

Not All Models Localize Linguistic Knowledge in the Same Place: A Layer-wise Probing on BERToids' Representations

Mohsen Fayyaz^{1*} Ehsan Aghazadeh^{1*} Ali Modarressi² Hosein Mohebbi²
Mohammad Taher Pilehvar³

¹ University of Tehran, Iran ² Iran University of Science and Technology, Iran

³ Tehran Institute for Advanced Studies, Khatam University, Iran

{mohsen.fayyaz77, eaghazade1998}@ut.ac.ir

{m.modarressi, hosein.mohebbi}@comp.iust.ac.ir

mp792@cam.ac.uk

Abstract

Most of the recent works on probing representations have focused on BERT, with the presumption that the findings might be similar to the other models. In this work, we extend the probing studies to two other models in the family, namely ELECTRA and XLNet, showing that variations in the pre-training objectives or architectural choices can result in different behaviors in encoding linguistic information in the representations. Most notably, we observe that ELECTRA tends to encode linguistic knowledge in the deeper layers, whereas XLNet instead concentrates that in the earlier layers. Also, the former model undergoes a slight change during fine-tuning, whereas the latter experiences significant adjustments. Moreover, we show that drawing conclusions based on the *weight mixing* evaluation strategy—which is widely used in the context of layer-wise probing—can be misleading given the norm disparity of the representations across different layers. Instead, we adopt an alternative information-theoretic probing with *minimum description length*, which has recently been proven to provide more reliable and informative results.

1 Introduction

With the impressive success of pre-trained language models, such as BERT (Devlin et al., 2019), and their significant advances in transfer learning, a wave of interest has recently been directed toward understanding the knowledge encoded in their representations (Rogers et al., 2020). One of the analytical tools which is widely used for this investigation is *probing*: training a shallow supervised classifier that attempts to predict specific linguistic properties or reasoning abilities, based on representations obtained from the model (Tenney et al.,

2019b,a; Hewitt and Manning, 2019; Talmor et al., 2020; Mohebbi et al., 2021; Ushio et al., 2021; Chen et al., 2021).

However, most of the previous studies have focused on BERT only, neglecting other models in the family. This leaves open the question of how training objectives (which are fundamentally different for some models) and architectural choices would impact the resulting representations and the knowledge encoded in them.

In this work, we carry out an analysis on three popular language models with totally different pre-training objectives: BERT (masked language modeling), XLNet (permuted language modeling, Yang et al., 2019), and ELECTRA (replaced token detection, Clark et al., 2020). We also show that the “weight mixing” evaluation strategy of Tenney et al. (2019a), which is widely used in the context of probing (de Vries et al., 2020; Kuznetsov and Gurevych, 2020; Choenni and Shutova, 2020, *inter alia*), might not be a reliable basis for drawing conclusions in the layer-wise cross model analysis as it does not take into account the norm disparity across the representations of different layers. Instead, we perform an information-theoretic probing analysis using Minimum Description Length proposed by Voita and Titov (2020).

Based on a series of experiments, we find that language models derived from BERT have different behaviors in encoding linguistic knowledge. Specifically, we show that, unlike BERT, XLNet encodes linguistic information in the earlier layers during pre-training, while ELECTRA tends to carry this information to the higher layers. We also extend our probing experiments to the fine-tuned setting to assess the extent of change in the encoded knowledge upon fine-tuning. Using Representation Similarity Analysis (Kriegeskorte et al., 2008, RSA), we show that representations from the higher layers in XLNet—that do not well encode

*Equal contribution.

specific linguistic knowledge—undergo substantially bigger changes during fine-tuning when compared to the other models. In summary, our main contributions are as follows:

- We point out that the weight mixing evaluation strategy in edge probing does not lead to reliable conclusions in layer-wise cross model analysis studies.
- By relying on an information-theoretic probing method, we carry out a probing analysis across three commonly used pre-trained models.
- We also extend our probing experiments to fine-tuned representations to examine how linguistic information changes during fine-tuning.
- To provide complementary results to validate our findings, we also employ RSA to measure the amount of change in the representations after fine-tuning.

2 Background and Pilot Analysis

In this section, we review BERT and two of its popular derivatives: XLNet and ELECTRA, highlighting their differences as well as two commonly used probing methods that are essential to our discussion. We then conduct a pilot experiment to show the limitations of the evaluation metric used in edge probing and then justify our choice of an information-theoretic alternative.

2.1 Models

There is a wide variety of models derived from BERT, which are generally categorized into *autoencoding* and *autoregressive* models. Here, we focus on the most prominent model in each category. The two models have totally different pre-training objectives and have both shown outstanding performance on standard NLP tasks. BERT¹ is also considered as our baseline.

BERT. BERT (Devlin et al., 2019) has multiple Transformer encoder layers (Vaswani et al., 2017) stacked on top of each other, which are pre-trained with two self-supervised training objectives; Masked Language Model (MLM) and Next Sentence Prediction (NSP). The former predicts randomly masked tokens in the input sentence,

¹Due to resource limitations, our evaluations are based on the *base* version (12-layer, 768-hidden size, 12-attention head, 110M parameters) of each model obtained from the HuggingFace’s Transformers library (Wolf et al., 2020).

whereas the latter checks whether two sentences could be considered consecutive.

