

RuSentEval: Linguistic Source, Encoder Force!

Vladislav Mikhailov^{1,2}, Ekaterina Taktasheva², Elina Sigdel², Ekaterina Artemova²

¹ SberDevices, Sberbank, Moscow, Russia

² HSE University, Moscow, Russia

Mikhaylov.V.Nikola@sberbank.ru

{etaktasheva, essigdel, elartemova}@hse.ru

Abstract

The success of pre-trained transformer language models has brought a great deal of interest on how these models work, and what they learn about language. However, prior research in the field is mainly devoted to English, and little is known regarding other languages. To this end, we introduce RuSentEval, an enhanced set of 14 probing tasks for Russian, including ones that have not been explored yet. We apply a combination of complementary probing methods to explore the distribution of various linguistic properties in five multilingual transformers for two typologically contrasting languages – Russian and English. Our results provide intriguing findings that contradict the common understanding of how linguistic knowledge is represented, and demonstrate that some properties are learned in a similar manner despite the language differences.

1 Introduction

Transformer language models (Vaswani et al., 2017) have achieved state-of-the-art results on a wide range of NLP tasks in multiple languages, demonstrated strong performance in zero-shot cross-lingual transfer (Pires et al., 2019), and even surpassed human solvers in NLU benchmarks such as SuperGLUE (Wang et al., 2019). The success has stimulated research in how these models work, and what they acquire about language. The majority of the introspection techniques are based on the concept of *probing tasks* (Adi et al., 2016; Shi et al., 2016; Conneau et al., 2018) which allow analyzing what linguistic properties are encoded in the intermediate representations. A rich variety of tasks has been introduced so far, ranging from token-level and sub-sentence probing (Liu et al., 2019; Tenney et al., 2019) to sentence-level probing (Alt et al., 2020). A prominent method

to explore the inner workings of the models involves training a lightweight classifier to solve a probing task over features produced by them, and assess their knowledge by the classifier’s performance. Recently, the methods have been greatly extended to latent subclass learning (Michael et al., 2020), correlation similarity measures (Wu et al., 2020), information-theoretic probing (Voita and Titov, 2020), investigation of individual neurons (Durrani et al., 2020; Suau et al., 2020), and many more.

Despite growing interest in the field, English remains the focal point of prior research (Belinkov and Glass, 2019; Rogers et al., 2021) leaving other languages understudied. To this end, several monolingual and cross-lingual probing suites have been assembled (see Section 2), with a few of them following SentEval toolkit (Conneau et al., 2018; Conneau and Kiela, 2018). However, most of them directly apply an English-oriented method which is not guaranteed to be universal across languages, or use the Universal Dependencies (UD) Treebanks (Nivre et al., 2016) which tend to be inconsistent (Alzetta et al., 2017; de Marneffe et al., 2017; Droganova et al., 2018).

This work proposes **RuSentEval**, a probing suite for evaluation of sentence embeddings for the Russian language. We adapted the method for English (Conneau et al., 2018) to complement the peculiarities of Russian. In contrast to closely related datasets (Ravishankar et al., 2019; Eger et al., 2020), RuSentEval is fully guided by linguistic expertise, relies on annotations obtained with the current state-of-the-art model for Russian morphosyntactic analysis (Anastasyev, 2020), and includes tasks that have not been explored yet.

The contributions are summarized as three-fold. First, we present an enhanced set of 14 probing tasks for Russian, organized by the type of linguistic properties. Second, we carry out a series of

probing experiments on two typologically different languages: English, which is an analytic Germanic language, and Russian, which is a fusional Slavic one. Keeping in mind that the English (Conneau et al., 2018) and Russian datasets are combined from different annotation schemas, we introspect five multilingual transformer-based encoders, including their distilled versions. We apply several probing methods to conduct the analysis from different perspectives, and support the findings with statistical significance. Besides, we establish count-based and neural-based baselines for the tasks. Finally, RuSentEval is publicly available¹, and we hope it will be used for evaluation and interpretation of language models and sentence embeddings for Russian.

2 Related Work

Introspection of pre-trained language models for languages other than English, specifically Russian, is usually conducted in the cross-lingual setting. The primary goal of such experiments is to explore how particular linguistic properties are distributed in a given collection of multilingual embedding and language models. LINSPECTOR (Şahin et al., 2020) is one of the first probing suites that covers a wide range of linguistic phenomena in 24 languages. The benchmark uses UniMorph 2.0 (Kirov et al., 2018) to design the tasks since UD do not provide a sufficient amount of data for the considered languages. Despite this, UD Treebanks have become the main source for collection of multilingual probing tasks, limiting the scope to morphology and syntax (Krasnowska-Kieraś and Wróblewska, 2019). Other sources for the assembly include multilingual datasets created by means of machine translation, such as XNLI (Conneau et al., 2020b), or datasets that labelled with similar annotation schemes for named entity recognition (NER) and semantic role labelling (SRL) tasks (Şahin et al., 2020).

A few prior works (Ravishankar et al., 2019; Eger et al., 2020) that follow SentEval toolkit (Conneau and Kiela, 2018) directly apply the method designed for English to multiple typologically diverse languages, which raises doubts if such strategy is universal across languages that exhibit unique peculiarities, e.g. a free word order and rich inflectional morphology. This can lead to low quality of the

¹<https://github.com/RussianNLP/rusenteval>

datasets and unreliable experimental results, particularly for Russian. Consider a few examples for (BShift) task in Russian: *Fedra zatem povestilas', a Tesej uznal pravdu. "Phaedra later hanged herself, and Theseus unraveled the truth."* (Ravishankar et al., 2019), and *Shestoe – zanimatsya nado svoim obrazovaniem* "The sixth point is that you **to take care need** of your education" (Eger et al., 2020). The sentences are labelled as positive, meaning that they exhibit incorrect word order. While the word order changes can lead to the syntax perturbations for English, both sentences are still acceptable in terms of syntax for Russian. Moreover, the dataset sizes tend to be inconsistent across languages due to using UD Treebanks which makes it difficult to compare the results (Eger et al., 2020).

Another line of research includes probing machine translation models (Mareček et al., 2020) over multiple languages, and probing for cross-lingual similarity by utilizing paired sentences in mutually intelligible languages (Choenni and Shutova, 2020a,b). Last but not least, such benchmarks as XGLUE (Liang et al., 2020) and XTREME (Hu et al., 2020) allow evaluating the current state of language transferring methods.

