

How Does Fine-tuning Affect the Geometry of Embedding Space: A Case Study on Isotropy

Sara Rajae[◇] and Mohammad Taher Pilehvar[♣]

[◇] Iran University of Science and Technology, Tehran, Iran

[♣] Tehran Institute for Advanced Studies, Khatam University, Tehran, Iran

sara_rajae@comp.iust.ac.ir

mp792@cam.ac.uk

Abstract

It is widely accepted that fine-tuning pre-trained language models usually brings about performance improvements in downstream tasks. However, there are limited studies on the reasons behind this effectiveness, particularly from the viewpoint of structural changes in the embedding space. Trying to fill this gap, in this paper, we analyze the extent to which the isotropy of the embedding space changes after fine-tuning. We demonstrate that, even though isotropy is a desirable geometrical property, fine-tuning does not necessarily result in isotropy enhancements. Moreover, local structures in pre-trained contextual word representations (CWRs), such as those encoding token types or frequency, undergo a massive change during fine-tuning. Our experiments show dramatic growth in the number of elongated directions in the embedding space, which, in contrast to pre-trained CWRs, carry the essential linguistic knowledge in the fine-tuned embedding space, making existing isotropy enhancement methods ineffective.

1 Introduction

Recently, several studies have focused on the remarkable potential of pre-trained language models, such as BERT (Devlin et al., 2019), in capturing linguistic knowledge. They have shown that pre-trained representations are able to encode various linguistic properties (Tenney et al., 2019a; Talmor et al., 2020; Goodwin et al., 2020; Wu et al., 2020; Zhou and Srikumar, 2021; Chen et al., 2021; Tenney et al., 2019b), among others, syntactic, such as part of speech (Liu et al., 2019a) and dependency tree (Hewitt and Manning, 2019), and semantic, such as word senses (Reif et al., 2019) and semantic dependency (Wu et al., 2021).

Despite their significant potential, pre-trained representations suffer from important weaknesses. Frequency and gender bias are two well-known problems in CWRs. While the former hurts the

semantic expressiveness of embedding space, the latter reflects the unwanted social bias in training data (Li et al., 2020; Garg et al., 2018; Gonen and Goldberg, 2019). The representation degeneration problem is another issue that limits their linguistic capacity. Gao et al. (2019) showed that the weight tying trick (Inan et al., 2017) in the pre-training procedure is mainly responsible for the degeneration problem in the embedding space. In such a case, the embeddings occupy a narrow cone in the space (Ethayarajh, 2019). Several approaches have been proposed to improve the isotropy of pre-trained models, which in turn boosts the representation power and downstream performance of CWRs (Zhang et al., 2020; Wang et al., 2020). However, previous studies have mainly focused on the anisotropy of pre-trained language models. Here, we investigate the impact of fine-tuning on isotropy. Specifically, We try to answer the following questions:

- Can the improved performance achieved by fine-tuning pre-trained language models (LMs) be attributed to the increased isotropy of the embedding space?
- Does isotropy enhancement (using methods that null out dominant directions) have the same positive outcome for the fine-tuned models as it has for the pre-trained ones?
- How does the distribution of CWRs change upon fine-tuning?

To answer these questions, we consider the semantic textual similarity (STS) as the target task and leverage the metric proposed by Mu and Viswanath (2018) for measuring isotropy. The pre-trained BERT and RoBERTa (Liu et al., 2019b) underperform static embeddings on STS, while fine-tuning significantly boosts their performance, suggesting the considerable change that CWRs un-

dergo during fine-tuning (Reimers and Gurevych, 2019; Rajae and Pilehvar, 2021).

Our analysis on the fine-tuned embedding space of BERT and RoBERTa demonstrates that word representations are highly anisotropic across all layers. An evaluation specifically carried out on the [CLS] tokens approves a similar pattern but to a greater extent. Moreover, experimental results show fading of local clustered areas in pre-trained CWRs during fine-tuning, which could be a possible reason for the improved performance. Interestingly, the fine-tuning procedure can change the linguistic knowledge encoded in dominant directions of embedding space from unnecessary information to the essential knowledge for the target task such that eliminating them toward making isotropic space hurts the performance of contextual representations.