XLNet. In contrast to BERT, which attempts to reconstruct the original sentence from corrupted input, XLNet is an auto-regressive model based on the Transformer-XL architecture (Dai et al., 2019) that leverages permutation language modeling to learn from a bi-directional context. This allows the model to consider dependencies between masked tokens in a sentence for prediction. Due to the permuted order of the tokens in this objective, the next predicted token could occur at any position, making it a more difficult task. To address this, XLNet uses a two-stream self-attention mechanism (one for content and one for the query) to autoregressively predict tokens. XLNet has also excluded NSP during pre-training. Though these modifications resulted in better performance than BERT, the costs of pre-training have increased in terms of FLOPs (Clark et al., 2020).

ELECTRA. Another model based on BERT is ELECTRA which falls into the category of autoencoder models. Clark et al. (2020) have introduced the replaced token detection pre-training objective to substitute BERT’s objectives. ELECTRA jointly trains two models: generator and discriminator. The generator receives corrupted inputs with masked tokens in the input and tries to reconstruct the input like BERT. The discriminator is trained to predict whether each token was replaced by the generator or not. The authors pointed out that this novel pre-training objective could lead to learning representations that outperform prior models while requiring less computational cost during training.

2.2 Edge Probing

Tenney et al. (2019b) introduced edge probing as a means to measure linguistic knowledge in word representations. This is done by training a simple classifier on top of the representations. The accuracy of this classifier is taken as a representative for the quality of the encoded information about the specific task. Edge probing consists of a set of span-level tasks, where the span is a part of the sentence indicated by the dataset. The probe can only access the representations in that specific span. Eight labeling tasks are considered, including syntactic tasks, such as dependency labeling, and semantic tasks, such as coreference resolution. Before giving the inputs to the classifier, edge probing structure pools the representations across layers to

make a fixed size vector for feeding to the classifier. The work of Tenney et al. (2019a) is one of the first studies which leverages edge probing to quantify where linguistic knowledge is captured within BERT.

Scalar Mixing Weights. To estimate the contribution of each layer to a given probing task, Tenney et al. (2019a) used a technique called scalar mixing weights (Peters et al., 2018) which associates a trainable scalar weight with each layer in the model. After learning these weights alongside the probing classifier, they interpret layers with higher weights as those having more information for the particular task.

2.2.1 Mixing Weights Reliability Issues

While Tenney et al. (2019a) made interesting conclusions on BERT using edge probing and the *scalar mixing weights* evaluation strategy, we argue that this procedure is not reliable for layer-wise comparison. Several recent studies have conducted their experiments based on edge probing and *weight mixing* evaluation strategy. In one such study, Toshniwal et al. (2020) concluded that XLNet relies heavily on the input embedding layer in mixing weight evaluation for the coreference arc prediction task (Toshniwal et al., 2020). We show that this conclusion might not be accurate given that the representation norms in XLNet drastically change throughout layers.

Specifically, Figure 1 shows the representation norms across different layers in BERT and XLNet. These norms are computed based on the representations of 500 tokens sampled from the OPUS dataset (Tiedemann, 2012). The results are the average of three runs, and the same tokens are given to both models in each run². In XLNet, the norm of the embedding layer is extremely smaller than that of other layers. This clearly shows that the concentration of edge probing’s weight on the embedding layer does not indicate the level of information encoded in that layer. Rather, the model tries to compensate for relatively small representation norms. On the contrary, BERT retains the same level of representation norms across different layers. However, even such minor differences in representation norms might affect the conclusions of edge probing.³ Given this issue with edge probing, we opted

²If a word was split into more than one token by the tokenizer, we used the first token of that word.

³To mitigate this issue, one can normalize the representa-




Figure 1: Comparison of the representations norm in different layers of XLNet and BERT when tested on Wikipedia examples. XLNet shows considerable norm disparities across different layers.

for a theoretically justified method, Minimum Description Length (MDL) probing, in our layer-wise analysis.

2.3 MDL Probing

Conventional probes (such as Conneau et al., 2018; Tenney et al., 2019b; Jawahar et al., 2019) leave unclear whether the classifier identifies linguistic knowledge in the representations or learns the task itself (Hewitt and Liang, 2019). Hence, researchers had to limit the size of the dataset (Zhang and Bowman, 2018) or the probe’s complexity (Liu et al., 2019) to make sure the probe is not learning the task itself. However, probes based on information theory enable us to obtain more interpretable and reliable probing results.

The goal of information-theoretic probing is to measure to what extent representations encode a specific linguistic knowledge and how much effort is required to extract it. Voita and Titov (2020) combined the final quality of the probe classifier and the difficulty of achieving it by reformulating probes to a data transmission problem. If N number of representations are given, we plan to send their corresponding labels with a minimum description length, where each label has K classes. In *uniform encoding*, we assume that each representation has a label with a probability of $1/K$ and transmit them as raw information without training, which results in the maximum codelength possible of $N \cdot \log_2(K)$. But suppose the representations

tions just before applying the scalar mixing weights to obtain more reliable results, especially in XLNet. We reported results with this modification in the appendix (Figure A.1).