3 Probing Tasks

Data The sentences for our probing tasks were extracted from the following publicly available resources: Russian Wikipedia articles² and news corpora such as Lenta.ru³ and the news segment of Taiga corpus (Shavrina and Shapovalova, 2017). We used rusenttokenize library⁴, a rule-based sentence segmenter for Russian, to split texts into sentences. Each sentence was tokenized with spaCy Russian Tokenizer⁵. The sentences were filtered by the 5-to-25 token range and annotated with the current state-of-the-art model for joint morphosyntactic analysis in Russian (Anastasyev, 2020). Besides, we performed two additional preprocessing steps. (1) We computed the IPM frequency of each sentence using the New Frequency Vocabulary of Russian Words (Lyashevskaya and Sharov, 2009) to control the word frequency distribution. The

²<https://dumps.wikimedia.org/ruwiki/latest/>

³<https://github.com/yutkin/Lenta.Ru-News-Dataset>

⁴<https://pypi.org/project/rusenttokenize>

⁵https://github.com/aatimofeev/spacy_russian_tokenizer

IPM values of each token lemma in the sentence (if present in the vocabulary) were averaged over a total number of token lemmas. The sentences with the IPM frequency of lower than 0.9 were filtered out. This allows discarding the majority of sentences that contain rare words, acronyms, abbreviations, or loanwords. (2) Syncretism is peculiar to fusional languages, that is, when a word can belong to multiple part-of-speech tags, or express multiple ambiguous morphosyntactic features (Baerman, 2007). Following (Şahin et al., 2020), we removed sentences in the semantic tasks (described below) where the target word has multiple part-of-speech tags. This step allows simplifying the probe interpretation and ensuring a fairer experimental setup in terms of the language comparison.

The total number of the annotated sentences after filtering is 3.6 mln, and they are publicly available. Each task consists of a 100k-sentence training set and 10k-sentence validation and test sets. There is no sentence overlap across the splits, and all sets are balanced by the number of instances per target class.

Surface properties tasks test if it is possible to recover information about surface properties from the contextualized representations. (**SentLen**) is a 6-way classification task aimed to predict a number of tokens given a sentence representation. Similar to (Adi et al., 2016; Conneau et al., 2018), we grouped sentences into 6 equal-width bins by length. The *word content* (**WC**) task tests if the information on the original words in a sentence can be inferred from its representation. We selected 1k lemmas from the source corpus vocabulary within the 1.5k-3k rank range when sorted by frequency, and sampled equal numbers of sentences that contain only one of these lemmas. The task is treated as a 1k-way classification that requires knowledge about lexical items and their inflectional paradigms.

Syntactic properties is a group of tasks that probe the encoder representations for syntactic properties. In the (**ConjType**) task, the sentences must be classified in terms of the type of connection between complex clauses. The objective of the classifier is to tell whether a sentence involves coordination or subordination.

(**ImpersonalSent**) is a binary classification task that aims to define if there is a lack of a grammatical subject in the main clause of a sentence. It is usually expressed by a singular third-person, reflexive, singular neuter, or invariable predicate form

(*smerkaetsya* “it is getting dark”; *zharko* “it is hot”; *pora idti* “it is time to go”), an adverbial predicate phrase (*bylo sovershenno tikho* “it was absolutely quiet”), and intransitive verbs typically combined with a noun phrase in the instrumental (*zapahlo rozami* “it smells roses”).

(**TreeDepth**) task tests whether the encoder representations store the information on the hierarchical and syntactic structure of sentences. Specifically, the goal is to probe for knowledge about the depth of the longest path from the root node to any leaf in the syntax tree. Similar to (Conneau et al., 2018), we obtained sentences where the tree depth and the sentence length are de-correlated. The tree depth values range from 5 to 9 which makes (**TreeDepth**) a 5-way classification task.

(**Gapping**) task deals with the detection of syntactic gapping that occurs in coordinated structures and elides a repeated predicate, typically from the second clause. We used data provided in the Shared Task on Automatic Gapping Resolution for Russian, or AGRR-2019 (Ponomareva et al., 2019). For instance, the sentence *Odin imel silu solntsa, drugoy – luny*. “One had the power of the Sun, the other (*had the power of*) the Moon.” comprises an omission of a repeating predicate in the non-initial clause with its semantics remaining expressed.

The *N-gram shift* task (**NShift**) is analogous to SentEval’s (**Bshift**) that tests the encoder’s sensitivity to incorrect word order. As opposed to English, only specific cases of word inversion in Russian lead to syntax perturbation. We, therefore, perturbed N-grams that correspond to a set of predefined morphosyntactic patterns. We used TF-IDF method from scikit-learn library (Pedregosa et al., 2011) to build an N-gram feature matrix that was further applied for the word order perturbation. For instance, we reversed adjacent words in prepositional phrases, numeral phrases, compound noun phrases, etc. Below is an example where the head of the prepositional phrase *v shkolu* ‘to school’ is inverted with the dependent noun:

Segodnya on ne poshel **shkolu v**.

“He did not go **school to** today.”

Semantic properties tasks rely on both syntactic and semantic structure of a sentence to recover a higher-level property. (**SubjNumber**) and (**SubjGender**) probe for the number and gender features of the subject in the main clause. Similarly, (**ObjNumber**) and (**ObjGender**) focus on

the number and gender of direct object in the main clause. In the following tasks, the aim is to probe for the morphosyntactic features of the predicate or the head of a predicative construction of the main clause: *predicate voice (PV)*, *predicate aspect (PA)*, and *predicate tense (PT)*. The latter is analogous to SentEval’s (**Tense**) task.

The semantic tasks test if the contextualized representations not only capture the morphosyntactic features but also encode higher-level, structural and syntactic-semantic information (namely, the syntax tree hierarchy and the actant structure of a predicate). Note that the boundary between the surface, syntactic and semantic tasks is relatively blurred.

4 Experimental Setup

4.1 Encoders

We run the experiments on the following 12-layer multilingual transformer encoders released by HuggingFace (Wolf et al., 2019):

M-BERT (Devlin et al., 2019) is trained on masked language modeling (MLM) and next sentence prediction tasks, over concatenated monolingual Wikipedia corpora in 104 languages.

XLM-R (Conneau et al., 2020a) is trained on ‘dynamic’ MLM task, over filtered CommonCrawl data in 100 languages (Wenzek et al., 2020).

MiniLM (Wang et al., 2020) is a distilled transformer of BERT architecture, but uses the XLM-RoBERTa tokenizer.

LABSE (Feng et al., 2020) employs a dual-encoder architecture that combines MLM and translation language modeling (Conneau and Lample, 2019).

M-BART (Liu et al., 2020) is a sequence-to-sequence transformer model with a BERT encoder, and an autoregressive GPT-2 decoder (Radford et al., 2019). We used only the encoder in the experiments.

4.2 Methods

Probing Classifiers We trained linear and non-linear classifiers over intermediate representations produced by the encoders⁶, using categorical cross-entropy loss, and Adam optimizer (Kingma and Ba, 2015). For the non-linear classifier (MLP), we used the Sigmoid activation function, the number of hidden states of 250, and the dropout rate of

⁶We used mean-pooled sentence representations to train the classifiers.

0.2. Training is run over 5 iterations with the L2-regularization parameter $\in [0.1, \dots, 1e^{-5}]$ tuned on the validation set, and the best classifier selected. The performance is evaluated by accuracy score.

Individual Neuron Analysis Neuron-level introspection technique (Durrani et al., 2020) allows identifying *top neurons* that contribute most to a probing task and observe how these neurons are distributed across layers of the encoder. Similarly, we trained a linear probing classifier over concatenated sentence representations and used the weights to measure the importance of each neuron. The classifier is trained using Elastic-net regularization with L1 and L2 λ ’s $\in [0.1, \dots, 1e^{-5}]$ tuned on the validation set. Refer to (Dalvi et al., 2019; Durrani et al., 2020) for more details.

