

# An Isotropy Analysis in the Multilingual BERT Embedding Space

Sara Rajae<sup>★</sup> and Mohammad Taher Pilehvar<sup>★</sup>

<sup>★</sup>Iran University of Science and Technology, Tehran, Iran

<sup>★</sup>Tehran Institute for Advanced Studies, Khatam University, Iran

sara\_rajae@comp.iust.ac.ir

mp792@cam.ac.uk

## Abstract

Several studies have explored various advantages of multilingual pre-trained models (such as multilingual BERT) in capturing shared linguistic knowledge. However, less attention has been paid to their limitations. In this paper, we investigate the multilingual BERT for two known issues of the monolingual models: anisotropic embedding space and outlier dimensions. We show that, unlike its monolingual counterpart, the multilingual BERT model exhibits no outlier dimension in its representations while it has a highly anisotropic space. There are a few dimensions in the monolingual BERT with high contributions to the anisotropic distribution. However, we observe no such dimensions in the multilingual BERT. Furthermore, our experimental results demonstrate that increasing the isotropy of multilingual space can significantly improve its representation power and performance, similarly to what had been observed for monolingual CWRs on semantic similarity tasks. Our analysis indicates that, despite having different degenerated directions, the embedding spaces in various languages tend to be partially similar with respect to their structures.<sup>1</sup>

## 1 Introduction

The multilingual BERT model (Devlin et al., 2019, mBERT), pre-trained on 104 languages with no supervision, has shown impressive ability in capturing linguistic knowledge across different languages (Pires et al., 2019). Many studies have explored the encoded knowledge in multilingual CWRs using probing tasks and under zero-shot setting (Wu and Dredze, 2019; K et al., 2020; Chi et al., 2020). Following the probing studies, in this paper, we investigate the multilingual embedding space of BERT, focusing on its geometry in terms of isotropy. Previous research has shown that many pre-trained

<sup>1</sup>Our code and datasets are publicly available at: <https://github.com/Sara-Rajae/Multilingual-Isotropy>

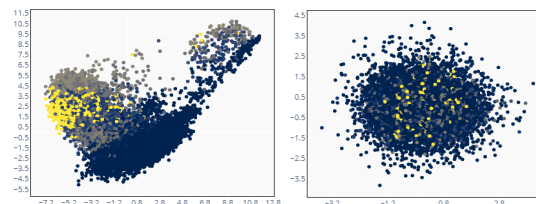


Figure 1: Degenerated (left) and isotropic (right) embedding spaces for Spanish plotted using PCA. Frequency-based distribution can be easily detected in the space (lighter colors indicate higher frequency). See Appendix C for more languages.

models, such as GPT-2 (Radford et al., 2019), BERT, and RoBERTa (Liu et al., 2019), have degenerated embedding spaces that downgrade their semantic expressiveness (Ethayarajh, 2019; Cai et al., 2021; Rajae and Pilehvar, 2021). Several proposals have been put forward to overcome this challenge (Gao et al., 2019; Zhang et al., 2020). However, to our knowledge, no study has so far been conducted on the degeneration problem in multilingual embedding spaces.

Using two well-known metrics, we evaluate isotropy in the mBERT embedding space for six different languages (including two low resources): English, Spanish, Arabic, Turkish, Sundanese, and Swahili. We find that the representation spaces are massively anisotropic in all these languages. However, unlike monolingual CWRs, where a few dimensions dominate the cosine similarity metric and have a high contribution to the anisotropic distribution (Timkey and van Schijndel, 2021), there is no dominant dimension that defines anisotropy in the multilingual representations. Extending our study to other structural properties of the multilingual space, we also investigate outliers, i.e., specific dimensions with consistently high values (Kovaleva et al., 2021). Our findings reveal that, as opposed to English BERT, the multilingual BERT space does not involve any major outliers. This indicates that the suggestion of Luo et al. (2021) on the role of

positional embeddings in the emergence of outliers may not be valid. Moreover, our analysis reveals that word frequency plays an important role in the distribution of the multilingual embedding space: words with similar frequencies create distinct local regions in the embedding space.

In analyzing multilingual space, we take a further step toward making the space isotropic. By applying a cluster-based isotropy enhancement method (Rajae and Pilehvar, 2021), we demonstrate that increasing isotropy of multilingual embedding space can result in significant performance improvements on semantic textual similarity tasks. Our frequency analysis and the remarkable performance improvement in the zero-shot setting denote that the feature space of mBERT has encoded a common linguistic knowledge into its dominant directions to some extent across different languages.