Dependencies	I think it will [help] ₂ [me] ₁ very much in my role . → obj (object)
NER	thirty eight years ago founded [the special Olympics] . → EVENT
SRL	Their father [called] ₁ later [to see if they were fine] ₂ . → ARG-M-PRP (Purpose)
Coreference	Thank [you] ₁ very much , [Tony] ₂ . → True
Rel. (SemEval)	NASA Kepler mission sends [names] ₁ into [space] ₂ . → Entity-Destination(e ₁ ,e ₂)

Table 1: Examples of sentences, spans, and target labels for each probing task.

show some degree of regularity with respect to the labels⁴. In that case, instead of sending the labels, we can train a classifier to predict the labels given the representations and transmit the classifier’s complexity (classifier codelength). Given that the classifier is usually not optimal, the final cross-entropy of the classifier over the data (data codelength) will be added to the classifier’s codelength, resulting in an alternative evaluation metric known as MDL. Also, since the number of targets N will affect the total sum of cross-entropy, and in turn the final codelength (MDL), it is preferred to use the compression evaluation metric, which is defined as:

$$c = \frac{N \cdot \log_2(K)}{\text{MDL}} \quad (1)$$

where the uniform codelength is divided by the MDL to eliminate the effect of N . Lower MDL indicates that the classifier predicts labels more accurately than a random guessing classifier, leading to higher compression. In contrast, when our classifier makes a random decision, MDL will be equal to the term in the numerator, resulting in no compression ($c = 1$). We will report compression instead of codelength in our experimental results.

To compute MDL, Voita and Titov (2020) proposed two methods. The first one is *variational coding*, which estimates the complexity of the probe using a Bayesian model. The second method is called *online coding*, which is based on the fact that if the regularity in data is strong, it can be revealed by only a small portion of the data. In this method, the probe will be progressively trained on different amounts of data from small to large portions, and the cross-entropy of each portion will be added together to form the final codelength. Since Voita and Titov (2020) showed that the two compression methods conform in results, we employ the *online coding* method due to its more straightforward implementation.

As the MDL probe is more stable and infor-

mative than other conventional probes (Voita and Titov, 2020), one can compare the codelength across layers of the same model or different models for a given probing task. The edge probing method does not allow this comparison since the mixing weights do not necessarily provide an accurate estimate of the richness of linguistic knowledge within each layer. In contrast, in MDL probing, each layer is probed separately, which gives us a direct estimate of the quality of the specific layer itself, rather than that relative to the other layers. MDL probing is, therefore, a better choice to have a layer-wise comparison among different models.

3 Probing Pre-trained Representations

We begin with probing the representations of each pre-trained model on a variety of core NLP tasks, including Dependency Labeling, Named Entity Recognition, Semantic Role Labeling, Coreference Resolution, and Relation Classification (see an example for each task in Table 1). Following Tenney et al. (2019b), we use OntoNotes 5.0 (Weischedel et al., 2011) for NER, SRL, and Coreference arc prediction, English Web Treebank portion of the Universal Dependencies (Silveira et al., 2014) for Dependencies, and SemEval 2010 Task 8 dataset (Hendrickx et al., 2010) for Relation classification. Statistics of the datasets are provided in Table A.1.

For probing evaluation, we perform MDL probing on frozen representations obtained from each model. In order to obtain a fixed-length representation, following Voita and Titov (2020), we project the representations to 256-dimensional vectors, and then apply a self-attention pooling on them.⁵

3.1 Results

For an overall cross-model comparison, we report the MDL probe compression and edge probing F1 score results in Table 2. For each model, the best MDL compression across different layers is reported. ELECTRA consistently achieves

⁴Highly regulated data would result in shorter codelength than weakly regulated ones (Voita and Titov, 2020).

⁵For tasks with two spans, we consider separate trainable weights for projection and attention pooling.

Task	BERT		XLNet		ELECTRA	
	F1 Score	Compression	F1 Score	Compression	F1 Score	Compression
Deps.	94.18	15.25	93.93	14.13	94.77	16.15
NER	95.61	16.87	95.51	15.46	96.07	16.88
SRL	90.91	13.94	90.56	13.32	91.69	14.44
Coref.	91.17	4.58	91.34	3.97	92.94	5.88
Rel.	80.63	3.04	82.07	2.97	82.41	3.37

Table 2: Cross-model MDL compression and edge probing micro-averaged F1 comparison. For each task we report the highest compression achieved among a model’s layers. Bold denotes the best performance on each task. In MDL probing, we employ logarithm with base 2 instead of natural logarithm in the training objective to have the codelength results in bits. Compression is the uniform codelength divided by the model’s codelength. F1 scores are obtained by scalar mixing weights similar to Tenney et al. (2019a) which are reported only for further comparison on **overall** probing performances.




Figure 2: MDL probing compression of BERT, XLNet, and ELECTRA across layers.

the highest compression and edge probing micro-averaged F1⁶ according to both metrics. Both results demonstrate how well the five tasks are en-

⁶In relation classification, we ignore the "Other" label for calculating the F1 score.

coded in the models’ representations during pre-training. Table 2 shows that ELECTRA can achieve the best quality in both MDL probing and edge probing compared to BERT and XLNet. Hence, ELECTRA seems to have the best pre-training objective for incorporating linguistic knowledge among the three models. On the other hand, XLNet displays comparable results to BERT, which is interesting given the relatively better fine-tuned performance of the former in a variety of downstream tasks.

Layer-wise analysis. Next, we use MDL probing to investigate how much linguistic knowledge is encoded in different layers in these models. Figure 2 shows layer-wise MDL probing compression results of BERT, XLNet, and ELECTRA on five probing tasks. Higher compression indicates better encoding of the task. As can be seen, ELECTRA attains the highest compression in different layers across most tasks, especially in the deeper layers. Notably, all models start with relatively low compressions and reach higher values in their middle layers. An interesting behavior shared among the three models is the decrease towards the final layer, which can be attributed to their pre-training objectives. The main difference between the models lies in the position in which the maximum amount of linguistic knowledge is accumulated.