Correlation Analysis Correlation-based analysis techniques posed in (Wu et al., 2020) allow measuring similarity of the encoder intermediate representations without any linguistic annotation. We apply neuron-level (`maxcorr`) and representation-level (`lincka`) similarity measures to investigate the encoders. `maxcorr` identifies pairs of neurons of the maximum correlation from two different layers. It is high when two layers have pairs of neurons with similar behavior. `lincka` provides a comparison of representations from different layers in a given collection of models. Two layers are assigned a high similarity if their representations behave similarly over all the neurons.

4.3 Baselines

We established a number of count-based and non-contextualized baseline features to train the probing classifiers as outlined in Section 4.2. We used N-gram range $\in [1, 4]$ and top-150k features in the vocabularies for each count-based baseline. The count-based features include **TF-IDF over character N-grams**, **TF-IDF over BPE tokens** (Sennrich et al., 2016), and **TF-IDF over SentencePiece tokens** (Kudo and Richardson, 2018). We applied multilingual BertTokenizer and XLMRobertaTokenizer by HuggingFace to segment sentences into BPE and SentencePiece tokens. The non-contextualized features are mean-pooled monolingual **fastText** sentence embeddings (Bojanowski et al., 2017). We used monolingual fastText models for English⁷ and Russian. The latter was trained

⁷<https://fasttext.cc/docs/en/crawl-vectors.html>

Probing Task	Language	M-BERT	LABSE	XLM-R	MiniLM	M-BART
Nshift	Ru	84.8 [8]	82.6 [5]	86.9 [9]	80.5 [9]	78.6 [12]
	En	81.8 [10]	84.4 [5]	85.7 [10]	79.3 [8]	83.8 [12]
ObjNumber	Ru	82.8 [6]	82.5 [2]	83.7 [10]	77.8 [10]	81.5 [7]
	En	86.2 [6]	85.4 [3]	86.0 [8]	85.2 [6]	85.9 [9]
SentLen	Ru	91.3 [2]	93.3 [1]	94.5 [2]	94.1 [2]	96.2 [4]
	En	96.3 [2]	96.6 [1]	95.8 [2]	96.1 [3]	97.3 [3]
SubjNumber	Ru	90.5 [7]	92.9 [3]	94.9 [11]	94.2 [12]	93.1 [10]
	En	87.8 [7]	90.7 [12]	86.9 [10]	85.6 [6]	87.3 [9]
Tense	Ru	99.5 [8]	99.8 [5]	99.8 [5]	98.2 [7]	99.6 [7]
	En	88.9 [8]	88.8 [6]	88.8 [9]	87.3 [5]	89.1 [9]
TreeDepth	Ru	44.7 [6]	46.1 [4]	46.5 [5]	44.8 [7]	45.8 [11]
	En	41.2 [5]	42.7 [5]	41.8 [7]	40.9 [7]	41.2 [12]
WC	Ru	84.8 [2]	85.8 [1]	82.6 [1]	72.8 [1]	88.0 [1]
	En	92.6 [1]	93.7 [1]	89.8 [1]	82.3 [1]	93.8 [1]

Table 1: Results of Logistic Regression classifier for each encoder over the shared English and Russian tasks. Languages: **Ru**=Russian, **En**=English.

over joint Russian Wikipedia and Lenta.ru news, and released by DeepPavlov (Burtsev et al., 2018).

5 Results

How is the linguistic knowledge of two contrasting languages distributed in pre-trained multilingual encoders? This section describes the results for shared English and Russian probing tasks, covering surface properties (**SentLen**, **WC**), syntax (**TreeDepth**, **NShift**), and semantics (**ObjNumber**, **SubjNumber**, and **Tense**). We also report the results for the remaining tasks and the baselines in Appendix B.

5.1 Layer-wise Supervised Probing

Table 1 presents the results of the linear classifier performance over the shared tasks. For the sake of space, we omit the results of the non-linear classifier and provide them in Appendix B. The best score for each task in each language is highlighted in grey, and the index number of the layer achieving the score is enclosed in square brackets. We observe that the surface and syntax tasks show similar trends for both languages, while a few semantic tasks reveal some differences. The overall pattern for the surface-level tasks (**SentLen**, **WC**) is that

the probing curves⁸ for both languages tend to be decaying after reaching the peak at the lower layers [1 – 4]. The exception is provided by **M-BART** which keeps the surface properties across the majority of layers. The baselines mostly perform on par with one another, with **fastText** and **TF-IDF Char** receiving the best score (see Appendix B).

(**Nshift**, **TreeDepth**) tasks demonstrate interesting distinctive features of **LABSE** and **M-BART** encoders with respect to the syntactic properties. Figure 1 depicts the probing curves over (**Nshift**) and (**Bshift**) tasks. First, **LABSE** is at most sensitive to the incorrect word order at the lower-to-middle layers [2 – 8], as opposed to other encoders which distribute the information at the middle-to-higher layers [8 – 12]. Second, **M-BART** shows confident knowledge in **TreeDepth** task for both languages at the middle-to-higher layers [7 – 12] and reaches the peak at layer 12, in contrast to the rest of the encoders which generally distribute the property at the middle layers [5 – 7] (see Appendix B). The baseline performance varies from achieving a low score (**TreeDepth**) to being close to random choice (**Nshift**).

The probing curves for (**PT**) and (**Tense**) tasks illustrate that the models encode the property in a

⁸We refer to *probing curves* as to the performance trajectory of a probing classifier




Figure 1: The probing results for each encoder on **NShift (Ru)**, left) and **Bshift (En)**, right) tasks. X-axis=Layer index number, Y-axis=Accuracy score.

very similar fashion, achieving the peak score at the middle layers [4, 5], and flattening the curves until the output layer (see Appendix B). In contrast, the behavior of the encoders on (**SubjNumber**) and (**ObjNumber**) is slightly different. The number of the subject for English is predominantly distributed at the middle-to-higher levels [5 – 12], while for Russian it is either at a steady pace from the very first layer (**M-BART**, **M-BERT**), or decaying at the lower and middle layers [3 – 5], and further increasing close to the higher levels [8–12] (**LABSE**, **MiniLM**). Other differences are found in the results for (**ObjNumber**) task. Specifically, the number of direct object for English is best inferred at the lower layers of **LABSE** [2, 3], as compared to the middle-to-higher layers of other encoders [6 – 9]. Despite this, the property is similarly distributed across the layers of the models. In the same manner as for English, **LABSE** concentrates the knowledge for the Russian task in layer 2 but decays once reaching the peak. However, the other models exhibit a distinct behavior as well, with the property best encoded at the middle layers [6, 7] (**M-BERT**, **M-BART**), or at layer 10 (**XLM-R**, **MiniLM**). Notably, the baselines receive a strong performance over the majority of the semantic tasks for both languages, meaning that they can be solved using sub-word features that can capture lexical, or morphosyntactic information.

