as our multilingual results show. This is supported by recent findings giving contradictory claims regarding, e.g., the importance of the Word Content probing task for downstream performances (Eger et al., 2019; Wang and Kuo, 2020; Perone et al., 2018). Our findings further add to contemporaneous work by Ravichander et al. (2020) and Elazar et al. (2020), who showed that probes do not necessarily identify linguistic properties required for solving an actual task, thus questioning a common interpretation of probing itself.

An important aspect to keep in mind for correlation analyses as we conducted is that results may heavily depend on the selection of encoders involved—in our case, we selected a number of recently proposed state-of-the-art models in con-

junction with weaker baseline models, for a diverse collection of encoders. While the small number of encoders we examined is a clear limitation of our approach, many of our results are significant (at relatively large p-values).

To the degree that the supervised probing tasks examined here will remain important tools for interpretation of sentence encoders in the future, our results indicate that multilingual probing is important for a fairer and more comprehensive comparison of encoders.

Acknowledgments

We thank the anonymous reviewers for their useful comments and suggestions. Wei Zhao and Benjamin Schiller gave useful feedback on earlier versions of this paper. Daniel Wehner and Martin Kerscher greatly helped to extend the SentEval code and collect the probing data for the low-resource languages. We finally thank Paul Meurer from the University of Bergen for making the Georgian National Corpus (GNC) available for our experiments in `ka`. The first author has been funded by the HMWK (Hessisches Ministerium für Wissenschaft und Kunst) as part of structural location promotion for TU Darmstadt in the context of the Hessian excellence cluster initiative “Content Analytics for the Social Good” (CA-SG). The second author has been supported by the German Federal Ministry of Education and Research (BMBF) under the promotional reference 03VP02540 (ArgumenText).

References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. *Fine-grained analysis of sentence embeddings using auxiliary prediction tasks*. *CoRR*, abs/1608.04207.
- Mikel Artetxe and Holger Schwenk. 2019. *Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond*. *TACL*, 7:597–610.

- Ron Artstein and Massimo Poesio. 2008. [Inter-coder agreement for computational linguistics](#). *Comput. Linguist.*, 34(4):555–596.
- Narendra Choudhary, Rajat Singh, Vijjini Anvesh Rao, and Manish Shrivastava. 2018. [Twitter corpus of resource-scarce languages for sentiment analysis and multilingual emoji prediction](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1570–1577, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Alexis Conneau and Douwe Kiela. 2018. [SentEval: An evaluation toolkit for universal sentence representations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Alexis Conneau, Germán Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. [What you can cram into a single vector: Probing sentence embeddings for linguistic properties](#). *CoRR*, abs/1805.01070.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Steffen Eger, Andreas Rücklé, and Iryna Gurevych. 2019. [Pitfalls in the evaluation of sentence embeddings](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 55–60, Florence, Italy. Association for Computational Linguistics.
- Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. 2020. [When bert forgets how to pos: Amnesic probing of linguistic properties and mlm predictions](#). *ArXiv*, abs/2006.00995.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. [Learning word vectors for 157 languages](#). In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. [Colorless green recurrent networks dream hierarchically](#). *CoRR*, abs/1803.11138.
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2015. [Skip-thought vectors](#). *CoRR*, abs/1506.06726.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqi, Sabrina J. Mielke, Arya McCarthy, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 2018. [UniMorph 2.0: Universal morphology](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Arne Köhn. 2016. [Evaluating embeddings using syntax-based classification tasks as a proxy for parser performance](#). In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 67–71, Berlin, Germany. Association for Computational Linguistics.
- Katarzyna Krasnowska-Kieraś and Alina Wróblewska. 2019. [Empirical linguistic study of sentence embeddings](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5729–5739, Florence, Italy. Association for Computational Linguistics.
- Tal Linzen, Grzegorz Chrupała, Yonatan Belinkov, and Dieuwke Hupkes, editors. 2019. *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. Association for Computational Linguistics, Florence, Italy.
- Lajanugen Logeswaran and Honglak Lee. 2018. [An efficient framework for learning sentence representations](#). *CoRR*, abs/1803.02893.
- Arya D. McCarthy, Miikka Silfverberg, Ryan Cotterell, Mans Hulden, and David Yarowsky. 2018. [Marrying universal dependencies and universal morphology](#). In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, pages 91–101, Brussels, Belgium. Association for Computational Linguistics.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. [Universal dependencies v1: A multilingual treebank collection](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1659–1666, Portorož, Slovenia. European Language Resources Association (ELRA).

- Christian Perone, Roberto Silveira, and Thomas Paula. 2018. [Evaluation of sentence embeddings in downstream and linguistic probing tasks](#). *CoRR*, abs/1806.06259.
- Abhilasha Ravichander, Yonatan Belinkov, and Eduard H. Hovy. 2020. [Probing the probing paradigm: Does probing accuracy entail task relevance?](#) *ArXiv*, abs/2005.00719.
- Vinit Ravishankar, Lilja Øvrelid, and Erik Velldal. 2019. [Probing multilingual sentence representations with x-probe](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (ReL4NLP-2019)*, pages 156–168, Florence, Italy. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Andreas Rücklé, Steffen Eger, Maxime Peyrard, and Iryna Gurevych. 2018. [Concatenated power mean word embeddings as universal cross-lingual sentence representations](#). *CoRR*, abs/1803.01400.
- Piotr Rybak and Alina Wróblewska. 2018. [Semi-supervised neural system for tagging, parsing and lematization](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 45–54, Brussels, Belgium. Association for Computational Linguistics.
- Gözde Gül Sahin, Clara Vania, Ilia Kuznetsov, and Iryna Gurevych. 2019. [LINSPECTOR: multilingual probing tasks for word representations](#). *CoRR*, abs/1903.09442.
- Christian Stab, Tristan Miller, Benjamin Schiller, Pranav Rai, and Iryna Gurevych. 2018. [Cross-topic argument mining from heterogeneous sources](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3664–3674, Brussels, Belgium. Association for Computational Linguistics.
- Milan Straka. 2018. [UDPipe 2.0 prototype at CoNLL 2018 UD shared task](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 197–207, Brussels, Belgium. Association for Computational Linguistics.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. [Learning general purpose distributed sentence representations via large scale multi-task learning](#). In *International Conference on Learning Representations*.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Bin Wang and C. C. Jay Kuo. 2020. [Sbert-wk: A sentence embedding method by dissecting bert-based word models](#).
- John Wieting and Douwe Kiela. 2019. [No training required: Exploring random encoders for sentence classification](#). *CoRR*, abs/1901.10444.