To better demonstrate the layer that most captures each task, we compute the center of gravity following Tenney et al. (2019a). The only difference is that we apply it on MDL probing compression, instead of scalar mixing weights, defined as:

$$\bar{E}_{\mathbf{c}}[\ell] = \frac{\sum_{\ell=0}^L \ell \cdot \mathbf{c}^{(\ell)}}{\sum_{\ell=0}^L \mathbf{c}^{(\ell)}} \quad (2)$$




Figure 3: Comparison of the MDL probing compression center of gravity in BERT, XLNet, and ELECTRA.

where $c^{(\ell)}$ is the compression score of layer ℓ . Figure 3 shows the center of gravity of compression. The most noticeable distinction among models is that XLNet’s linguistic knowledge is concentrated in earlier layers than BERT, while ELECTRA’s knowledge is mostly accumulated in deeper layers.

We hypothesize that the difficulty of the objectives has a direct effect on \bar{E}_c , which indicates the expected position with the most encoded linguistic knowledge. In particular, recovering input tokens in the final layers of the model in the pre-training objective of BERT and XLNet is a surface task. Some of the linguistic knowledge might diminish in the final layers since highly contextualized representations have to be transformed into a less contextualized level to predict the original inputs (Voita et al., 2019). Whereas the pre-training objective in ELECTRA might be considered as a more semantic task, in which detecting replaced tokens requires more context-aware representations.

4 Probing Fine-tuned Representations

After observing that the amount and distribution of encoded linguistic knowledge can differ in these models, we aim to investigate how this information might be affected after the fine-tuning process. To this end, we repeated MDL probing as described in Section 3 on the fine-tuned representations. Next, to validate our results, we will expand our experiments to complementary analyses and measure the final quality of these representations




Figure 4: The change in centers of gravity after fine-tuning on MNLI dataset in five linguistic tasks.

on downstream tasks at different layers. We opted for the MNLI (Williams et al., 2018) dataset for fine-tuning all models. We used the same hyper-parameters for fine-tuning all three models: 32 as the batch size, max length of 128, the learning rate of $2e-5$, and five epochs of training.

Center of Gravity in fine-tuned models. Figure 4 demonstrates the difference in the average layer that most captures the information related to a particular task between the pre-trained and the fine-tuned models. We measure the difference of the two centers of gravity to evaluate the extent to which the concentration of knowledge (\bar{E}_c) shifts for each model on a specific probing task after fine-tuning:

$$\Delta \bar{E}_c = \bar{E}_c^{finetuned} - \bar{E}_c^{pretrained} \quad (3)$$

First, we show that the concentration of information in fine-tuned models is usually in earlier layers compared to the pre-trained models. This can be attributed to the significant loss of linguistic knowledge in the final layers of fine-tuned models in favor of the specific information of the fine-tuning task. We show that XLNet in most tasks falls back to earlier layers than the two other models because it forgets the most linguistic knowledge in the final layers. This suggests that XLNet is going through a more extensive change in its representations which we investigate in the following sections. We also indicate that ELECTRA is falling back more than BERT. We hypothesize that since ELECTRA focuses its most capable representations in the final layers, and those are the layers that change the most




Figure 5: Similarity of the representations of BERT, XLNet, and ELECTRA, before and after fine-tuning on MNLI dataset. Additional plots for CoLA and SST-2 are provided in the appendix (Figures A.3 and A.5).

in fine-tuning, its center of gravity changes more than BERT. Full MDL compression and codelength results are reported in the appendix (Table A.2).

Global RSA. After fine-tuning each model, we leverage Representational Similarity Analysis (RSA) to investigate the overall amount of changes in the representations of each layer. RSA is a technique borrowed from neuroscience (Kriegeskorte et al., 2008) which is used for comparing two different representation spaces. To be specific, we sampled 5000 English sentences from OPUS dataset (Tiedemann, 2012) and used Global RSA⁷, introduced by Chrupała et al. (2020). This is a better choice since we used average pooling in our fine-tuning, discussed next. As a measure of intra- and inter-space similarities, we used cosine similarity and Pearson correlation, respectively.

Figure 5 shows the results of the RSA measure applied to the representations of our models. We observe that XLNet has changed drastically during fine-tuning, while in BERT and ELECTRA, only the top layers are primarily affected. We also see that BERT shows a conservative pattern in the fine-tuning process which is consistent with findings of Merchant et al. (2020). As seen in Figure 2, we hypothesize that higher layers in XLNet are more open to change since they have relatively less specific knowledge. On the contrary, ELECTRA,

⁷Using the average pooled representations as stimuli instead of individual tokens.




Figure 6: Layer-wise comparison of the performance scores on MNLI dataset across the representations of the pre-trained and fine-tuned models.

which enjoys more linguistic information in its representations, especially in the higher layers, does not need to change very much.

Quality of the representations for downstream tasks. With the observations on RSA curves, we were interested in knowing the impact of the extent of these changes on downstream performance. To this end, we evaluated the quality of the representations for downstream tasks in both pre-trained and fine-tuned models. We trained separate classifiers on the unweighted average of representations for each layer.⁸ We used Adam optimizer with a learning rate of $5e-4$, and binary cross-entropy loss function for optimization.