## 2 Background

The representation degeneration problem in LMs has attracted lots of attention in recent years. Several regularizer-based methods have been proposed to make the space isotropic by adding an extra constraint to the loss function during pre-training (Gao et al., 2019; Zhang et al., 2020; Wang et al., 2020). Because of the re-training cost, other light approaches have been presented as a post-processing step (Li et al., 2020; Rajae and Pilehvar, 2021). While analyzing the isotropy of embedding space is a well-studied area in English space, there are limited related studies on the multilingual embedding space. In this line, Vulić et al. (2020) investigated the structural similarity of different language embedding spaces by evaluating their isomorphism. Xu and Koehn (2021) showed the positive effect of isotropic space on the degree of isomorphism, which in turn results in improved performance in cross-lingual alignment algorithms. However, a focused study on the isotropy of multilingual embedding space has not been conducted. In this work, we provide more insights on the anisotropic distribution of multilingual embeddings and their notable differences from their English counterpart.

### 2.1 Isotropy

Geometrically, in an anisotropic space, embeddings occupy a narrow cone. This brings about an over-estimation of the similarity between embeddings (Gao et al., 2019). To quantify isotropy, we utilize two well-known metrics based on cosine similarity

	BERT		mBERT				
	En	En	Es	Ar	Tr	Su	Sw
$I_{Cos}(\mathcal{W})$	0.34	0.24	0.27	0.27	0.25	0.25	0.27
$I_{PC}(\mathcal{W})$	2.4E-5	6.4E-5	5.0E-5	1.6E-5	2.5E-4	1.2E-4	7.8E-5

Table 1: The isotropy of BERT and mBERT on a sub-set of Wikipedia, reporting based on  $I_{Cos}(\mathcal{W})$  and  $I_{PC}(\mathcal{W})$ .

and principal components (PCs).

**Cosine Similarity.** Ethayarajh (2019) used cosine similarity between random embeddings as an approximation of isotropy in the space. As mentioned before, random embeddings with an isotropic distribution have near-zero cosine similarities. The metric can be formulated as follows:

$$I_{Cos}(\mathcal{W}) = \frac{1}{N} \sum_{i=1, x_i \neq y_i}^N Cos(x_i, y_i) \quad (1)$$

where  $x_i \in X, y_i \in Y$ ,  $X$  and  $Y$  are the sets of randomly sampled embeddings, and  $\mathcal{W}$  is the embedding matrix.  $N$  is the number of sampled pairs that is set to 1000 in our experiments. Lower  $I_{Cos}(\mathcal{W})$  values indicate higher isotropy.

**Principal Components.** Mu and Viswanath (2018) proposed a metric based on principal components (PCs), approximated as follows:

$$I_{PC}(\mathcal{W}) \approx \frac{\min_{u \in U} F(u)}{\max_{u \in U} F(u)}, F(u) = \sum_{i=1}^M \exp(u^T w_i) \quad (2)$$

where  $w_i$  is the  $i^{th}$  word embedding,  $M$  is the number of all representations in the space,  $U$  is the set of eigenvectors of the embedding matrix, and  $F(u)$  is the partition function described in Equation 2. Arora et al. (2016) proved that  $F(u)$  could be approximated using a constant for isotropic embedding spaces. Therefore,  $I_{PC}(\mathcal{W})$  would be close to one in an isotropic embedding space.

## 3 Analysis

For all experiments, we opted for the multilingual BERT model (mBERT), which has a 12-layer transformer-based architecture similar to English BERT-base, and the representations are obtained from the last layer<sup>2</sup>. We selected English, Spanish,

<sup>2</sup>To broaden our insights on the geometry of multilingual space, we expand our analyses to XLM-R model. The related results are reported in Appendix A.

	$I_{Cos}(\mathcal{W})$	First	Second	Third
BERT	0.34	0.385	0.005	0.005
English	0.24	0.041	0.029	0.020
Spanish	0.27	0.033	0.029	0.018
Arabic	0.27	0.033	0.025	0.022
Turkish	0.25	0.036	0.024	0.024
Sundanese	0.25	0.036	0.016	0.016
Swahili	0.27	0.025	0.018	0.014

Table 2: The contribution of top-three dimensions to the expected cosine similarity ( $I_{Cos}(\mathcal{W})$ ) in BERT and mBERT models.

Arabic, Turkish, Sundanese, and Swahili. Our selection of these languages was to ensure that our analysis covers both high and low-resource languages. As our evaluation benchmark, we chose a subset of Wikipedia articles in the selected languages. The analysis experiments in Sections 3.1 to 3.3 have been conducted on the same dataset.

