

# Linguistic Profiling of a Neural Language Model

Alessio Miaschi<sup>\*◊</sup>, Dominique Brunato<sup>◊</sup>, Felice Dell’Orletta<sup>◊</sup>, Giulia Venturi<sup>◊</sup>

<sup>\*</sup>Department of Computer Science, University of Pisa

<sup>◊</sup>Istituto di Linguistica Computazionale “Antonio Zampolli”, Pisa

ItaliaNLP Lab – [www.italianlp.it](http://www.italianlp.it)

[alessio.miaschi@phd.unipi.it](mailto:alessio.miaschi@phd.unipi.it), [{name.surname}@ilc.cnr.it](mailto:{name.surname}@ilc.cnr.it)

## Abstract

In this paper we investigate the linguistic knowledge learned by a Neural Language Model (NLM) before and after a fine-tuning process and how this knowledge affects its predictions during several classification problems. We use a wide set of probing tasks, each of which corresponds to a distinct sentence-level feature extracted from different levels of linguistic annotation. We show that BERT is able to encode a wide range of linguistic characteristics, but it tends to lose this information when trained on specific downstream tasks. We also find that BERT’s capacity to encode different kind of linguistic properties has a positive influence on its predictions: the more it stores readable linguistic information of a sentence, the higher will be its capacity of predicting the expected label assigned to that sentence.

## 1 Introduction

Neural Language Models (NLMs) have become a central component in NLP systems over the last few years, showing outstanding performance and improving the state-of-the-art on many tasks (Peters et al., 2018; Radford et al., 2018; Devlin et al., 2019). However, the introduction of such systems has come at the cost of interpretability and, consequently, at the cost of obtaining meaningful explanations when automated decisions take place.

Recent work has begun to study these models in order to understand whether they encode linguistic phenomena even without being explicitly designed to learn such properties (Marvin and Linzen, 2018; Goldberg, 2019; Warstadt et al., 2019). Much of this work focused on the analysis and interpretation of attention mechanisms (Tang et al., 2018; Jain and Wallace, 2019; Clark et al., 2019) and on the definition of *probing models* trained to predict simple linguistic properties from unsupervised representations.

Probing models trained on different contextual representations provided evidences that such models are able to capture a wide range of linguistic phenomena (Adi et al., 2016; Perone et al., 2018; Tenney et al., 2019b) and even to organize this information in a hierarchical manner (Belinkov et al., 2017; Lin et al., 2019; Jawahar et al., 2019). However, the way in which this knowledge affects the decisions they make when solving specific downstream tasks has been less studied.

In this paper, we extended prior work by studying the linguistic properties encoded by one of the most prominent NLM, BERT (Devlin et al., 2019), and how these properties affect its predictions when solving a specific downstream task. We defined three research questions aimed at understanding: (i) what kind of linguistic properties are already encoded in a pre-trained version of BERT and where across its 12 layers; (ii) how the knowledge of these properties is modified after a fine-tuning process; (iii) whether this implicit knowledge affects the ability of the model to solve a specific downstream task, i.e. Native Language Identification (NLI). To tackle the first two questions, we adopted an approach inspired to the ‘linguistic profiling’ methodology put forth by van Halteren (2004), which assumes that wide counts of linguistic features automatically extracted from parsed corpora allow modeling a specific language variety and detecting how it changes with respect to other varieties, e.g. complex vs simple language, female vs male-authored texts, texts written in the same L2 language by authors with different

---

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

L1 languages. Particularly relevant for our study, is that multi-level linguistic features have been shown to have a highly predictive role in tracking the evolution of learners’ linguistic competence across time and developmental levels, both in first and second language acquisition scenarios (Lubetich and Sagae, 2014; Miaschi et al., 2020).

Given the strong informative power of these features to encode a variety of language phenomena across stages of acquisition, we assume that they can be also helpful to dig into the issues of interpretability of NLMs. In particular, we would like to investigate whether features successfully exploited to model the evolution of language competence can be similarly helpful in profiling how the implicit linguistic knowledge of a NLM changes across layers and before and after tuning on a specific downstream task. We chose the NLI task, i.e. the task of automatically classifying the L1 of a writer based on his/her language production in a learned language (Malmasi et al., 2017). As shown by Cimino et al. (2018), linguistic features play a very important role when NLI is tackled as a sentence–classification task rather than as a traditional document–classification task. This is the reason why we considered the sentence-level NLI classification as a task particularly suitable for probing the NLM linguistic knowledge. Finally, we investigated whether and which linguistic information encoded by BERT is involved in discriminating the sentences correctly or incorrectly classified by the fine-tuned models. To this end, we tried to understand if the linguistic knowledge that the model has of a sentence affects the ability to solve a specific downstream task involving that sentence.

**Contributions** In this paper: (i) we carried out an in-depth linguistic profiling of BERT’s internal representations (ii) we showed that contextualized representations tend to lose their precision in encoding a wide range of linguistic properties after a fine-tuning process; (iii) we showed that the linguistic knowledge stored in the contextualized representations of BERT positively affects its ability to solve NLI downstream tasks: the more BERT stores information about these features, the higher will be its capacity of predicting the correct label.