Results are shown in Figure 6. Based on the performance scores for pre-trained representations, we observe that XLNet encodes most essential information for the downstream task in the shallower layers, BERT in the middle ones, and ELECTRA

⁸Unlike the other two models, BERT involves [CLS] token for NSP objective during pre-training step. Hence, to have a fair comparison between our fine-tuning and feature extraction experiments, we opted for the mean pooling strategy consistently throughout the paper.

in the deeper layers. Interestingly, these patterns are well aligned with the MDL probing curves in Figure 2 which indicates that pre-trained representations with more linguistic knowledge are more suitable for downstream tasks. We also report our results on CoLA (Warstadt et al., 2019) and SST-2 (Socher et al., 2013) datasets in Figures A.2 and A.4. Our results show that the observations are consistent across these downstream tasks. In addition, we show that XLNet significantly improves performance in its second half of layers, while ELECTRA undergoes smaller adjustments. We observe that, before fine-tuning, the last layers of XLNet have fairly similar or even lower performance than BERT. However, when fine-tuned, XLNet compensates for the performance deficit by injecting more task-specific information in those layers, helping the model to outperform BERT. Finally, we demonstrate that the changes in layers and their extent are similar to what we saw in the RSA results in Figure 5, which indicates that the changes in RSA were actually made to achieve higher quality in the fine-tuning task.

5 Related Work

While extensive research has been devoted to probing BERT (Hewitt and Manning, 2019; Tenney et al., 2019a; Merchant et al., 2020), other popular models, such as XLNet and ELECTRA which have significant differences in their pre-training objectives, are less thoroughly investigated. Only a few probing studies similar to our work exist which cover other models within the family. Mosbach et al. (2020) mostly focused on the impact of pooling strategy in probing sentence-level tasks. Based on the accuracy, they found that fine-tuning can affect the linguistic knowledge encoded in the representations. Moreover, Durrani et al. (2020) investigated the distribution of linguistic knowledge across individual neurons. In particular, they found that neurons in XLNet are more localized in encoding individual linguistic information compared to BERT, where neurons are shared across multiple properties. By adopting the method of Hewitt and Manning (2019), Aspillaga et al. (2021) investigated whether pre-trained language models encode semantic information, for instance by checking their representations against the lexico-semantic structure of WordNet (Miller, 1994).

The above studies mainly rely on the accuracy metric for their probing evaluation, which is re-

cently shown to fail in adequately reflecting the differences among representations (Voita and Titov, 2020). To our knowledge, this is the first time that an information-theoretic probing method is employed for conducting a cross-model and layer-wise analytical study.

6 Conclusions

In this paper, we aimed to extend probing studies on BERT to the other models in the family to investigate how training objectives and architectural choices would affect the resulting representations and the linguistic knowledge encoded in them. To this end, we leveraged MDL probing method, which has recently proven to provide more reliable and informative results when compared with conventional probes. To the best of our knowledge, this is the first time MDL probing has been employed to analyze such state-of-the-art pre-trained language models as BERT. By probing three state-of-the-art language models, i.e., BERT, XLNet, and ELECTRA, we found considerable differences in the extent and distribution of the core linguistic knowledge in their representations. Specifically, we demonstrate that XLNet accumulates linguistic knowledge in the earlier layers than BERT, whereas that of ELECTRA is mainly in the final layers.

Moreover, from probing and employing RSA similarity measure on fine-tuned models, we illustrate that XLNet is more susceptible to forgetting linguistic knowledge in final layers and undergoes substantial adjustments to its representations when compared to the other models. Based on differential downstream performance observations for before and after fine-tuning, we confirm that the changes in representations are proportional to the provided gain in the downstream task, which consequently indicates that XLNet injects more information during fine-tuning into its representations than the two other models.

In summary, through probing and measurement tools, we demonstrate that BERT’s derivative models, especially those with different objectives and structural choices, express different behaviors in their representations. We hope our analysis helps make more informed choices in the selection and fine-tuning of these state-of-the-art models.

Acknowledgments

Our work is in part supported by Tehran Institute for Advanced Studies (TeIAS), Khatam University.

References

- Carlos Aspillaga, Marcelo Mendoza, and Alvaro Soto. 2021. [Inspecting the concept knowledge graph encoded by modern language models](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2984–3000, Online.
- Boli Chen, Yao Fu, Guangwei Xu, Pengjun Xie, Chuanqi Tan, Mosha Chen, and Liping Jing. 2021. [Probing BERT in hyperbolic spaces](#). In *International Conference on Learning Representations*.
- Rochelle Choenni and Ekaterina Shutova. 2020. What does it mean to be language-agnostic? probing multilingual sentence encoders for typological properties. *arXiv preprint arXiv:2009.12862*.
- Grzegorz Chrupała, Bertrand Higy, and Afra Alishahi. 2020. [Analyzing analytical methods: The case of phonology in neural models of spoken language](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4146–4156, Online.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- Alexis Conneau, German 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, Melbourne, Australia. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota.
- Nadir Durrani, Hassan Sajjad, Fahim Dalvi, and Yonatan Belinkov. 2020. [Analyzing individual neurons in pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4865–4880, Online.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2010. [SemEval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 33–38, Uppsala, Sweden.
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China.
- 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, Minneapolis, Minnesota.
- 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.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter Bandettini. 2008. [Representational similarity analysis - connecting the branches of systems neuroscience](#). *Frontiers in Systems Neuroscience*, 2:4.
- Iliia Kuznetsov and Iryna Gurevych. 2020. [A matter of framing: The impact of linguistic formalism on probing results](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 171–182, Online.
- 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.
- Amil Merchant, Elahe Rahimtoroghi, Ellie Pavlick, and Ian Tenney. 2020. [What happens to BERT embeddings during fine-tuning?](#) In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 33–44, Online.