In what follows, we first assess isotropy as a desirable property in the multilingual space and evaluate the contribution of individual dimensions to this property. We also expand our study to the outliers and word-frequency bias in CWRs. The former is a weak point of language models, and the latter is a well-known bias in the monolingual embedding space. Lastly, we assess the effect of isotropy enhancement on the quality of the multilingual embeddings in the semantic similarity task.

### 3.1 Probing isotropy

As the first step, we quantify the isotropy of the mBERT and BERT embedding spaces using the two metrics. For mBERT, we separately assess the isotropy of each language in the embedding space.

Based on the presented results in Table 1, the average cosine similarity between random embeddings ( $I_{Cos}(\mathcal{W})$ ) is much higher than zero, denoting anisotropic distribution in all considered languages. Measuring isotropy using  $I_{PC}(\mathcal{W})$  also confirms the anisotropy issue in mBERT’s space as well as the monolingual BERT model.

Aligned with the numerical results, the illustration of multilingual CWRs in the left column of Figure 1 gives us a clear perspective of the degenerated distribution in space.

### 3.2 Sensitivity to Rogue Dimensions

As we discussed before, cosine similarity is a widely used metric to measure the degree of isotropy in embedding space where a near-zero

similarity demonstrates isotropic distribution. In this section, we evaluate the contribution of individual dimensions to the cosine similarity of two randomly chosen embeddings. Evaluating dimension-wise contribution sheds more light on the sensitivity of similarity-based metrics to individual dimensions and their role in (an)isotropic distribution.

The dimension-based cosine similarity between  $x$  and  $y$  vectors with  $d$  dimensions can be defined as follows (Timkey and van Schijndel, 2021):

$$Cos(x, y) = \sum_{i=1}^d \frac{x_i y_i}{\|x\| \|y\|} \quad (3)$$

where  $CC_i = x_i y_i / \|x\| \|y\|$  is the contribution of  $i^{th}$  dimension to the cosine similarity.

We compute the average cosine similarity,  $I_{Cos}(\mathcal{W})$ , by randomly sampling 1000 token pairs and report the average contribution of the top-three dimensions to the average cosine similarity.

The results are reported in Table 2. Aligned with the findings of Timkey and van Schijndel (2021), we observe that only one dimension has a considerable contribution to the cosine similarity metric in the BERT embedding space. Therefore, the anisotropic distribution is dominated by the mentioned dimension and is not a global property of the whole space. Unlike the monolingual BERT, multilingual BERT has no rogue dimensions. Hence, the anisotropic structure of the multilingual space cannot be attributed to certain dimensions.

### 3.3 Outlier Dimensions

Pre-trained LMs exhibit consistent outliers, peculiar dimensions with large values, in their Layer-Norms’ weights and consequently, in their contextual representations across all layers. Through several experiments, Kovaleva et al. (2021) have demonstrated that disabling these outliers can notably impair the performance of pre-trained and fine-tuned LMs. Trying to find the root cause of outliers, Luo et al. (2021) have shown that removing positional embeddings makes the outliers disappear, concluding that the positional information is responsible for the emergence of outliers.

In this part, we investigate rogue dimensions in the multilingual embedding space of BERT. We check outliers by averaging over 10000 randomly selected representations and calculate the mean and standard deviation ( $\sigma$ ) of dimensions’ distribution. Following the Kovaleva et al. (2021)’s suggestion, we consider a dimension as an outlier if its value