## 2 Related Work

In the last few years, several methods have been devised to obtain meaningful explanations regarding the linguistic information encoded in NLMs (Belinkov and Glass, 2019). They range from techniques to examine the activations of individual neurons (Karpathy et al., 2015; Li et al., 2016; Kádár et al., 2017) to more domain specific approaches, such as interpreting attention mechanisms (Raganato and Tiedemann, 2018; Kovaleva et al., 2019; Vig and Belinkov, 2019), studying correlations between representations (Saphra and Lopez, 2019) or designing specific *probing tasks* that a model can solve only if it captures a precise linguistic phenomenon using the *contextual* word/sentence embeddings of a pre-trained model as training features (Conneau et al., 2018; Zhang and Bowman, 2018; Hewitt and Liang, 2019; Miaschi and Dell’Orletta, 2020). These latter studies demonstrated that NLMs are able to encode a variety of language properties in a hierarchical manner (Belinkov et al., 2017; Blevins et al., 2018; Tenney et al., 2019b) and even to support the extraction of dependency parse trees (Hewitt and Manning, 2019). Jawahar et al. (2019) investigated the representations learned at different layers of BERT, showing that lower layer representations are usually better for capturing surface features, while embeddings from higher layers are better for syntactic and semantic properties. Using a suite of probing tasks, Tenney et al. (2019a) found that the linguistic knowledge encoded by BERT through its 12/24 layers follows the traditional NLP pipeline: POS tagging, parsing, NER, semantic roles and then coreference. Liu et al. (2019), instead, quantified differences in the transferability of individual layers between different models, showing that higher layers of RNNs (ELMo) are more task-specific (less general), while transformer layers (BERT) do not exhibit this increase in task-specificity.

## 3 Our Approach

To probe the linguistic knowledge encoded by BERT and understand how it affects its predictions in several classification problems, we relied on a suite of 68 probing tasks, each of which corresponds to a distinct feature capturing lexical, morpho–syntactic and syntactic properties of a sentence. Specifically,

Level of Annotation	Linguistic Feature	Label
<b>Raw Text Properties (<i>RawText</i>)</b>		
Raw Text	Sentence Length	sent_length
	Word Length	char_per_tok
<b>Vocabulary Richness (<i>Vocabulary</i>)</b>		
Vocabulary	Type/Token Ratio for words and lemmas	trr_form, trr_lemma
<b>Morphosyntactic information (<i>POS</i>)</b>		
POS tagging	Distribution of UD and language-specific POS	upos_dist_*, xpos_dist_*
	Lexical density	lexical_density
<b>Inflectional morphology (<i>VerbInflection</i>)</b>		
	Inflectional morphology of lexical verbs and auxiliaries	xpos_VB-VBD-VBP-VBZ, aux_*
<b>Verbal Predicate Structure (<i>VerbPredicate</i>)</b>		
	Distribution of verbal heads and verbal roots	verbal_head_dist, verbal_root_perc
	Verb arity and distribution of verbs by arity	avg_verb_edges, verbal_arity_*
<b>Global and Local Parsed Tree Structures (<i>TreeStructure</i>)</b>		
Dependency Parsing	Depth of the whole syntactic tree	parse_depth
	Average length of dependency links and of the longest link	avg_links_len, max_links_len
	Average length of prepositional chains and distribution by depth	avg_prep_chain_len, prep_dist_*
	Clause length	avg_token_per_clause
	<b>Order of elements (<i>Order</i>)</b>	
	Relative order of subject and object	subj_pre, obj_post
<b>Syntactic Relations (<i>SyntacticDep</i>)</b>		
	Distribution of dependency relations	dep_dist_*
<b>Use of Subordination (<i>Subord</i>)</b>		
	Distribution of subordinate and principal clauses	principal_prop_dist, subordinate_prop_dist
	Average length of subordination chains and distribution by depth	avg_subord_chain_len, subordinate_dist_l
	Relative order of subordinate clauses	subordinate_post

Table 1: Linguistic Features used in the probing tasks.

we defined three sets of experiments. The first consisted in probing the linguistic information learned by a pre-trained version of BERT (BERT-base, cased) using gold sentences annotated according to the Universal Dependencies (UD) framework (Nivre et al., 2016). In particular, we defined a probing model that uses BERT contextual representations for each sentence of the dataset and predicts the actual value of a given linguistic feature across the internal layers. The second set of experiments consisted in investigating variations in the encoded linguistic information between the pre-trained model and 10 different fine-tuned ones obtained training BERT on as many Native Language Identification (NLI) binary tasks. To do so, we performed again all probing tasks using the 10 fine-tuned models. For the last set of experiments, we investigated how the linguistic competence contained in the models affects the ability of BERT to solve the NLI downstream tasks.

### 3.1 Data

We used two datasets: (i) the UD English treebank (version 2.4) for probing the linguistic information learned before and after a fine-tuning process; (ii) a dataset used for the NLI task, which is exploited both for fine-tuning BERT on the downstream task and for reproducing the probing tasks in the third set of experiments.

**UD dataset** It includes three UD English treebanks: UD\_English-ParTUT, a conversion of a multilingual parallel treebank consisting of a variety of text genres, including talks, legal texts and Wikipedia articles (Sanguinetti and Bosco, 2015); the Universal Dependencies version annotation from the GUM corpus (Zeldes, 2017); the English Web Treebank (EWT), a gold standard universal dependencies corpus for English (Silveira et al., 2014). Overall, the final dataset consists of 23,943 sentences.

**NLI dataset** We used the 2017 NLI shared task dataset, i.e. the TOEFL11 corpus (Blanchard et al., 2013). It contains test responses from 13,200 test takers (one essay and one spoken response transcription per test taker) and includes 11 native languages (L1s) with 1,200 test takers per L1. We selected only written essays and we created pairwise subsets of essays written by Italian L1 native speakers and essays for all the other languages. At the end of this process, we obtained 10 datasets of 2,400 documents (33,756 sentences in average): 1,200 for the Italian L1 speakers and 1,200 for each of the other L1s included in the TOEFL11 corpus.