- George A. Miller. 1994. [WordNet: A lexical database for English](#). In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Hosein Mohebbi, Ali Modarressi, and Mohammad Taher Pilehvar. 2021. Exploring the role of BERT token representations to explain sentence probing results. *arXiv preprint arXiv:2104.01477*.
- Marius Mosbach, Anna Khokhlova, Michael A. Hedderich, and Dietrich Klakow. 2020. [On the interplay between fine-tuning and sentence-level probing for linguistic knowledge in pre-trained transformers](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 68–82, Online.
- 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, New Orleans, Louisiana.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in BERTology: What we know about how BERT works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Chris Manning. 2014. [A gold standard dependency corpus for English](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 2897–2904, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment tree-bank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. 2020. [olmpics-on what language model pre-training captures](#). *Transactions of the Association for Computational Linguistics*, 8:743–758.
- 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.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019b. [What do you learn from context? Probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*, pages 2214–2218, Istanbul, Turkey. European Languages Resources Association (ELRA).
- Shubham Toshniwal, Haoyue Shi, Bowen Shi, Lingyu Gao, Karen Livescu, and Kevin Gimpel. 2020. [A cross-task analysis of text span representations](#). In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 166–176, Online.
- Asahi Ushio, Luis Espinosa Anke, Steven Schockaert, and Jose Camacho-Collados. 2021. [BERT is to NLP what AlexNet is to CV: Can pre-trained language models identify analogies?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3609–3624, Online.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In *Advances in neural information processing systems*, pages 5998–6008.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. [The Bottom-up Evolution of Representations in the Transformer: A Study with Machine Translation and Language Modeling Objectives](#). 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 4387–4397.
- Elena Voita and Ivan Titov. 2020. [Information-theoretic probing with minimum description length](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 183–196, Online.
- Wietse de Vries, Andreas van Cranenburgh, and Malvina Nissim. 2020. [What’s so special about BERT’s layers? a closer look at the NLP pipeline in monolingual and multilingual models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4339–4350, Online.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Ralph Weischedel, Sameer Pradhan, Lance Ramshaw, Martha Palmer, Nianwen Xue, Mitchell Marcus, Ann Taylor, Craig Greenberg, Eduard Hovy, Robert Belvin, et al. 2011. *OntoNotes release 4.0. LDC2011T03, Philadelphia, Penn.: Linguistic Data Consortium*.

- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). 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 1112–1122, New Orleans, Louisiana.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.
- 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, Brussels, Belgium.

A Appendices

A.1 Edge Probing Normalized Mixing Weights Results




Figure A.1: Edge probing mixing weights results for BERT, XLNet, and ELECTRA. This is a modified version of mixing weights where we normalize the representations before applying mixing weights to eliminate the norms disparity effect.

A.2 Results for CoLA Dataset

Figures A.2 and A.3 report the performance and RSA results before and after fine-tuning on CoLA dataset. The results are consistent with figure 6.




Figure A.2: Layer-wise comparison of the performance scores on CoLA dataset across the representations of the pre-trained and fine-tuned models.




Figure A.3: Comparison of the representations of BERT, XLNet, and ELECTRA base, with their respective fine-tuned model on CoLA dataset.

A.3 Results for SST-2 Dataset

Figures A.4 and A.5 report the performance and RSA results before and after fine-tuning on SST-2 dataset.

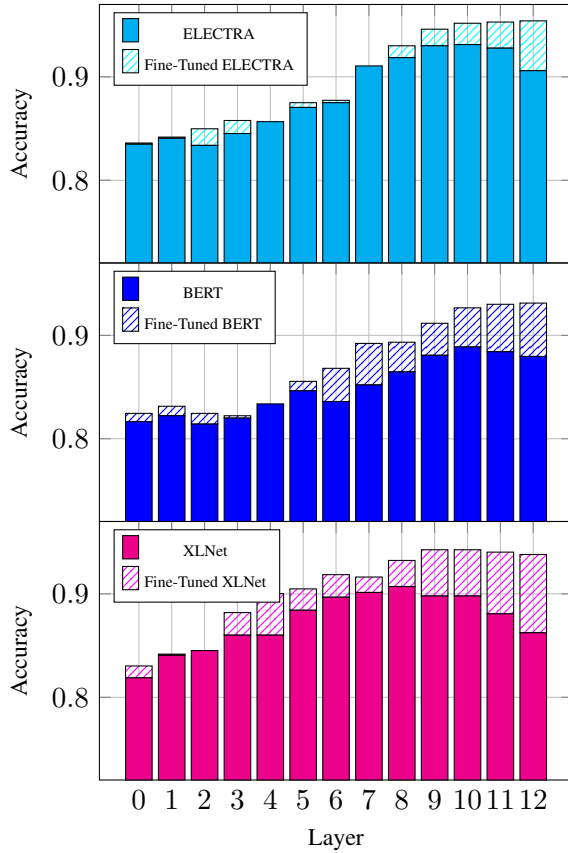


Figure A.4: Layer-wise comparison of the performance scores on SST-2 dataset across the representations of the pre-trained and fine-tuned models.




Figure A.5: Comparison of the representations of BERT, XLNet, and ELECTRA base, with their respective fine-tuned model on SST-2 dataset.