2 Related Work

Following the research line in understanding the reasons behind the outstanding performance of pre-trained language models and their capabilities, most recent investigations on fine-tuning have been done through probing tasks and by evaluating the encoded linguistic knowledge (Merchant et al., 2020; Mosbach et al., 2020; Talmor et al., 2020; Yu and Ettinger, 2021). These studies demonstrate that most changes in fine-tuning are applied to the upper layers, such that those layers encode task-specific knowledge, while lower layers are responsible for the core linguistic phenomenon (Durrani et al., 2021). Moreover, the results show that some linguistic information is surprisingly eliminated by this procedure (Mosbach et al., 2020). Studies on the multi-head attention structure suggest a similar trend in their patterns during fine-tuning; in higher layers, attention weights experience more significant changes (Hao et al., 2020). More detailed analysis on self-attention modules indicates that dense and value projection matrices have heavily been affected by fine-tuning (Radiya-Dixit and Wang, 2020). However, geometric analysis on the embedding space and changes applied to the structure of embeddings during fine-tuning are aspects that have not been properly understood. Furthermore, evaluating the fine-tuning effect on frequency bias in CWRs is another aspect that distinguishes our work from previous studies.

3 Methodology

3.1 Background

Fine-tuning is a straightforward yet quite effective process for taking advantage of the linguistic knowledge encoded in pre-trained models and for achieving high performance on different downstream tasks (Peters et al., 2019). The [CLS] embedding or other strategies in calculating sentence representations (e.g., max- or mean-pooling) can be considered as the input to the classifier layer, which is jointly trained with the parameters of the pre-trained model on a specific task (Devlin et al., 2019).

Isotropy is a geometrical assessment of the distribution of data points in a feature space, which is ideally uniform (Gao et al., 2019). An embedding space is considered isotropic if the word embeddings are not biased towards a specific direction (feature). In other words, in isotropic space, word embeddings are uniformly distributed in the space, leading to low correlation and near-zero cosine similarity for randomly sampled words.

Contextual embedding spaces are known to lack the desirable isotropy property (Rajae and Pilehvar, 2021; Ethayarajh, 2019). Gao et al. (2019) called the defect *the representation degeneration problem* and attributed it mainly to the weight tying trick (Press and Wolf, 2017) and the language modeling as the objective of the training. Under such a circumstance, random word embeddings are highly similar to one another while shaping a narrow cone in the space. Clearly, anisotropic distribution hurts the expressiveness of the embedding space, especially for semantic downstream tasks.

Cosine similarity-based metrics have usually been employed for assessing the isotropy of embedding spaces where a near-zero cosine similarity between random embeddings indicates isotropic distribution. However, Rajae and Pilehvar (2021) demonstrated that these metrics might not be reliable for calculating isotropy since, in some cases, the cosine similarity of random words is zero while their distribution is not uniform. Hence, we utilize another metric based on Principal Components (PCs).

As we mentioned before, anisotropic embedding spaces have unusual elongations toward different directions. Using the eigenvectors calculated in Principal Component Analysis (PCA) procedure, we can find the most elongated directions of the

space, which are the reason for anisotropic distribution. The distribution is more uniform and isotropic if the extent of elongation is similar across different directions (the most and the least elongated directions). With this in mind, [Mu and Viswanath \(2018\)](#) proposed a measurement to quantify the embedding space isotropy employing PCs as follows:

$$I(\mathcal{W}) = \frac{\min_{u \in U} F(u)}{\max_{u \in U} F(u)} \quad (1)$$

where U is the set of all eigenvectors of the word embedding matrix, and $F(u)$ is the following partition function:

$$F(u) = \sum_{i=1}^N e^{u^T w_i} \quad (2)$$

where N is the number of word embeddings and w_i is the i^{th} word embedding. [Arora et al. \(2016\)](#) demonstrated that for a perfectly isotropic embedding space, $F(u)$ could be approximated by a constant. The value of $I(\mathcal{W})$ is closer to one for the more isotropic embedding spaces.