### 3.2 Probing Tasks and Linguistic Features

Our experiments are based on the probing tasks approach defined in Conneau et al. (2018), which aims to capture linguistic information from the representations learned by a NLM. In our study, each probing task

Models	RawText	Vocabulary	POS	VerbInflection	VerbPredicate	TreeStructure	Order	SyntacticDep	Subord	All
BERT	0.68	0.78	0.68	0.72	0.60	0.78	0.72	0.69	0.71	0.69
Baseline	0.52	0.23	0.27	0.35	0.48	0.70	0.51	0.34	0.48	0.38

Table 2: BERT  $\rho$  scores (average between layers) for all the linguistic features (*All*) and for the 9 groups corresponding to different linguistic phenomena. Baseline scores are also reported.




Figure 1: BERT average layerwise  $\rho$  scores.

consists in predicting the value of a specific linguistic feature automatically extracted from the parsed sentences in the NLI and UD datasets. The set of features is based on the ones described in Brunato et al. (2020) which are acquired from raw, morpho-syntactic and syntactic levels of annotation and can be categorised in 9 groups corresponding to different linguistic phenomena. As shown in Table 1, these features model linguistic phenomena ranging from raw text ones, to morpho-syntactic information and inflectional properties of verbs, to more complex aspects of sentence structure modeling global and local properties of the whole parsed tree and of specific subtrees, such as the order of subjects and objects with respect to the verb, the distribution of UD syntactic relations, also including features referring to the use of subordination and to the structure of verbal predicates.

### 3.3 Models

**NLM** We relied on the pre-trained English version of BERT (BERT-base cased, 12 layers, 768 hidden units) for both the extraction of contextual embeddings and the fine-tuning process for the NLI downstream task. To obtain the embeddings representations for our sentence-level tasks we used for each of its 12 layers the activation of the first input token ( $[CLS]$ ), which somehow summarizes the information from the actual tokens, as suggested in Jawahar et al. (2019).

**Probing model** As mentioned above, each of our probing tasks consists in predicting the actual value of a given linguistic feature given the inner sentence representations learned by a NLM for each of its layers. Therefore, we used a linear Support Vector Regression (LinearSVR) as probing model.

## 4 Profiling BERT

Our first experiments investigated what kind of linguistic phenomena are encoded in a pre-trained version of BERT. To this end, for each of the 12 layers of the model (from input layer  $-12$  to output layer  $-1$ ), we firstly represented each sentence in the UD dataset using the corresponding sentence embeddings according to the criterion defined in Sec. 3.3. We then performed for each sentence representation our set of 68 probing tasks using the LinearSVR model. Since most of our probing features are strongly correlated with sentence length, we compared the probing model results with the ones obtained with a baseline computed by measuring the Spearman’s rank correlation coefficient ( $\rho$ ) between the length of the UD dataset sentences and the corresponding probing values. The evaluation is performed with a 5-fold cross validation and using Spearman correlation ( $\rho$ ) between predicted and gold labels as evaluation metric.

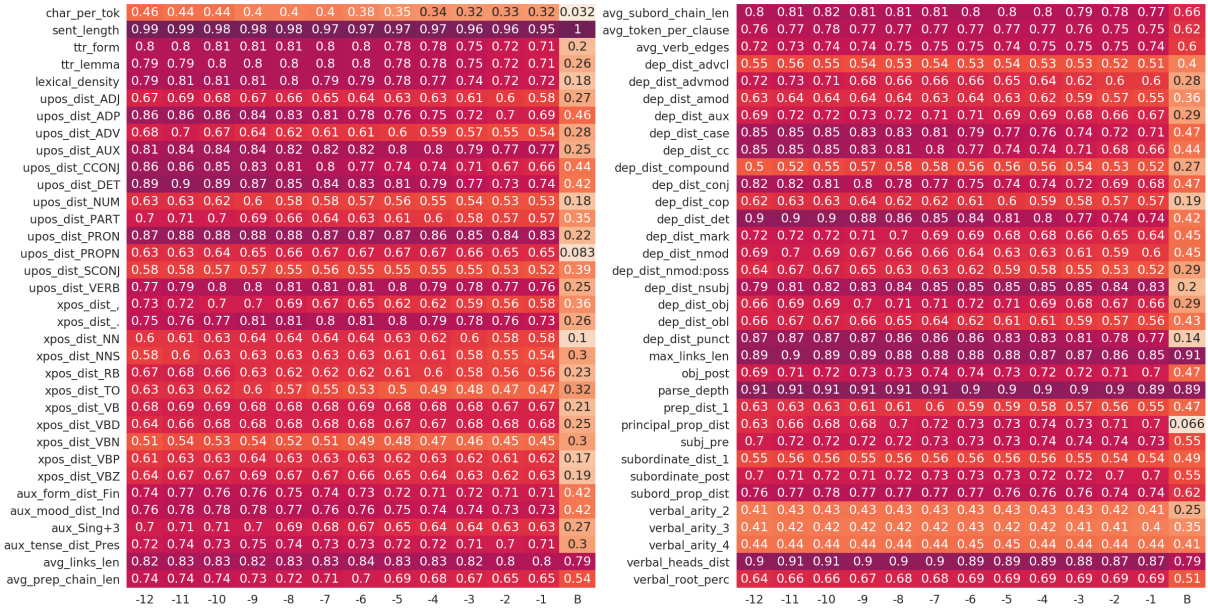


Figure 2: Layerwise  $\rho$  scores for the 68 linguistic features. Absolute baseline scores are reported in column B.