Bootstrap We estimate statistical significance of the supervised probing by means of layer-wise bootstrap procedure (Berg-Kirkpatrick et al., 2012). Typically, we observe that when the probing curve rises, the layer-wise difference is statistically sig-

nificant, while at the peak of the probing curve it turns insignificant. We can treat these observations in the following way: starting from the peak of the probing curve, the models acquire the knowledge needed for the task, and do not generally lose it in the higher layers.

5.2 Neuron-level Analysis

In contrast to layer-wise supervised probing (Section 5.1) which introspects each layer independently, individual neuron analysis allows exploring the distribution of top neurons selected from *the entire encoder*. This provides an alternative perspective to which layers contribute predominantly towards the probing tasks. Note that we trained distinct classifiers regularized with Elastic-net (Zou and Hastie, 2005) to estimate the importance of the neurons⁹. Figure 2 presents the results for (**SentLen**) tasks in Russian and English. Different from other models which predominantly capture the property by neurons at the lower layers [1 – 4], **M-BART** distributes the information across all the layers. Let us analyze the results with respect to some syntactic and semantic levels. A similar behavior of the encoders by language is observed over (**NShift**) and (**TreeDepth**) tasks (see Appendix C). **M-BERT**, **XLM-R**, and **MiniLM** generally capture the sensitivity to illegal word order by neurons at the middle and higher layers, while it is contributed by fewer layers of **LABSE (Ru)**: [2 – 4]; **En**: [4 – 7]), and **M-BART** which surprisingly stores the property at layer 12 for each language. Another interesting finding is that the depth of the

⁹We selected top-20% neurons from each encoder using neuron ranking algorithm described in (Durrani et al., 2020)




Figure 2: The distribution of top neurons over **SentLen** tasks for both languages: **Ru**=Russian, **En**=English. X-axis=Layer index number, Y-axis=Number of neurons selected from the layer.

syntax tree is typically spread across most of the layers of each encoder. An exception to this pattern is **M-BART** which locates the knowledge at the middle-to-higher layers [7 – 11].

The neuron distributions for (**ObjNumber**) task are akin by language and slightly different by encoder, showing that the number of the direct object is learned by at least 5-7 layers of different encoder depth. At the same time, the properties for (**SubjNumber**) and (**PT, Tense**) tasks are captured differently by the encoders. **M-BERT**, **XLM-R** and **MiniLM** reveal similar behavior by task, whereas **LABSE** specializes the knowledge at either the lower layers (**Ru, SubjNumber**; and **En, Tense**: [1 – 4]), or the higher ones (**En, SubjNumber**; and **Ru, PT**: [10 – 12]). On the other hand, the most contributing neurons of **M-BART** are predominantly spread across all the layers.

5.3 Correlation Analysis

For analyzing the encoders by means of the correlation-based techniques, we used 1k stratified sentences from each test set of the shared tasks. We obtained the sentence representations and computed the measures applying the publicly-available code (Wu et al., 2020). Figure 3 shows heatmaps of similarities between layers of the encoders under neuron-level and representation-level measures for English. Notably, the heatmaps are very alike to the ones for Russian, which we enclose in Appendix D. `maxcorr` (Figure 3a) demonstrates that different layers of a single encoder have similar individual neurons, but the inter-encoder neuron similarities are greatly low. On the contrary, `lincka` (Figure 3b) induces considerably high similarities across the encoders, meaning that they produce sim-

ilar sentence representations. However, **M-BERT** and **M-BART** show lower similarity with other encoders, particularly at the lower and higher layers (**M-BERT**: [1 – 3], [10 – 12]; **M-BART**: [10, 11]). Besides, **M-BART** and **M-BERT** demonstrate low pairwise similarity, being fairly different at the lower-to-higher layers [3 – 11].

6 Discussion

Three encoders exhibit similar behavior, but two other differ from them in capturing linguistic properties The most striking distinction based upon the results is that **M-BART** and **LABSE** exhibit different behavior for both languages, as opposed to **M-BERT**, **XLM-R**, and **MiniLM** (Section 5.1, 5.2). Specifically, **M-BART** generally tends to distribute the surface properties across all the layers, unlike other encoders where the information is specialized by the lower ones. The syntactic properties tend to be localized at the higher layers [11 – 12] which is as well demonstrated in other Russian tasks (**Gapping, ImpersonalSent**). In contrast to other models that capture the semantic properties at the middle-to-higher layers, **LABSE** typically displays the knowledge at the lower-to-middle layers. The analysis of the understudied encoders contradicts the common findings on transformer models that syntactic information is stored at the middle layers, while semantic knowledge is most prominent at the higher layers (Rogers et al., 2021). Therefore, there is still room for exploring transformer-based models that differ by architecture type, or by a set of pre-training objectives.




Figure 3: Similarity heatmaps of layers in the encoders under neuron-level (`maxcorr`) and representation-level (`lincka`) measures.

MiniLM inherits linguistic properties from M-BERT Teacher models are usually compared with their distilled versions on a range of downstream tasks (Tsai et al., 2019), and little is explored on what language properties they bequeath to their students. We find that **MiniLM** is likely to mimic the behavior of **XLM-R** rather than of **M-BERT**, most probably due to using the same tokenization method. The similarity is most demonstrated by the individual neuron analysis (see Appendix C) and `lincka` (see Figure 3b). Along with that, the model receives comparative performance under layer-wise supervised probing (see Appendix B).

Surface and syntactic information is learned in a similar manner The probing curves demonstrate that the surface and semantic properties of the two languages are similarly distributed in the encoders. The surface properties are generally captured at the lower layers, and the pattern of the curves is decaying towards the output layer. The syntactic properties are predominantly inferred at the middle or higher layers, and the semantic tasks reveal a number of differences described in Section 5. Notably, the results obtained under layer-wise supervised probing (Section 5.1) are supported by the individual neuron analysis (Section 5.2), and the representation-level analysis (Section 5.3).

Encoders may have similar distributions of neurons by task, but different individual neurons Two neuron-level introspection methods allow

drawing the following finding. Despite that top-neuron distributions in the encoder layers share similar patterns in the majority of probing tasks (Section 5.2), `maxcorr` induces high intra-encoder similarities and low inter-encoder similarities (Section 5.3). In other words, the neurons can similarly localize particular properties in the layers and yet behave differently across the models.

7 Conclusion

This paper introduces RuSentEval, an enhanced probing suite of 14 probing tasks that cover various linguistic phenomena of the Russian language. We explored five multilingual transformer encoders over the probing tasks on two typologically contrasting languages – Russian and English. The experiments are conducted using a combination of complementary probing methods, including layer-wise supervised probing, individual neuron-analysis, neuron- and representation-level similarity measures. Particularly, the behavior of the encoders under probing classifiers is reflected in distributions of top neurons with respect to a task, and the similarity is supported by linear centered kernel alignment method. We found that despite the language distinctions, the surface and syntactic properties are learned in a fairly similar manner, and the semantic knowledge is captured differently in a majority of tasks. We believe that the findings make the ongoing studies on cross-lingual transfer even more promising, specifically from English to Russian or vice versa. In contrast to prior works on

how linguistic knowledge is represented in transformers, the analysis of the understudied models reveals that syntax and semantics can be differently represented across the layers. Besides, we found that different encoders often have similar distributions of neurons that contribute most to a probing task, and yet differ under neuron-level similarities. We also observed that distilled models inherit linguistic properties from their teachers, and receive comparative performance on a number of probing tasks. An exciting direction for future work is to investigate the correlation between probing and high-level downstream tasks, in order to identify which linguistic properties anticipate the behavior of a model in action.