A.4 Datasets Statistics

The number of labels and targets of the NLP linguistic tasks we used in probing are reported in table A.1.

Task	Labels	Number of targets
Dependency Labeling	49	203919 / 25110 / 25049
Named Entity Recognition	18	128738 / 20354 / 12586
Semantic Role Labeling	66	598983 / 83362 / 61716
Coreference Resolution	2	207830 / 26333 / 27800
Relation Classification	19	6851 / 1149 / 2717

Table A.1: Dataset statistics for all five core tasks used in probing. Numbers of targets are given for train / dev / test sets.

Tasks	BERT		XLNet		ELECTRA	
	pre-trained	fine-tuned	pre-trained	fine-tuned	pre-trained	fine-tuned
Dependencies	5.99 (186.7)	6.04 (185.1)	5.01 (223.1)	5.03 (222.3)	6.35 (176.0)	6.33 (176.5)
	8.39 (133.3)	7.89 (141.7)	9.39 (119.1)	8.41 (132.9)	9.33 (119.8)	9.26 (120.8)
	9.49 (117.8)	8.92 (125.3)	11.37 (98.3)	10.28 (108.8)	10.71 (104.4)	10.70 (104.5)
	10.75 (104.0)	10.29 (108.7)	13.18 (84.8)	12.44 (89.8)	10.97 (101.9)	10.93 (102.3)
	12.75 (87.7)	12.12 (92.3)	13.98 (80.0)	13.14 (85.1)	11.75 (95.2)	11.55 (96.8)
	13.57 (82.4)	12.88 (86.8)	14.00 (79.9)	13.41 (83.4)	13.54 (82.5)	13.57 (82.4)
	14.45 (77.4)	13.45 (83.1)	14.13 (79.1)	13.54 (82.6)	14.42 (77.5)	14.36 (77.9)
	14.95 (74.8)	13.81 (81.0)	13.57 (82.4)	13.11 (85.3)	15.65 (71.5)	15.04 (74.3)
	15.25 (73.3)	13.61 (82.2)	13.30 (84.1)	12.19 (91.7)	16.15 (69.2)	15.15 (73.8)
	14.34 (78.0)	12.46 (89.7)	12.72 (87.9)	10.93 (102.3)	15.67 (71.3)	14.27 (78.3)
	13.17 (84.9)	11.60 (96.4)	11.63 (96.2)	9.32 (120.0)	15.76 (70.9)	14.03 (79.7)
	12.08 (92.5)	10.83 (103.3)	10.42 (107.3)	6.77 (165.1)	15.44 (72.4)	13.89 (80.5)
	11.06 (101.1)	9.76 (114.6)	7.26 (154.1)	3.08 (362.5)	13.62 (82.1)	11.98 (93.4)
Entities	8.37 (62.6)	8.28 (63.3)	9.33 (56.2)	9.41 (55.7)	8.72 (60.1)	8.63 (60.7)
	11.37 (46.1)	10.86 (48.3)	12.14 (43.2)	11.67 (44.9)	11.74 (44.6)	11.47 (45.7)
	12.19 (43.0)	11.77 (44.5)	13.56 (38.7)	13.57 (38.6)	13.55 (38.7)	13.24 (39.6)
	12.95 (40.5)	12.48 (42.0)	14.97 (35.0)	14.92 (35.1)	14.79 (35.4)	14.30 (36.7)
	14.15 (37.0)	13.65 (38.4)	15.46 (33.9)	15.79 (33.2)	15.25 (34.4)	15.02 (34.9)
	15.28 (34.3)	14.76 (35.5)	15.30 (34.3)	15.91 (33.0)	16.77 (31.3)	15.99 (32.8)
	15.50 (33.8)	14.92 (35.1)	14.85 (35.3)	15.63 (33.5)	16.88 (31.1)	15.79 (33.2)
	16.29 (32.2)	15.32 (34.2)	14.29 (36.7)	15.61 (33.6)	16.60 (31.6)	15.63 (33.5)
	16.77 (31.3)	15.58 (33.7)	14.20 (36.9)	15.13 (34.6)	16.76 (31.3)	15.44 (34.0)
	16.87 (31.1)	15.33 (34.2)	13.72 (38.2)	14.09 (37.2)	16.46 (31.9)	14.83 (35.3)
	16.59 (31.6)	14.72 (35.6)	13.36 (39.2)	12.92 (40.6)	16.29 (32.2)	14.22 (36.9)
	16.06 (32.7)	14.55 (36.0)	12.77 (41.0)	10.46 (50.1)	15.47 (33.9)	13.59 (38.6)
	15.64 (33.5)	13.60 (38.5)	10.84 (48.4)	5.38 (97.4)	14.10 (37.2)	11.21 (46.8)
SRL	7.93 (445.7)	7.87 (449.3)	6.91 (511.4)	6.93 (510.3)	8.07 (438.4)	8.03 (440.1)
	9.67 (365.7)	9.44 (374.7)	10.70 (330.5)	10.33 (342.2)	10.55 (335.0)	10.48 (337.4)
	10.42 (339.2)	10.26 (344.6)	11.90 (297.1)	11.68 (302.6)	11.64 (303.7)	11.51 (307.3)