As a first analysis, we probed BERT’s linguistic competence with respect to the 9 groups of probing features. Table 2 reports BERT (average between layers) and baseline scores for all the linguistic features and for the 9 groups corresponding to different linguistic phenomena. As a general remark, we can notice that the scores obtained by BERT’s internal representations always outperform the ones obtained with the correlation baseline. For both BERT and the baseline, the best results are obtained for groups including features highly sensitive to sentence length. For instance, this is the case of syntactic features capturing global aspects of sentence structure (*Tree structure*). However, differently from the baseline, the abstract representations of BERT are also very good at predicting features related to other linguistic information such as morpho-syntactic (*POS*, *Verb inflection*) and syntactic one, e.g. the structure of verbal predicate and the order of nuclear sentence elements (*Order*).

We then focused on how BERT’s linguistic competence changes across layers. These results are reported in Figure 1, where we see that the average layerwise  $\rho$  scores are lower in the last layers both for all distinct groups and for all features together. As suggested by Liu et al. (2019), this could be due to the fact that the representations that are better-suited for language modeling (output layer) are also those that exhibit worse probing task performance, indicating that Transformer layers trade off between encoding general and probed features. However, there are differences between the considered groups: competences about raw texts features (*RawText*) and the distribution of POS are lost in the very first layers (by layer -10), while the knowledge about the order of subject/object with respect to the verb, the use of subordination, as well as features related to verbal predicate structure is acquired in the middle layers.

Interestingly, if we consider how the knowledge of each feature changes across layers (Figure 2), we observe that not all features belonging to the same group have an homogeneous behaviour. This is for example the case of the two features included in the *RawText* group: word length (*char\_per\_tok*) achieves quite lower scores across all layers with respect to the *sent\_length* feature. Similarly, the knowledge about POS differs when we consider more granular distinctions. For instance, within the broad categories of verbs and nouns, worse predictions are obtained by sub-specific classes of verbs based on tense, person and mood features (see especially past participle, *xpos\_dist\_VBN*), and by inflected nouns both singular and plural (*\_NN*, *\_NNS*). Within the broad set of features extracted from syntactic annotation, we also see that different scores are reported for features referring e.g. to types of dependency relations: those linking a functional POS to its head (e.g. *dep\_dist\_case*, *dep\_dist\_cc*, *dep\_dist\_conj*, *dep\_dist\_det*) are




Figure 3: Hierarchical clustering of the 68 probing tasks based on layerwise  $\rho$  values. Bold numbers correspond to the ranking of each probing feature based on the correlation with sentence length.

	<b>KOR</b>	<b>TEL</b>	<b>HIN</b>	<b>JPN</b>	<b>CHI</b>	<b>TUR</b>	<b>ARA</b>	<b>GER</b>	<b>FRE</b>	<b>SPA</b>
<b>Baseline</b>	59.05	51.32	54.09	56.27	55.68	55.66	52.92	59.29	56.03	52.61
<b>BERT</b>	85.74	85.18	84.75	84.19	82.78	79.29	76.38	72.78	72.50	70.03

Table 3: NLI classification results in terms of accuracy. We used the Zero Rule algorithm as baseline. Note that, for each task, sentences of the 10 languages are paired with the Italian ones (e.g. KOR = KOR-ITA).

better predicted than others relations, such as *dep\_dist\_amod*, *advcl*. Besides, within the *VerbPredicate* group, lower  $\rho$  scores are obtained by features encoding sub-categorization information about verbal predicates, such as the distribution of verbs by arity (*verbal\_arity\_2,3,4*), which also remains almost stable across layers.

Since we observed these not homogeneous scores within the groups we defined a priori, we investigated how BERT hierarchically encodes across layers all the features. To this end, we clustered the 68 linguistic characteristics according to layerwise probing results: specifically, we performed hierarchical clustering using Euclidean distance as distance metric and Ward variance minimization as clustering method. Interestingly enough, Figure 3 shows that the traditional division of features with respect to the linguistic annotation levels has not been maintained. On the contrary, BERT puts together features from all linguistic groups into clusters of different size. In addition, these clusters gather features that are differently ranked according to the baseline scores (ranking positions are bolded in the figure). For example, the first cluster includes features with similar  $\rho$  scores, and both highly and lower ranked by the baseline. All these features model aspects of global sentence structure, e.g. *sent\_length*, functional POSs (e.g. *upos\_dist\_DET*, *\_ADP*, *\_CCONJ*), parsed tree structures (e.g. *parse\_depth*, *verbal\_heads\_dist*, *avg\_links\_len*), nuclear elements of the sentence such as subjects (*dep\_dist\_nsubj*), verbs (*\_VERBS*), pronouns (*\_PRON*).