3.2 Methodology

We study the changes applied to the embedding space by fine-tuning from the perspective of isotropy. In this regard, we take several approaches explained as follows.

Zero-mean. This method simply transfers all the embeddings to the center.

Clustering+ZM. Here, we first cluster embeddings and then separately make each cluster zero-mean ([Cai et al., 2021](#)).

These two approaches give us a precise picture of the extent of isotropy in the fine-tuned embedding space, globally and locally, since making zero-mean is a prerequisite for measuring isotropy ([Mu and Viswanath, 2018](#)).

Global app. This is a simple and effective post-processing algorithm for improving the isotropy of embedding space proposed by [Mu and Viswanath \(2018\)](#). In this method, after making embeddings zero-mean, a few top dominant directions calculated using PCA are being discarded.

Cluster-based app. Based on the clustered structure of pre-trained LMs ([Michael et al., 2020](#); [Reif et al., 2019](#)), this method can significantly improve the performance of contextual embedding spaces as well as their isotropy ([Rajae and Pilehvar, 2021](#)).

Here, we first cluster embeddings and then make each cluster zero-mean individually. At the last step, dominant directions are calculated in each cluster and discarded.

The last two approaches help us make the embedding space isotropic and potentially attain performance improvement. Moreover, they give us an insight into the changes of clustered structure of pre-trained models during fine-tuning.

3.3 Target Task

To analyze changes in a fine-tuned model, we choose Semantic Textual Similarity (STS) as the target task considering STS-Benchmark dataset ([Cer et al., 2017](#)). STS is a semantic regression task in which the model needs to determine the similarity of two sentences in a paired sample. The label is a continuous range in 0 to 5.

The interesting point about STS, which makes it a reasonable choice for our analyses, is that the performance of pre-trained LMs is drastically low on this task ([Reimers and Gurevych, 2019](#)). In fact, BERT and RoBERTa’s contextual representations under-perform static embeddings, such as Glove ([Pennington et al., 2014](#)) in this task. Moreover, the [CLS] token, which is usually considered a sentence representation for classification tasks, has a lower performance than simple averaging over all tokens of a sentence. However, fine-tuning, whether with [CLS] token or mean-pooling method, can dramatically enhance the performance ([Reimers and Gurevych, 2019](#)).

3.4 Experimental Setup

We analyze the influence of fine-tuning on the embedding space of the base versions of BERT and RoBERTa. Both models have similar transformer-based architectures, while RoBERTa has been trained with more training data and a slight difference in the optimization procedure. For the pre-trained setting, we use the models as feature extractors (the weights are frozen in this phase). Applying the mean-pooling method over the word embeddings, we obtain a sentence representation for every sample and consider the cosine similarity of the sentence representations as the textual similarity score. In the fine-tuning scenario, we fine-tune the models with a Siamese architecture introduced by [Reimers and Gurevych \(2019\)](#) that uses cosine similarity and the mean-pooling method for sentence representation. In our experiments, the batch size is set to 32, the learning rate is set to 7E-5, and

	Baseline		Zero-mean		Clustering+ZM		Global		Cluster-based	
	Perf.	Isotropy	Perf.	Isotropy	Perf.	Isotropy	Perf.	Isotropy	Perf.	Isotropy
Pre-trained [†]	54.14	1.1E-5	59.70	1.1E-4	67.73	0.31	69.20	0.59	74.01	0.83
Fine-tuned [†]	84.41	4.1E-3	84.94	6.6E-3	80.10	0.11	82.14	0.22	64.43	0.60
Pre-trained [‡]	33.99	2.5E-6	37.66	8.3E-2	60.32	0.69	65.99	0.86	73.86	0.95
Fine-tuned [‡]	81.08	3.3E-4	81.34	6.1E-3	76.03	0.05	79.71	0.18	60.96	0.28

Table 1: Spearman correlation performance and isotropy for five different settings in the pre-trained and fine-tuned BERT[†] and RoBERTa[‡]. Unlike the pre-trained models, increased isotropy does not bring about improved performance for the fine-tuned models.

the models are fine-tuned for 3 epochs. Following our previous work (Rajae and Pilehvar, 2021), we set the number of clusters and discarded dominant directions in Global and Cluster-based approaches to 27 and 12, respectively, for both models.