	11.38 (310.8)	11.27 (313.8)	12.95 (273.0)	12.98 (272.5)	12.07 (292.9)	11.94 (296.0)
	12.45 (284.0)	12.40 (285.2)	13.32 (265.5)	13.39 (264.0)	12.77 (276.8)	12.59 (280.8)
	13.25 (266.8)	12.91 (273.8)	13.23 (267.3)	13.36 (264.6)	13.79 (256.4)	13.55 (261.0)
	13.73 (257.6)	13.26 (266.7)	13.05 (270.9)	13.25 (266.8)	14.13 (250.2)	13.91 (254.2)
	13.94 (253.7)	13.30 (265.9)	12.50 (282.8)	12.66 (279.3)	14.44 (244.8)	14.03 (252.1)
	13.88 (254.8)	13.06 (270.7)	12.05 (293.4)	12.13 (291.4)	14.32 (246.8)	13.59 (260.1)
	13.27 (266.5)	12.35 (286.2)	11.57 (305.7)	11.30 (312.8)	13.94 (253.6)	13.09 (270.1)
	12.54 (281.8)	11.64 (303.7)	11.06 (319.8)	10.02 (352.7)	13.78 (256.6)	12.79 (276.4)
	12.00 (294.6)	11.23 (314.8)	10.36 (341.2)	8.16 (433.3)	13.52 (261.4)	12.40 (285.2)
	11.46 (308.5)	10.58 (334.3)	8.40 (420.9)	4.60 (769.0)	12.47 (283.5)	11.32 (312.3)
Coreference	2.74 (74.2)	2.75 (73.8)	2.65 (76.6)	2.63 (77.3)	2.73 (74.4)	2.75 (73.8)
	2.96 (68.5)	2.94 (69.0)	3.03 (67.0)	2.91 (69.9)	3.06 (66.4)	3.06 (66.2)
	3.11 (65.2)	3.06 (66.4)	3.31 (61.3)	3.24 (62.6)	3.23 (62.9)	3.22 (63.0)
	3.29 (61.6)	3.31 (61.3)	3.63 (55.9)	3.63 (56.0)	3.33 (60.9)	3.35 (60.5)
	3.59 (56.5)	3.57 (56.8)	3.88 (52.3)	3.87 (52.4)	3.57 (56.8)	3.57 (56.9)
	3.76 (54.0)	3.76 (54.0)	3.92 (51.8)	4.00 (50.7)	4.04 (50.2)	3.93 (51.6)
	3.95 (51.4)	3.83 (53.0)	3.95 (51.4)	4.03 (50.3)	4.29 (47.3)	4.22 (48.0)
	4.14 (49.1)	3.98 (51.0)	3.94 (51.5)	4.07 (49.9)	4.56 (44.5)	4.43 (45.8)
	4.43 (45.8)	4.18 (48.5)	3.93 (51.6)	4.04 (50.3)	5.06 (40.1)	4.78 (42.4)
	4.58 (44.3)	4.13 (49.2)	3.97 (51.1)	3.95 (51.3)	5.44 (37.3)	4.93 (41.2)
	4.49 (45.2)	4.00 (50.7)	3.78 (53.7)	3.54 (57.3)	5.88 (34.5)	5.09 (39.9)
	4.32 (47.0)	3.85 (52.7)	3.50 (58.0)	3.11 (65.2)	5.68 (35.7)	4.95 (41.0)
	4.09 (49.6)	3.67 (55.3)	2.82 (72.1)	2.01 (101.1)	4.76 (42.7)	4.26 (47.6)
Relations	1.51 (18.9)	1.52 (18.7)	1.66 (17.1)	1.65 (17.2)	1.63 (17.4)	1.61 (17.6)
	1.80 (15.8)	1.73 (16.4)	1.95 (14.5)	1.85 (15.4)	1.85 (15.3)	1.86 (15.3)
	1.91 (14.9)	1.85 (15.4)	2.16 (13.1)	2.04 (13.9)	2.00 (14.2)	2.01 (14.1)
	2.12 (13.4)	2.02 (14.1)	2.51 (11.3)	2.43 (11.7)	2.18 (13.0)	2.15 (13.2)
	2.31 (12.3)	2.29 (12.4)	2.59 (11.0)	2.58 (11.0)	2.34 (12.2)	2.33 (12.2)
	2.48 (11.5)	2.42 (11.8)	2.68 (10.6)	2.67 (10.7)	2.80 (10.1)	2.69 (10.6)
	2.51 (11.3)	2.56 (11.1)	2.89 (9.8)	2.91 (9.8)	2.89 (9.8)	2.83 (10.0)
	2.74 (10.4)	2.76 (10.3)	2.90 (9.8)	3.00 (9.5)	3.04 (9.4)	2.94 (9.7)
	2.98 (9.5)	2.95 (9.6)	2.95 (9.6)	3.07 (9.3)	3.37 (8.4)	3.26 (8.7)
	3.04 (9.4)	3.03 (9.4)	2.95 (9.6)	2.91 (9.8)	3.34 (8.5)	3.19 (8.9)
	2.93 (9.7)	2.93 (9.7)	2.97 (9.6)	2.69 (10.6)	3.32 (8.6)	3.14 (9.0)
	2.77 (10.3)	2.86 (9.9)	2.82 (10.1)	2.36 (12.1)	2.98 (9.5)	2.83 (10.0)
	2.57 (11.1)	2.69 (10.6)	2.33 (12.2)	1.71 (16.6)	2.57 (11.0)	2.30 (12.3)

Table A.2: Cross-model MDL compression in pre-trained and fine-tuned models on MNLI dataset. The corresponding codelengths are presented in the brackets. Layers are 0 to 12 from top to bottom.