## 5 The Impact of Fine-Tuning on Linguistic Knowledge

Once we have probed the linguistic knowledge encoded by BERT across its layers, we investigated how it changes after a fine-tuning process. To do so, we started with the same pre-trained version of the model used in the previous experiment and performed a fine-tuning process for each of the 10 subsets






Figure 4: Layerwise mean  $\rho$  scores for the pre-trained and fine-tuned models.

char_per_tok	1.5	1.8	2.5	2.9	2.8	1.5	2.6	4.5 *	4.1	4.6 *	2.7	2.9	2.5	2.7	1.9	2.1	4.7	4.2	5.2	7.8	avg_subord_chain_len	2.7	2.9	2.5	2.7	1.9	2.1	4.7	4.2	5.2	7.8
sent_length	1.8 *	2.2	2.1	1.9	1.5	1.2	2.7 *	2.6 *	2.2	5.2 *	2.4	3.4	3.2	3.4	2.9	2.6	4.7	5.6	4.5	6.9	avg_token_per_clause	2.4	3.4	3.2	3.4	2.9	2.6	4.7	5.6	4.5	6.9
ttr_form	7.5	6.7	7.7	7.6	6.3	6	13	10	9.9	20	2.3	3.6	2.5	3	2.1	3.3	4.3	4.5	4.9	6.8	avg_verb_edges	1.9	2	1.9	1.9	0.9	2	3.3	3.7	6.2	7.2
ttr_lemma	7.2	6.5	7.2	7.7	5.9	5.7	12	9.2	9.4	19	dep_dist_advcl	5.1	5.6	3	3.7	0.82	3.1	4.2	4.7	4.8	6.3										
lexical_density	4.7	6.5	5.5	5.6	5.2	5.5	8.2	7.8	9.3	12	dep_dist_advmod	3.9	6.1	5	2.6	5.2	4.6	7.6	8.9	8.7	7.8 *										
upos_dist_ADJ	5	6.9	6.1	5	5.5	4.5	7.7	8.3	9.1	8.8 *	dep_dist_amod	7.2	7.2	3.5	7.8	4.2	4.6	3.1	7.5	6.4	11										
upos_dist_ADP	4.6	5.6	3	3.3	0.52	3.3	4.7	5.2	4.7	9.4	dep_dist_aux	2.2	5.4	4.7	1.3	2.4	3.6	7.5	8.1	7.4	10										
upos_dist_ADV	6.6	8.1	4	7.8	4.5	6.1	6.4	11	8.5	15	dep_dist_case	2.3	2.2	2.7	0.79	1.9	1.9	1.2	5.4	8.7	8.7										
upos_dist_AUX	2	2	2.9	0.61	1.8	1.8	1.1	5.5	8.6	8.6	dep_dist_cc	2.7	4.5	3.1	2	3.1	2.8	4.1	5.1	5.3	4.4 *										
upos_dist_CCONJ	0.72	5.3	5	1.1	2.6	4.2	9.5	9.2	7.5	8.4	dep_dist_compound	3.2	2.6	3.2	1.1	3.1	2.6	2.2	5.8	10	9.9										
upos_dist_DET	3.9	6.7	3.8	5.2	5	5.2	7.6 *	8.3	8.6 *	13 *	dep_dist_conj	5.8	8 *	4	6.4	4	4.8	7.6	11 *	9.1	16										
upos_dist_NUM	5.5	7	3.6	3	4.2	5.1	7.7	7.3	9.8	11	dep_dist_cop	0.59	5.5	5.2	0.89	2.7	4.2	9.7	9.7	7.5	8.3										
upos_dist_PART	3.3	3.9	4.3	4.5	4.2	4.6	7.1 *	5	5.7	5.3	dep_dist_det	3.1	4.4	2.8	2.6	1.1	2.6	5.4	4.1	6.4	9.6										
upos_dist_PRON	3.4	4.4	4.3	3.2	3	4	4.3 *	5.1 *	5.2	4.7 *	dep_dist_mark	1.2	4.5	3.8	1.1	1.8	2.4	6.1	6.2	5.3	6.8										
upos_dist_PROPN	2.3	2.4	2	1.7	-0.78	1.2	3.9	3.7	5.2	7.7	dep_dist_nmod	5.3	7.8	6.5	4.1	8.3	8	11 *	10	6.9	4.5										
upos_dist_VERB	4.9	6.6	3.9	4.8	3.8	4.5	8	7.6	8	12	dep_dist_nsubj	3	4.3	4.1	3.8	2.7	4.1	6.2	5.3	6.3	9.3										
xpos_dist_	2.7	5.2	4.8	3.5	1.8	1.1	8.6 *	6.8	4.8	15 *	dep_dist_obj	3.7	6.9	5.2	4.4	4.9	5	7.9	7.2	7.5	13										
xpos_dist_	8	6.4	7.9	6.8	5.6	7.9	11	11	9.5	21	dep_dist_obl	0.95	4	2	0.77	1.6	2	4.4	5.4	5.5	9.2										
xpos_dist_NN	5.8	10	7.4	5.1	7.2	7.1	11	12	12	11 *	dep_dist_punct	6.2	7.5	8.9	5.6	5.3	5.9	15	12	13	21										
xpos_dist_NNS	5.7	9.8	7.9	6.9	5.9	7.7	11	10	13 *	7.5 *	max_links_len	1.2	1.5	1.9	1.6	1.2	0.97	1.9 *	2.1	1.8	4.7 *										
xpos_dist_RB	4.8	4.7	2.6	2.9	0.87	2.6	3.9	4.5	4.9	7.3	obj_post	3.1	4.9	4.1	3.7	3.3	3.6	5.6	5.5	5.9	11										
xpos_dist_TO	4.4	6.4	3.6	1.9	3.1	3.9	7.6	6.3	8.1	11	parse_depth	1.4	2.5	2	1.6	1.1	1	3.3	3.1	3	5.6										
xpos_dist_VB	6.5	7.8	4	6.3	3.9	5.7	7.3	7.3	8.6	12 *	prep_dist_1	1.3	3.3	2.7	1	1.6	2.2	4.6	4.3	4.1	5.8										
xpos_dist_VBD	2.3	3.3 *	2.9 *	3.4	2	3 *	6.9 *	6.5 *	3.9 *	8.8 *	principal_prop_dist	2.6	3.8	4.2	6.1	2.9	3	7.8	7.9	7.2	16										
xpos_dist_VBN	4.6	8	2.7	5.2	2.7	4.7	5.2	8.4	6.3	9.9 *	subj_pre	1.6	2.3	2.9	2.8	1.8	2	2.9	2.9	2.5	6.5 *										
xpos_dist_VBP	5.5	6.2	3.8	8.4	4.2	5.2	7	8.3	5.8	11	subordinate_dist_1	1.2	1.3	1.4	1.4	1.2	0.74	2.5	2.3	2.8	4.5										
xpos_dist_VBZ	5.7 *	8 *	5.2 *	8 *	4.4 *	6.2 *	13 *	14 *	6.6 *	17 *	subordinate_post	3	2.7	3.1	3	1.2	2.3	4.5	4.3	5.5	9.6										
aux_form_dist_Fin	5	6.1	3.1	5.9	3.8	4.7	5.3	8	6.5	12	subord_prop_dist	2.1	2.2	2.4	2.6	1.5	1.5	4.5	4.4	4.9	8.3										
aux_mood_dist_Ind	5.7	7.3	3.1	7	4.3	3.8	6.4	9.1	7	15	verbal_arity_2	3	4.2	2.9	4.3	1.6	3	6.2	5.1	4.9	9.7										
aux_Sing+3	5.4	6.1	3.2	6.6	4.7	4.4	8.3	11	7.1	16 *	verbal_arity_3	1.6	1.5	2.2	2.5	1	2.2	1.9	2.4	2.3	3.6										
aux_tense_dist_Pres	5.1	6.4	4.6	7.2	4.3	3.9	11	11	5.5	14	verbal_arity_4	1.5	2.1	1.5	2	2.2	2.5	2.8	3.2	3.7	5.1										
avg_links_len	2.2	2.1	2.6	2.6	2.9	2.5	4	4	3.7	6.4	verbal_heads_dist	2.6	3.3	2.2	2.9	2	1.9	3.9	3.9	4.5	6.6										
avg_prep_chain_len	1.4	4.5	3.6	1.2	1.6	2.4	5.5	5.7	4.7	6.9	verbal_root_perc	0.32	0.73	0.84	1.5	0.83	0.26	0.83	1.7	0.93	5.2 *										
	KOR	TEL	HIN	JPN	CHI	TUR	ARA	GER	FRE	SPA		KOR	TEL	HIN	JPN	CHI	TUR	ARA	GER	FRE	SPA										