4 Findings

The embedding space of fine-tuned models is still highly anisotropic. Figure 1 depicts our experimental results on evaluating the isotropy in the models’ embedding spaces using $I(\mathcal{W})$. We take the pre-trained embedding space as a baseline and compare its isotropy to the fine-tuned space (all representations) and the [CLS] tokens in all layers. The results demonstrate that performance enhancements achieved after fine-tuning cannot be attributed to the increased isotropy of the embedding space. Although fine-tuning improves isotropy, specifically in the upper layers, the distribution of embeddings is still highly non-uniform. Moreover, in most layers, the [CLS] tokens’ representations are much more anisotropic than all representations in the fine-tuned space. These patterns hold for both BERT and RoBERTa, while the latter tends to be more anisotropic. We also note that although different random seeds change the reported numbers, the difference between isotropy of [CLS], fine-tuned, and pre-trained embedding spaces remain.

Adjusting the fine-tuned embedding space for isotropy hurts its performance. Several studies have shown that isotropy has theoretical and practical benefits. A natural question that arises here is if increasing the isotropy of a fine-tuned embedding space would lead to further performance improvements? To examine this hypothesis, we fine-tuned the models with the Siamese architecture and considered the settings explained in Section 3.2. Results are listed in Table 1. Clearly, as opposed to

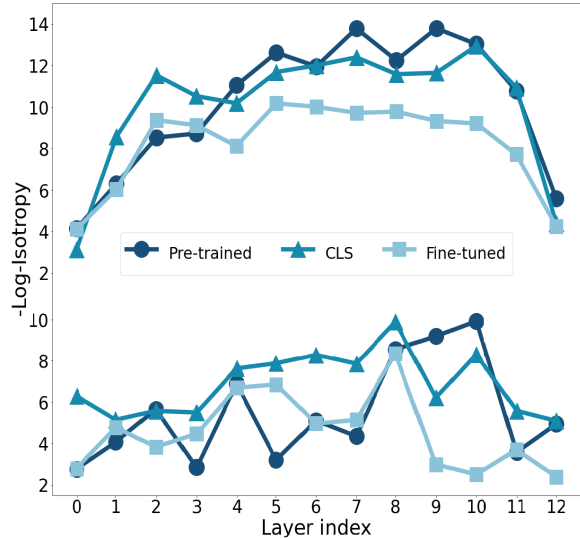


Figure 1: Negative \log of isotropy for [CLS] tokens, and all the tokens in the pre-trained and fine-tuned embedding space in all layers of BERT (bottom) and RoBERTa (top) using $I(\mathcal{W})$ on STS-B dev set. Higher values indicate lower isotropy.

pre-trained models, increasing isotropy of the fine-tuned embedding space does not enhance performance. Instead, we observe a drop in performance. This can be attributed to the fact that fine-tuning concentrates information about the target task in the dominant directions, whether it is obtained during the fine-tuning procedure or just brought up from the encoded knowledge in the pre-trained model.

The fine-tuned models heavily rely on a few top directions to solve the target task. To investigate the sensitivity of the fine-tuned model to the linguistic knowledge encoded in different directions, we discarded the least dominant directions and evaluated the performance of representations. The results of the experiment have been presented in Table 2. By eliminating the 100 and 700 least dominant directions from a total of 768 directions,

	Global App.				Cluster-based App.					
	Baseline		100 least dir.		700 least dir.		100 least dir.		700 least dir.	
	Perf.	Isotropy	Perf.	Isotropy	Perf.	Isotropy	Perf.	Isotropy	Perf.	Isotropy
BERT	84.41	4.1E-3	84.93	2.2E-3	82.93	2.2E-3	77.87	0.10	75.10	0.16
RoBERTa	81.08	3.3E-4	81.66	3.2E-4	78.59	1.4E-2	73.19	0.13	71.39	0.13

Table 2: Spearman correlation performance and isotropy after removing the least dominant directions in Global and Cluster-based approaches on STS dev set. The results suggest the low sensitivity of the fine-tuned models to eliminating more than 90% of directions with lower elongations.