Acknowledgments

Ekaterina Artemova works within the framework of the HSE University Basic Research Program.

References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2016. Fine-grained Analysis of Sentence Embeddings Using Auxiliary Prediction Tasks. *arXiv preprint arXiv:1608.04207*.
- Christoph Alt, Aleksandra Gabryszak, and Leonhard Hennig. 2020. Probing Linguistic Features of Sentence-Level Representations in Relation Extraction. pages 1534–1545.
- Chiara Alzetta, Felice Dell’Orletta, Simonetta Montemagni, and Giulia Venturi. 2017. Dangerous Relations in Dependency Treebanks. pages 201–210.
- Daniil Anastasyev. 2020. Exploring Pretrained Models For Joint Morpho-syntactic Parsing of Russian.
- Matthew Baerman. 2007. Syncretism. *Language and Linguistics Compass*, 1(5):539–551.
- Yonatan Belinkov and James Glass. 2019. Analysis Methods in Neural Language Processing: a Survey. *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. An Empirical Investigation of Statistical Significance in NLP. pages 995–1005.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Mikhail Burtsev, Alexander Seliverstov, Rafael Airapetyan, Mikhail Arkhipov, Dilyara Baymurzina, Nickolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yurii Kuratov, Denis Kuznetsov, et al. 2018. DeepPavlov: Open-source Library for Dialogue Systems. pages 122–127.
- Rochelle Choenni and Ekaterina Shutova. 2020a. Cross-neutralising: Probing for Joint Encoding of Linguistic Information in Multilingual Models. *arXiv e-prints*, pages arXiv–2010.
- Rochelle Choenni and Ekaterina Shutova. 2020b. What Does it Mean to be Language-agnostic? Probing Multilingual Sentence Encoders for Typological Properties. *arXiv preprint arXiv:2009.12862*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. Unsupervised Cross-lingual Representation Learning at Scale. pages 8440–8451.
- Alexis Conneau and Douwe Kiela. 2018. SentEval: An Evaluation Toolkit for Universal Sentence Representations.
- 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. *arXiv preprint arXiv:1805.01070*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual Language Model Pre-training. pages 7059–7069.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Holger Schwenk, Ves Stoyanov, Adina Williams, and Samuel R Bowman. 2020b. XNLI: Evaluating Cross-lingual Sentence Representations. pages 2475–2485.
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Yonatan Belinkov, Anthony Bau, and James Glass. 2019. What is One Grain of Sand in the Desert? Analyzing Individual Neurons in Deep NLP Models. 33:6309–6317.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. pages 4171–4186.
- Kira Droganova, Olga Lyashevskaya, and Daniel Zeman. 2018. Data Conversion and Consistency of Monolingual Corpora: Russian UD Treebanks. pages 52–65.
- Nadir Durrani, Hassan Sajjad, Fahim Dalvi, and Yonatan Belinkov. 2020. Analyzing Individual Neurons in Pre-trained Language Models. pages 4865–4880.
- Steffen Eger, Johannes Daxenberger, and Iryna Gurevych. 2020. How to Probe Sentence Embeddings in Low-Resource Languages: On Structural Design Choices for Probing Task Evaluation. pages 108–118.

- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic BERT Sentence Embedding. *arXiv preprint arXiv:2007.01852*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. XTREME: A Massively Multilingual Multi-task Benchmark for Evaluating Cross-lingual Generalisation. pages 4411–4421.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sabrina J Mielke, Arya D McCarthy, Sandra Kübler, et al. 2018. UniMorph 2.0: Universal Morphology.
- Katarzyna Krasnowska-Kieraś and Alina Wróblewska. 2019. Empirical Linguistic Study of Sentence Embeddings. pages 5729–5739.
- Taku Kudo and John Richardson. 2018. SentencePiece: A Simple and Language Independent Subword Tokenizer and Detokenizer for Neural Text Processing. pages 66–71.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, et al. 2020. XGLUE: A New Benchmark Dataset for Cross-lingual Pre-training, Understanding and Generation. pages 6008–6018.
- Nelson F Liu, Matt Gardner, Yonatan Belinkov, Matthew E Peters, and Noah A Smith. 2019. Linguistic Knowledge and Transferability of Contextual Representations. pages 1073–1094.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Olga Lyashevskaya and Sergey Sharov. 2009. The Frequency Dictionary of Modern Russian Language. *Azbukovnik, Moscow*.
- David Mareček, Hande Celikkanat, Miikka Silfverberg, Vinit Ravishankar, and Jörg Tiedemann. 2020. Are Multilingual Neural Machine Translation Models Better at Capturing Linguistic Features? *The Prague Bulletin of Mathematical Linguistics*, (115):143–162.
- Marie-Catherine de Marneffe, Matias Grioni, Jenna Kanerva, and Filip Ginter. 2017. Assessing the Annotation Consistency of the Universal Dependencies Corpora. pages 108–115.
- Julian Michael, Jan A Botha, and Ian Tenney. 2020. Asking without Telling: Exploring Latent Ontologies in Contextual Representations. *arXiv preprint arXiv:2004.14513*.
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. Universal Dependencies v1: A Multilingual Treebank Collection. pages 1659–1666.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12:2825–2830.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How Multilingual is Multilingual BERT? pages 4996–5001.
- Maria Ponomareva, Kira Droганova, Ivan Smurov, and Tatiana Shavrina. 2019. AGRR 2019: Corpus for Gapping Resolution in Russian. pages 35–43.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. *OpenAI blog*, 1(8):9.
- Vinit Ravishankar, Lilja Øvrelid, and Erik Velldal. 2019. Probing Multilingual Sentence Representations With XProbe. *ACL 2019*, page 156.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2021. A Primer in BERTology: What we Know about how BERT Works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Gözde Gül Şahin, Clara Vania, Ilia Kuznetsov, and Iryna Gurevych. 2020. Linspector: Multilingual Probing Tasks for Word Representations. *Computational Linguistics*, 46(2):335–385.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. pages 1715–1725.
- Tatiana Shavrina and Olga Shapovalova. 2017. To the Methodology Corpus Construction for Machine Learning: “TAIGA” Syntax Tree Corpus and Parser. *CORPUS LINGUISTICS*, 2017, page 78.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. Does String-based Neural MT Learn Source Syntax? pages 1526–1534.
- Xavier Suau, Luca Zappella, and Nicholas Apostoloff. 2020. Finding Experts in Transformer Models. *arXiv preprint arXiv:2005.07647*.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim,

- Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, et al. 2019. What do you Learn from Context? Probing for Sentence Structure in Contextualized Word Representations. *arXiv preprint arXiv:1905.06316*.
- Henry Tsai, Jason Riesa, Melvin Johnson, Naveen Arivazhagan, Xin Li, and Amelia Archer. 2019. Small and Practical BERT Models for Sequence Labeling. pages 3623–3627.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. pages 5998–6008.
- Elena Voita and Ivan Titov. 2020. Information-Theoretic Probing with Minimum Description Length. *arXiv preprint arXiv:2003.12298*.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. SuperGLUE: A Stickier Benchmark for General-purpose Language Understanding Systems. pages 3266–3280.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep Self-attention Distillation for Task-agnostic Compression of Pre-trained Transformers. *arXiv preprint arXiv:2002.10957*.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Édouard Grave. 2020. CCNet: Extracting High Quality Monolingual Datasets from Web Crawl Data. pages 4003–4012.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv*, pages arXiv–1910.
- John Wu, Yonatan Belinkov, Hassan Sajjad, Nadir Durani, Fahim Dalvi, and James Glass. 2020. Similarity Analysis of Contextual Word Representation Models. pages 4638–4655.
- Hui Zou and Trevor Hastie. 2005. Regularization and Variable Selection via the Elastic Net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2):301–320.

A Examples from RuSentEval Tasks

Table 2 provides with examples from RuSentEval tasks.

Task	Example	Label
ConjType	On otmetil , chto podobnyye progulki nebezopasny . 'He noted that such walks are unsafe.'	SCONJ
Gapping	Ya yezdila dvazhdy , sestra – trizhdy . 'I went [there] twice, my sister [went there] three times'	1
ImpersonalSent	Rabotal takzhe kak kontsertmeyster i lektor . '[He] also worked as an accompanist and lecturer.'	0
NShift	Kogda etogo poluchilos' ne , on ubezhal . 'When it work out didn't , he ran away'	I
ObjGender	Rossiyskiy duet dopustil odnu oshibku . 'The Russian duo made one mistake .'	F
ObjNumber	Serial poluchil neskol'ko prestizhnykh nagrad . 'The series has received several prestigious awards .'	NNS
PA	On nikak ne ob'yasn il svoyu pozitsiyu . 'He did not explain his position in any way .'	PERF
PT	Molodyye spetsialisty poluchayut yezhemesyachnuyu doplatu k zarplate . 'Young professionals receive a monthly supplement to their salary .'	PRES
PV	Srok vozmozhnoy prem'yery lenty poka ne nazyvayetsya 'The date of a possible premiere of the film has not yet been announced .'	PASS
SentLen	Ya ne videla boleye zlogo cheloveka . 'I haven't seen a more angry man.'	0
SubjGender	On nosit beluyu dlinnuyu rubashku i dlinnyye seryye bryuki. ' He wears a white long shirt and long grey trousers.'	M
SubjNumber	On byl lyubimtsem vsej moskovskoy i peterburgskoy aristokratii . ' He was a favorite of the entire Moscow and St. Petersburg aristocracy.'	NN
TreeDepth	I vot v pervuyu ochered' my khoteli by pogovorit' ob etom . 'And first of all, we would like to talk about this.'	5
WC	Proshluyu noch' ya sovsem ne spal . 'I didn't sleep at all last night.'	spat' 'to sleep'

Table 2: Examples from RuSentEval tasks.

B Layer-wise Supervised Probing

B.1 Results on All Tasks

The results reported in the main body of the paper are obtained with Logistic Regression classifier over the shared tasks (see Section 5.1). We present detailed results for both linear and non-linear classifiers on all Russian and English tasks in Tables 3–6. Tables 7–9 show the results of the baselines. Table 11 outlines statistical description of the tasks.

Probing Task	M-BERT	LABSE	XLM-R	MiniLM	M-BART
ConjType	98.8 [7]	99.3 [4]	99.3 [6]	98.6 [5]	98.8 [7]
Gapping	85.2 [7]	89.7 [8]	94.1 [8]	91.1 [9]	85.6 [12]
ImpersonalSent	91.6 [7]	92 [6]	92.6 [4]	88.4 [6]	85.7 [12]
NShift	81.8 [10]	82.6 [5]	86.9 [9]	80.5 [9]	78.6 [12]
ObjGender	70.1 [6]	70.4 [2]	69.4 [5]	64.1 [9]	71.8 [1]
ObjNumber	82.8 [6]	82.5 [2]	83.7 [10]	77.8 [10]	81.5 [7]
PA	91.2 [6]	93.8 [4]	94.4 [5]	89.4 [5]	95.9 [10]
PT	99.5 [8]	99.8 [5]	99.8 [5]	98.2 [7]	99.6 [7]
PV	77.5 [5]	76.3 [5]	76.8 [5]	71.4 [3]	77.7 [3]
SentLen	91.3 [2]	93.3 [1]	94.5 [2]	94.1 [2]	96.2 [4]
SubjGender	79.1 [9]	79.2 [2]	79.4 [11]	78.7 [10]	77.7 [7]
SubjNumber	90.5 [7]	92.9 [3]	94.9 [11]	94.2 [12]	93.1 [10]
TreeDepth	44.7 [6]	46.1 [4]	46.5 [5]	44.8 [7]	45.8 [11]
WC	84.8 [2]	85.8 [1]	82.6 [1]	72.8 [1]	88.0 [1]

Table 3: Results of Logistic Regression classifier by the encoder for RuSentEval tasks.

Probing Task	M-BERT	LABSE	XLM-R	MiniLM	M-BART
BShift	84.8 [8]	84.4 [5]	85.7 [10]	79.3 [8]	83.8 [12]
CoordInv	66.0 [8]	68.9 [8]	68.6 [8]	63.3 [8]	69.4 [12]
ObjNumber	86.2 [6]	85.4 [3]	86.0 [8]	85.2 [6]	85.9 [9]
SOMO	57.4 [8]	60.8 [7]	60.0 [8]	56.1 [9]	62.3 [12]
SentLen	96.3 [2]	96.6 [1]	95.8 [2]	96.1 [3]	97.3 [3]
SubjNumber	87.8 [7]	90.7 [12]	86.9 [10]	85.6 [6]	87.3 [9]
Tense	88.9 [8]	88.8 [6]	88.8 [9]	87.3 [5]	89.1 [9]
TopConst	88 [6]	79.9 [5]	78.5 [5]	76.5 [5]	79.5 [8]
TreeDepth	41.2 [5]	42.7 [5]	41.8 [7]	40.9 [7]	41.2 [12]
WC	92.6 [1]	93.7 [1]	89.8 [1]	82.3 [1]	93.8 [1]

Table 4: Results of Logistic Regression classifier by the encoder for SentEval tasks.