Figure 5: Differences between BERT-base and fine-tuned models  $\rho$  scores (multiplied by 100) computed using the output layer representations (-). Statistically significant variations (Wilcoxon Rank-sum test) are marked (\*).

built from the original NLI corpus (Sec. 3.1). We decided to use 50% of each NLI subset for training (40% and 10% for training and development set) and the remaining 50% for testing the accuracy of the newly generated models.

Table 3 reports the results for the 10 binary NLI tasks. As we can notice, BERT achieves good results for all downstream tasks, meaning that it is able to discriminate the L1 of a native speaker on a sentence-level regardless of the L1 pairs taken into account. The best performance is achieved by the model that was fine-tuned on the Korean and Italian pairwise subset, while the lowest scores are obtained with the model trained on the subset consisting of essays written by Spanish and Italian L1 speakers (SPA-ITA). Interestingly, these results seem to reflect typological distances among L1 pairs, with higher scores for languages that are more distant from Italian (Korean, Telugu or Hindi) and lower scores for L1s belonging to the same language family (FRE-ITA or SPA-ITA).

After fine-tuning the model on NLI, we performed again the suite of probing tasks on the UD dataset using the 10 newly generated models and following the same approach discussed in Section 4. Figure 4 reports layerwise mean  $\rho$  correlation values for all probing tasks obtained with BERT-base and the other fine-tuned models. It can be noticed that the representations learned by the NLM tend to lose their



Figure 6: % of probing features for which the MSE of the sentences correctly classified by BERT-base (*Pre-train*) and the fine-tuned models (*Fine-tune*) is lower than that of the incorrectly ones. Results are reported for layers -12, -7 and -1.

precision in encoding our set of linguistic features after the fine-tuning process. This is particularly noticeable at higher layers and it possibly suggests that the model is storing task-specific information at the expense of its ability to encode general knowledge about the language. Again, this is particularly evident for the models fine-tuned on the classification of language pairs belonging to the same family, SPA-ITA above all. To study which phenomena are mainly involved in this loss, we computed the differences between the probing tasks results obtained before and after the fine-tuning process. We focused in particular on the scores obtained on the output layer representations (layer -1), since it is the most task-specific (Kovaleva et al., 2019). For each subset, Figure 5 reports the difference between the score of each linguistic feature obtained with the pre-trained model and the fine-tuned one. Not surprisingly, the loss of linguistic knowledge reflects the typological trend observed for overall classification performance. In fact, when the task is to distinguish Italian vs German, French and Spanish L1, BERT loses much of its encoded knowledge for almost all the considered features. This is particularly evident for the morpho-syntactic features (i.e. distribution of *upos\_dist* and *xpos\_dist*) and for features related to lexical variety (i.e. *ttr\_form*, *ttr\_lemma*). It seems that for typologically similar languages BERT needs more task-specific knowledge mostly encoded at the level of morpho-syntactic information rather than the structural level. On the contrary, the drop is less pronounced and in most cases not significant for models fine-tuned on the classification of more distant languages (e.g. models fine-tuned on KOR-ITA or TUR-ITA). In this case, the quite stable performance on the probing tasks may suggest that those features were still useful to perform the downstream task. Interestingly, the class of features that decreases significantly in all models are those encoding the knowledge about the tense of verbs. This is particularly the case of the third-person singular verbs in the present tense (*xpos\_dist\_VBZ*) and of verbs in the past tense (*xpos\_dist\_VBD*). A possible explanation could be related to the prompts of essays, which are the same across the NLI dataset. Thus, the textual genre could have favored a quite homogeneous use of verbal morphology features by students of all L1s. This makes this class of features less useful for the identification of native languages.