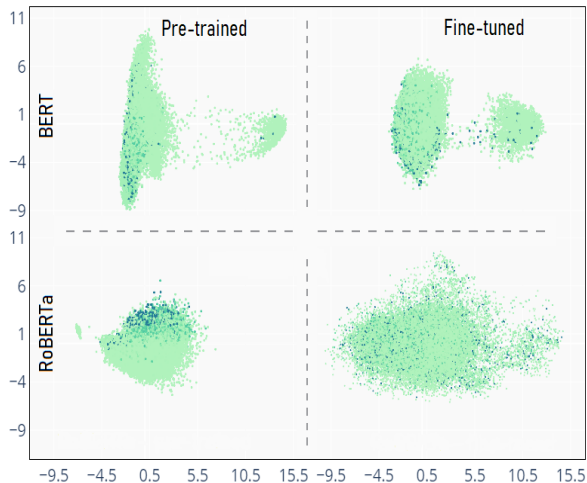


Figure 2: Illustration of pre-trained and fine-tuned CWRs colored based on their frequency in BERT and RoBERTa (using Wikipedia dump as corpus). The more frequent words have darker colors. As can be observed, the embedding space is still anisotropic after fine-tuning, while the frequency-based distribution of CWRs has been remedied.

we observe a slight drop in the performance compared to removing 12 top dominant directions. This suggests that the top dominant directions carry essential knowledge about the target task. We leave further investigation of this interesting behavior to future work.

The clustered structure of the embedding space changes during fine-tuning. The results of the Clustering+ZM setting and Cluster-based approach in Table 1 show that the clustered structure of the pre-trained embedding space (Cai et al., 2021) has faded in the fine-tuned CWRs. These two settings can improve the STS performance of the pre-trained model by increasing isotropy. However, applying them to fine-tuned CWRs leads to performance reduction. Moreover, as can be seen in Figure 2, the local areas that encode frequency

information in the pre-trained CWRs have been removed by fine-tuning, which can be a reason for the high performance of fine-tuned representations.

The number of elongated dominant directions significantly increases after fine-tuning. The results of Global and Cluster-based approaches in Table 1 reveal that with equal numbers of removed directions, the fine-tuned embedding space is less isotropic compared to the pre-trained one. This means that to have similar embedding spaces in terms of isotropy, we need to eliminate more dominant directions from the fine-tuned embedding space.

5 Conclusions

In this paper, we explored the effect of fine-tuning on the structure of the embedding space of BERT and RoBERTa. Our analysis demonstrates that the remarkable performance usually gained as a result of fine-tuning is not due to its enhancement of isotropy in the embedding space. Similarly to their pre-trained counterparts, fine-tuned CWRs have elongated directions towards different dimensions across all layers, and the number of these directions tends to increase by fine-tuning. We have also found that fine-tuning changes the nature of the linguistic knowledge encoded in dominant directions such that removing them hurts the performance (unlike pre-trained models for which removing such directions often result in performance improvements). Moreover, the clustered structure of pre-trained models is entirely modified upon fine-tuning, producing unbiased embedding space from the viewpoint of word frequency.

As future work, we plan to experiment with more target tasks and different fine-tuning strategies to expand our knowledge about the fine-tuning procedure. Furthermore, we aim at exploring the type of linguistic knowledge encoded in specific dimensions or subspaces in the semantic space.