Probing Task	M-BERT	LABSE	XLM-R	MiniLM	M-BART
ConjType	98.6 [5]	99.4 [4]	99.2 [5]	98.9 [5]	98.8 [7]
Gapping	89.7 [10]	90.0 [9]	96.0 [8]	92.0 [11]	83.1 [9]
ImpersonalSent	93.6 [7]	90.9 [5]	92.4 [7]	88.7 [7]	89.4 [9]
NShift	81.5 [9]	82.7 [5]	87.6 [9]	81.1 [10]	78.8 [12]
ObjGender	69.1 [6]	70.1 [2]	69.5 [9]	65.1 [10]	72.2 [1]
ObjNumber	83.9 [6]	82.5 [2]	84.8 [10]	78.7 [10]	83.0 [7]
PA	90.9 [7]	93.5 [5]	94.6 [5]	89.7 [5]	95.5 [8]
PT	99.4 [4]	99.9 [10]	99.8 [5]	98.4 [6]	99.6 [9]
PV	77.7 [4]	76.5 [5]	78.4 [4]	72.5 [4]	82.2 [1]
SentLen	93.5 [2]	95.2 [1]	97.1 [2]	96.7 [1]	98.2 [5]
SubjGender	79.5 [9]	80.0 [2]	81.0 [11]	80.2 [12]	78.1 [7]
SubjNumber	90.3 [5]	93.0 [3]	96.3 [12]	95.8 [11]	94.5 [7]
TreeDepth	43.6 [6]	45.4 [4]	44.8 [7]	46.7 [8]	46.0 [7]
WC	80.8 [3]	82.7 [1]	78.5 [1]	69.9 [1]	84.4 [1]

Table 5: Results of MLP classifier by the encoder for RuSentEval tasks.

Probing Task	M-BERT	LABSE	XLM-R	MiniLM	M-BART
BShift	83.1 [8]	84.7 [6]	85.8 [9]	79.4 [7]	84.4 [12]
CoordInv	65.2 [8]	68.1 [8]	68.8 [8]	63.7 [8]	67.8 [10]
ObjNumber	86.5 [6]	86.4 [10]	86.4 [8]	85.0 [6]	86.8 [8]
SOMO	56.5 [8]	60.5 [7]	58.8 [8]	54.8 [9]	61.7 [11]
SentLen	97.0 [2]	98.4 [1]	98.0 [3]	98.5 [1]	98.8 [4]
SubjNumber	86.5 [10]	90.9 [11]	86.6 [6]	86.0 [6]	87.5 [9]
Tense	89.2 [9]	89.0 [6]	88.5 [10]	87.9 [5]	89.4 [9]
TopConst	82.0 [7]	80.6 [5]	79.5 [5]	77.8 [6]	80.6 [8]
TreeDepth	41.9 [6]	43.1 [5]	43.2 [7]	42.3 [6]	45.3 [10]
WC	91.2 [1]	92.7 [1]	88.9 [1]	80.3 [1]	93.0 [1]

Table 6: Results of MLP classifier by encoder for each SentEval task.

Probing Task	fastText	TF-IDF Char	TF-IDF BPE	TF-IDF SP
ConjType	88.1	96.9	95.4	95.5
Gapping	84.1	82.7	80.4	80.6
ImpersonalSent	78.7	69.4	53.8	56.3
NShift	53.2	53.0	51.0	50.5
ObjGender	70.1	71.0	35.4	38.9
ObjNumber	82.3	76.4	56.8	55.0
PA	90.8	80.7	53.4	54.2
PT	95.1	97.7	53.8	53.7
PV	69.2	78.2	36.0	37.0
SentLen	40.4	64.0	42.9	42.2
SubjGender	78.7	74.4	34.8	38.0
SubjNumber	95.0	90.4	63.7	64.4
TreeDepth	35.7	32.7	26.5	24.8
WC	70.8	49.2	22.0	13.0

Table 7: Results of Logistic Regression classifier by the baseline feature for each RuSentEval task.

Probing Task	fastText	TF-IDF Char	TF-IDF BPE	TF-IDF SP
BShift	50.0	51.1	49.9	50.1
CoordInv	52.2	54.9	50.2	50.1
ObjNumber	72.8	79.4	68.1	69.0
SOMO	49.9	49.9	50.4	49.7
SentLen	65.2	54.1	42.3	44.6
SubjNumber	76.6	79.2	68.1	71.6
Tense	81.2	84.2	70.8	74.2
TopConst	59.8	58.3	23.0	23.4
TreeDepth	30.0	28.3	23.3	23.2
WC	18.1	47.3	20.0	24.0

Table 8: Results of Logistic Regression classifier by the baseline feature for each SentEval task.

Probing Task	fastText	TF-IDF Char	TF-IDF BPE	TF-IDF SP
ConjType	88.4	97.3	95.6	95.5
Gapping	82.7	86.1	80.2	68.8
ImpersonalSent	78.6	70.5	52.9	56.6
NShift	52.7	52.7	50.0	50.6
ObjGender	70.0	70.9	35.1	37.2
ObjNumber	82.8	77.2	56.6	54.8
PA	91.2	80.9	51.8	53.5
PT	96.0	97.6	54.4	54.1
PV	68.5	78.5	35.3	36.8
SentLen	42.4	73.7	42.7	42.4
SubjGender	80.0	75.2	34.0	38.8
SubjNumber	96.2	90.8	61.8	64.4
TreeDepth	29.5	35.6	32.8	23.9
WC	71.2	53.8	20.0	11.0

Table 9: Results of MLP classifier by the baseline feature for each RuSentEval task.

Probing Task	fastText	TF-IDF Char	TF-IDF BPE	TF-IDF SP
BShift	48.2	50.6	50.0	49.3
CoordInv	50.1	54.0	50.0	51.7
ObjNumber	70.9	77.1	68.1	70.0
SOMO	50.1	49.9	50.2	50.2
SentLen	49.1	62.5	41.8	43.5
SubjNumber	72.8	80.5	66.4	71.3
Tense	74.7	85.0	70.5	73.8
TopConst	58.0	59.7	22.2	23.0
TreeDepth	23.0	29.5	23.0	22.1
WC	63.3	54.4	18.0	22.0

Table 10: Results of MLP classifier by the baseline feature for each SentEval task.