## 6 Are Linguistic Features useful for BERT’s predictions?

As a last research question we investigated whether the implicit linguistic knowledge affects BERT’s predictions when solving the NLI downstream task. To answer this question we have split each NLI subset into two groups, i.e. sentences correctly classified according to the L1 and those incorrectly classified. For the two groups of each NLI subset, we performed the probing tasks using the pre-trained BERT-base and the specific NLI fine-tuned model. For each sentence of the two groups, we calculated the variation between the actual and predicted feature value obtaining two lists of absolute errors. We used the Wilcoxon Rank-sum test to verify whether the two lists were selected from samples with the same distribution. As a general remark, we observed that much more than half of features vary in a statistically significant way between correctly and incorrectly classified sentences. This suggests that BERT’s linguistic competence on the two groups of sentences is very different. To deepen the analysis of



this difference, we calculated the accuracy achieved by BERT in terms of Mean Square Error (MSE) only for the set of features varying in a significant way. Figure 6 reports the percentage of features for which the MSE of the sentences correctly classified (*MSE Pos*) is lower than that of the incorrectly ones (*MSE Neg*). This percentage is significantly higher, thus showing that BERT’s capacity to encode different kind of linguistic information could have an influence on its predictions: the more BERT stores readable linguistic information into the representations it creates, the higher will be its capacity of predicting the correct L1. Moreover, we noticed that this is true also (and especially) using the pre-trained model. In other words, this result suggests that the evaluation of the linguistic knowledge encoded in a pre-trained version of BERT on a specific input sequence could be an insightful indicator of its ability in analyzing that sentence with respect to a downstream task.

Interestingly, if we analyze the average length of correct and incorrect classified sentences, the correct ones are much more longer than the others for all tasks (from 3 tokens more for SPA-ITA to 9 for TEL-ITA). This is quite expected for the NLI task, since a higher number of linguistic events possibly occurring in longer sentences are needed to classify the L1 of a sentence (Dell’Orletta et al., 2014). At the same time, longer sentences make more complex the probing tasks because the output space is larger for almost all them. This is an additional evidence that BERT’s linguistic knowledge is not strictly related to sentence complexity, but rather to the model’s ability to solve a specific downstream task. To confirm this hypothesis and verify whether such tendency does not only depend on sentence length, we trained another LinearSVR that takes as input the sentence length and predict our probing tasks according to correctly or incorrectly classified NLI sentences. Table 4 reports the average Spearman’s correlation coefficients between gold and predict probing features for the two classes of sentences. Results showed that, for all the considered language pairs, the LinearSVR achieved higher accuracy for the probing tasks computed with respect to the incorrectly NLI classified sentences. This is an additional evidence that deeper linguistic knowledge is needed for BERT to correctly classify the L1 of a sentences.

Model	ARA	CHI	TUR	SPA	GER	FRE	JPN	KOR	TEL	HIN
Correct	0.226	0.225	0.236	0.223	0.215	0.224	0.276	0.239	0.234	0.229
Incorrect	0.248	0.251	0.249	0.235	0.244	0.239	0.290	0.255	0.258	0.257

Table 4: Average  $\rho$  scores for sentences correctly and incorrectly classified using only sentence length as input feature.

## 7 Conclusion

In this paper we studied what kind of linguistic properties are stored in the internal representations learned by BERT before and after a fine-tuning process and how this implicit knowledge correlates with the model predictions when it is trained on a specific downstream task. Using a suite of 68 probing tasks, we showed that the pre-trained version of BERT encodes a wide range of linguistic phenomena across its 12 layers, but the order in which probing features are stored in the internal representations does not necessarily reflect the traditional division with respect to the linguistic annotation levels. We also found that BERT tends to lose its precision in encoding our set of probing features after the fine-tuning process, probably because it is storing more task-related information for solving NLI. Finally, we showed that the implicit linguistic knowledge encoded by BERT positively affects its ability to solve the tested downstream tasks.

## 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*.
- Yonatan Belinkov and James Glass. 2019. Analysis methods in neural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Yonatan Belinkov, Lluís Màrquez, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. 2017. Evaluating layers of representation in neural machine translation on part-of-speech and semantic tagging tasks. In

*Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1–10.

Daniel Blanchard, Joel Tetreault, Derrick Higgins, Aoife Cahill, and Martin Chodorow. 2013. Toefl11: A corpus of non-native english. *ETS Research Report Series*, 2013(2):i–15.

Terra Blevins, Omer Levy, and Luke Zettlemoyer. 2018. Deep rnns encode soft hierarchical syntax. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 14–19.

Dominique Brunato, Andrea Cimino, Felice Dell’Orletta, Giulia Venturi, and Simonetta Montemagni. 2020. Profiling-ud: a tool for linguistic profiling of texts. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 7147–7153, Marseille, France, May. European Language Resources Association.