References

- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. 2016. [A latent variable model approach to PMI-based word embeddings](#). *Transactions of the Association for Computational Linguistics*, 4:385–399.
- Xingyu Cai, Jiayi Huang, Yuchen Bian, and Kenneth Church. 2021. [Isotropy in the contextual embedding space: Clusters and manifolds](#). In *International Conference on Learning Representations*.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Boli Chen, Yao Fu, Guangwei Xu, Pengjun Xie, Chuanqi Tan, Moshu Chen, and Liping Jing. 2021. [Probing BERT in hyperbolic spaces](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nadir Durrani, Hassan Sajjad, and Fahim Dalvi. 2021. [How transfer learning impacts linguistic knowledge in deep NLP models?](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4947–4957, Online. Association for Computational Linguistics.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). 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 55–65, Hong Kong, China. Association for Computational Linguistics.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tiejun Liu. 2019. [Representation degeneration problem in training natural language generation models](#). In *International Conference on Learning Representations*.
- Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. [Word embeddings quantify 100 years of gender and ethnic stereotypes](#). *Proc. Natl. Acad. Sci. USA*, 115(16):E3635–E3644.
- Hila Gonen and Yoav Goldberg. 2019. [Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them](#). 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 609–614, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Goodwin, Koustuv Sinha, and Timothy J. O’Donnell. 2020. [Probing linguistic systematicity](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1958–1969, Online. Association for Computational Linguistics.
- Yaru Hao, Li Dong, Furu Wei, and Ke Xu. 2020. [Investigating learning dynamics of BERT fine-tuning](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 87–92, Suzhou, China. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. 2017. [Tying word vectors and word classifiers: A loss framework for language modeling](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. [On the sentence embeddings from pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9119–9130, Online. Association for Computational Linguistics.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019a. [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. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, M. Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). *ArXiv*, abs/1907.11692.

- 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. Association for Computational Linguistics.
- Julian Michael, Jan A. Botha, and Ian Tenney. 2020. [Asking without telling: Exploring latent ontologies in contextual representations.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6792–6812, Online. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Jiaqi Mu and Pramod Viswanath. 2018. [All-but-the-top: Simple and effective postprocessing for word representations.](#) In *International Conference on Learning Representations*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation.](#) In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. 2019. [To tune or not to tune? adapting pre-trained representations to diverse tasks.](#) In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 7–14, Florence, Italy. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models.](#) In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Evani Radiya-Dixit and Xin Wang. 2020. [How fine can fine-tuning be? learning efficient language models.](#) In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2435–2443. PMLR.
- Sara Rajaei and Mohammad Taher Pilehvar. 2021. [A cluster-based approach for improving isotropy in contextual embedding space.](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 575–584, Online. Association for Computational Linguistics.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B. Viegas, Andy Coenen, Adam Pearce, and Been Kim. 2019. [Visualizing and measuring the geometry of bert.](#) In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks.](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- 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 rediscovered the classical NLP pipeline.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. 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, and Ellie Pavlick. 2019b. [What do you learn from context? probing for sentence structure in contextualized word representations.](#) In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Lingxiao Wang, Jing Huang, Kevin Huang, Ziniu Hu, Guangtao Wang, and Quanquan Gu. 2020. [Improving neural language generation with spectrum control.](#) In *International Conference on Learning Representations*.
- Zhaofeng Wu, Hao Peng, and Noah A. Smith. 2021. [Infusing Finetuning with Semantic Dependencies.](#) *Transactions of the Association for Computational Linguistics*, 9:226–242.
- Zhiyong Wu, Yun Chen, Ben Kao, and Qun Liu. 2020. [Perturbed masking: Parameter-free probing for analyzing and interpreting BERT.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4166–4176, Online. Association for Computational Linguistics.
- Lang Yu and Allyson Ettinger. 2021. [On the interplay between fine-tuning and composition in transformers.](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2279–2293, Online. Association for Computational Linguistics.
- Zhong Zhang, Chongming Gao, Cong Xu, Rui Miao, Qinli Yang, and Junming Shao. 2020. [Revisiting representation degeneration problem in language](#)

[modeling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 518–527, Online. Association for Computational Linguistics.

Yichu Zhou and Vivek Srikumar. 2021. [DirectProbe: Studying representations without classifiers](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5070–5083, Online. Association for Computational Linguistics.