	RuSentEval				SentEval			
	Train	Dev	Test	Overall	Train	Dev	Test	Overall
	SentLen				SentLen			
sample size	100k	10k	10k	120k	100k	10k	10k	120k
tokens	1.45kk	144.8k	144.8k	1.74kk	1.66kk	165.4k	165.4k	1.99kk
unique tokens	116.7k	34.0k	33.7k	126.5k	34.8k	9.6k	10.0k	36.8k
tokens/sentence	14.47	14.48	14.48	14.47	16.59	16.54	16.55	16.59
label distribution	16.7/16.7/16.7/16.7/16.7/16.7				16.7/16.7/16.7/16.7/16.7/16.7			
	WC				WC			
sample size	100k	10k	10k	120k	100k	10k	10k	120k
tokens	1.19kk	117.9k	118.9k	1.43kk	1.5kk	149.2k	149.8k	1.8kk
unique tokens	106.6k	30.6k	30.8k	115.7k	37.4k	13.4k	13.4k	40.1k
tokens/sentence	11.89	11.79	11.89	11.88	15.02	14.92	14.98	15.00
label distribution	0.1/label, 1000 labels				0.1/label, 1000 labels			
	NShift				BShift			
sample size	100k	10k	10k	120k	100k	10k	10k	120k
tokens	1.49kk	125.0k	127.4k	1.74kk	1.38kk	137.4k	136.4k	1.65kk
unique tokens	138.2k	29.2k	29.6k	146.9k	36.2k	12.8k	12.7k	38.7k
tokens/sentence	14.88	12.74	12.51	14.51	13.78	13.74	13.64	13.77
label distribution	50/50				50/50			
	SubjNumber				SubjNumber			
sample size	100k	10k	10k	120k	100k	10k	10k	120k
tokens	1.04kk	99.9k	103.2k	1.25kk	1.41kk	140.3k	141.7k	1.7kk
unique tokens	100.3k	21.7k	23.1k	108.3k	38.5k	14.4k	14.5k	41.3k
tokens/sentence	10.42	9.99	10.32	10.38	14.14	14.03	14.17	14.13
label distribution	50/50				50/50			
	ObjNumber				ObjNumber			
sample size	100k	10k	10k	120k	100k	10k	10k	120k
tokens	946.7k	103.3k	100.2k	1.15kk	1.4kk	140.5k	139.9k	1.68kk
unique tokens	86.0k	23.8k	22.8k	95.7k	38.0k	14.3k	13.9k	40.9k
tokens/sentence	9.47	10.33	10.02	9.59	13.96	14.05	13.99	13.97
label distribution	50/50				50/50			
	PT				Tense			
sample size	100k	10k	10k	120k	100k	10k	10k	120k
tokens	1.13kk	112.8k	113.0k	1.36kk	1.32kk	131.1k	129.6k	1.58kk
unique tokens	114.5k	26.8k	26.7k	126.6k	35.9k	13.1k	13.2k	38.6k
tokens/sentence	11.30	11.28	11.30	11.30	13.20	13.11	12.96	13.17
label distribution	50/50				50/50			
	TreeDepth				TreeDepth			
sample size	100k	10k	10k	120k	100k	10k	10k	120k
tokens	1.57kk	157.2k	157.6k	1.88kk	1.35kk	135.0k	134.7k	1.62kk
unique tokens	150.3k	38.4k	38.3k	165.7k	34.8k	12.5k	12.6k	37.1k
tokens/sentence	15.72	15.72	15.76	15.73	13.47	13.50	13.47	13.47
label distribution	13.6/22.6/32.9/20.7/10.3				15.4/11.9/7.0/ 13.6/16.5/17.9/17.7			

Table 11: Comparative data statistics for the shared Russian and English probing tasks. Label distribution by target class is presented in %.

B.2 Probing Trajectories

Figures 4–9 show the probing curves of Logistic Regression classifier over the shared tasks (see Section 5.1). The results obtained with MLP classifier are pretty consistent with the ones presented in this Appendix.




Figure 4: The probing results of Logistic Regression classifier for each encoder on **ObjNumber**. **Ru** is at the left; and **En** is at the right. X-axis=Layer number, Y-axis=Accuracy score.




Figure 5: The probing results of Logistic Regression classifier for each encoder on **SentLen**. **Ru** is at the left; and **En** is at the right. X-axis=Layer number, Y-axis=Accuracy score.




Figure 6: The probing results of Logistic Regression classifier for each encoder on **SubjNumber**. **Ru** is at the left; and **En** is at the right. X-axis=Layer number, Y-axis=Accuracy score.




Figure 7: The probing results of Logistic Regression classifier for each encoder on **TreeDepth**. **Ru** is at the left; and **En** is at the right. X-axis=Layer number, Y-axis=Accuracy score.




Figure 8: The probing results of Logistic Regression classifier for each encoder on **Tense**. **Ru (PT)** is at the left; and **En (Tense)** is at the right. X-axis=Layer number, Y-axis=Accuracy score.




Figure 9: The probing results of Logistic Regression classifier for each encoder on **WC**. **Ru** is at the left; and **En** is at the right. X-axis=Layer number, Y-axis=Accuracy score.

C Individual Neuron Analysis

Figures 10–14 depict top neuron distributions for the tasks (see Section 5.2).

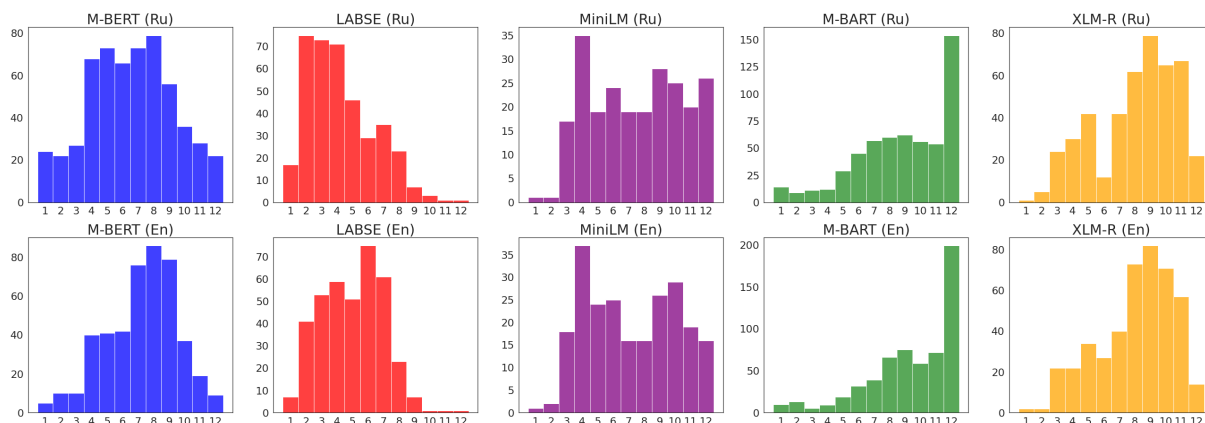


Figure 10: The distribution of top neurons over **NShift** tasks for both languages: **Ru**=Russian, **En**=English. X-axis=Layer index number, Y-axis=Number of neurons selected from the layer.




Figure 11: The distribution of top neurons over **ObjNumber** tasks for both languages: **Ru**=Russian, **En**=English. X-axis=Layer index number, Y-axis=Number of neurons selected from the layer.




Figure 12: The distribution of top neurons over **SubjNumber** tasks for both languages: **Ru**=Russian, **En**=English. X-axis=Layer index number, Y-axis=Number of neurons selected from the layer.




Figure 13: The distribution of top neurons over **TreeDepth** tasks for both languages: **Ru**=Russian, **En**=English. X-axis=Layer index number, Y-axis=Number of neurons selected from the layer.




Figure 14: The distribution of top neurons over **Tense** tasks for both languages: **Ru**=Russian, **En**=English. X-axis=Layer index number, Y-axis=Number of neurons selected from the layer.

D Correlation Methods

Heatmaps (Figure 15) show similarities of the encoders under neuron-level and representation-level correlation-based similarity measures on the Russian tasks (see Section 5.3).




Figure 15: Similarity heatmaps of layers in the encoders under neuron-level (`maxcorr`) and representation-level (`lincka`) measures for Russian.