Andrea Cimino, Felice Dell’Orletta, Dominique Brunato, and Giulia Venturi. 2018. Sentences and documents in native language identification. In *CLiC-it*.

Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. What does BERT look at? an analysis of BERT’s attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy, August. 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. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136.

Felice Dell’Orletta, Martijn Wieling, Giulia Venturi, Andrea Cimino, and Simonetta Montemagni. 2014. Assessing the readability of sentences: Which corpora and features? In *Proceedings of the Ninth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 163–173, Baltimore, Maryland, June. Association for Computational Linguistics.

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.

Yoav Goldberg. 2019. Assessing bert’s syntactic abilities. *arXiv preprint arXiv:1901.05287*.

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.

John Hewitt and Christopher D Manning. 2019. A structural probe for finding syntax in word representations. 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 4129–4138.

Sarthak Jain and Byron C. Wallace. 2019. Attention is not Explanation. 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 3543–3556, Minneapolis, Minnesota, June. Association for Computational Linguistics.

Ganesh Jawahar, Benoît Sagot, Djamé Seddah, Samuel Unicomb, Gerardo Iñiguez, Márton Karsai, Yannick Léo, Márton Karsai, Carlos Sarraute, Éric Fleury, et al. 2019. What does bert learn about the structure of language? In *57th Annual Meeting of the Association for Computational Linguistics (ACL), Florence, Italy*.

Akos Kádár, Grzegorz Chrupała, and Afra Alishahi. 2017. Representation of linguistic form and function in recurrent neural networks. *Computational Linguistics*, 43(4):761–780.

Andrej Karpathy, Justin Johnson, and Li Fei-Fei. 2015. Visualizing and understanding recurrent networks. *arXiv preprint arXiv:1506.02078*.

Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of BERT. 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 4365–4374, Hong Kong, China, November. Association for Computational Linguistics.

- Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 2016. Visualizing and understanding neural models in nlp. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 681–691.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. Open sesame: Getting inside BERT’s linguistic knowledge. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy, August. Association for Computational Linguistics.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. Linguistic knowledge and transferability of contextual representations. 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 1073–1094, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Shannon Lubetich and Kenji Sagae. 2014. Data-driven measurement of child language development with simple syntactic templates. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2151–2160.
- Shervin Malmasi, Keelan Evanini, Aoife Cahill, Joel Tetreault, Robert Pugh, Christopher Hamill, Diane Napolitano, and Yao Qian. 2017. A report on the 2017 native language identification shared task. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 62–75, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Rebecca Marvin and Tal Linzen. 2018. Targeted syntactic evaluation of language models. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Alessio Miaschi and Felice Dell’Orletta. 2020. Contextual and non-contextual word embeddings: an in-depth linguistic investigation. In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 110–119, Online, July. Association for Computational Linguistics.
- Alessio Miaschi, Sam Davidson, Dominique Brunato, Felice Dell’Orletta, Kenji Sagae, Claudia Helena Sanchez-Gutierrez, and Giulia Venturi. 2020. Tracking the evolution of written language competence in 12 spanish learners. In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*. Association for Computational Linguistics, July.
- 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. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1659–1666.
- Christian S Perone, Roberto Silveira, and Thomas S Paula. 2018. Evaluation of sentence embeddings in downstream and linguistic probing tasks. *arXiv preprint arXiv:1806.06259*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alessandro Raganato and Jörg Tiedemann. 2018. An analysis of encoder representations in transformer-based machine translation. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. Association for Computational Linguistics.
- Manuela Sanguinetti and Cristina Bosco. 2015. Parttut: The turin university parallel treebank. In *Harmonization and Development of Resources and Tools for Italian Natural Language Processing within the PARLI Project*, pages 51–69. Springer.
- Naomi Saphra and Adam Lopez. 2019. Understanding learning dynamics of language models with svcca. 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 3257–3267.
- Natalia Silveira, Timothy Dozat, Marie-Catherine De Marneffe, Samuel R Bowman, Miriam Connor, John Bauer, and Christopher D Manning. 2014. A gold standard dependency corpus for english. In *LREC*, pages 2897–2904.

- Gongbo Tang, Rico Sennrich, and Joakim Nivre. 2018. An analysis of attention mechanisms: The case of word sense disambiguation in neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 26–35, Belgium, Brussels, October. Association for Computational Linguistics.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. BERT rediscovers the classical NLP pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy, July. Association for Computational Linguistics.
- 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. 2019b. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*.
- Hans van Halteren. 2004. Linguistic profiling for author recognition and verification. In *Proceedings of the Association for Computational Linguistics*, pages 200–207.
- Jesse Vig and Yonatan Belinkov. 2019. Analyzing the structure of attention in a transformer language model. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 63–76, Florence, Italy, August. Association for Computational Linguistics.
- Alex Warstadt, Yu Cao, Ioana Grosu, Wei Peng, Hagen Blix, Yining Nie, Anna Alsop, Shikha Bordia, Haokun Liu, Alicia Parrish, et al. 2019. Investigating bert’s knowledge of language: Five analysis methods with npis. 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 2870–2880.
- Amir Zeldes. 2017. The GUM corpus: Creating multilayer resources in the classroom. *Language Resources and Evaluation*, 51(3):581–612.
- Kelly Zhang and Samuel Bowman. 2018. Language modeling teaches you more than translation does: Lessons learned through auxiliary syntactic task analysis. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 359–361.