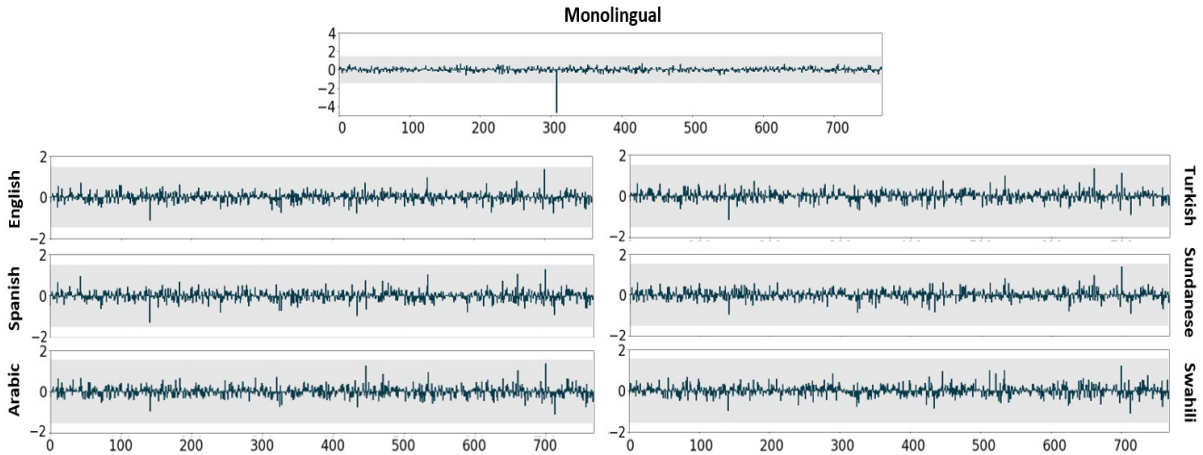


Figure 2: The average representation in English BERT (top) and mBERT (bottom). The shaded area denotes  $3\sigma$ . While an outlier has emerged in the former, we do not see any major outliers in the multilingual space.

is at least  $3\sigma$  larger/smaller than the mean of the distribution.

Results are shown in Figure 2. On top, the outlier dimension of the mean representation can be easily detected in the English BERT. However, interestingly, multilingual BERT exhibits no major outliers in its embedding space across different languages. It can be concluded that, contrary to the suggestion of Luo et al. (2021), positional embeddings cannot be responsible for outliers, given that both multi- and monolingual spaces are constructed using the same training procedure involving positional encodings. We leave further investigation of outliers in contextual embedding spaces to future work.

Putting together the results of the previous sections, we observe that the mBERT embedding space is highly anisotropic, despite not having any outliers or dominant dimensions in the cosine similarity metric.

### 3.4 Word frequency Bias

It has been shown that frequency plays an important role in the distribution of CWRs. Frequency-similar words make distinct local regions in the embedding space (Gao et al., 2019), with high-frequency and rare words being around the center and far from the origin, respectively (Li et al., 2020). Frequency-based distribution is a factor that hampers the expressiveness of the embedding space. So, it is essential to investigate frequency bias in the multilingual embedding space.

In this experiment, we analyze English, Spanish, and Arabic since there are enough resources to properly define word frequency in these languages. We randomly sample 500 sentences from the cor-

responding Wikipedia datasets and obtain a word representation by averaging over all its sub-token representations.

Figure 1 shows the distribution of word representations per word frequency<sup>3</sup>. Every point represents a word embedding dyed based on its frequency. As can be observed on the left, multilingual CWRs are biased toward their frequency, where words with similar frequencies create clustered regions. A similar pattern can be observed for the English BERT CWRs (Rajae and Pilehvar, 2021), with the only difference that in mBERT, low-frequency words are distributed near the origin and frequent words are far from it.

### 3.5 Isotropy Enhancement

Making the embedding space isotropic has theoretical and empirical benefits (Gao et al., 2019). Several approaches have been proposed to improve isotropy in monolingual CWRs. Some require a re-training of the model with additional objectives to address the degeneration problem (Gao et al., 2019; Zhang et al., 2020), whereas others are applied as a light post-processing (Mu and Viswanath, 2018). To investigate the effect of isotropy enhancement for the multilingual embedding space, we opted for our cluster-based approach (Rajae and Pilehvar, 2021), which is a recent example from the latter category. The proposed method splits the space into several clusters and discards dominant directions for each cluster. The approach also allows us to investigate the similarity of the clustered structure of the embedding space across different

<sup>3</sup>We used the wordfreq library (<https://pypi.org/project/wordfreq/>). See Appendix B.

	Ar-Ar	Ar-En	Es-Es	Es-En	Es-En-WMT	Tr-En	En-En
<b>Baseline</b>	51.76 (8E-5)	10.61 (1E-4)	64.15 (3E-5)	31.26 (5E-4)	11.39 (1E-4)	17.78 (1E-4)	60.82 (2E-6)
<b>Individual</b>	64.26 (0.60)	23.10 (0.57)	70.88 (0.54)	46.23 (0.50)	13.47 (0.50)	25.59 (0.55)	71.99 (0.54)
<b>Zero-shot</b>	52.76 (6E-5)	19.36 (0.04)	65.69 (8E-4)	43.82 (0.09)	13.68 (8E-3)	19.89 (0.03)	-

Table 3: STS performance (Spearman correlation percentage) on multi- and cross-lingual datasets using mBERT. Isotropy is reported based on  $I_{PC}(\mathcal{W})$  in parentheses. Applying the cluster-based method can improve the performance on the multi- and cross-lingual datasets in both Individual and Zero-shot settings.

languages under a zero-shot setting. More details on this method can be found in Appendix E.

We consider the multilingual and cross-lingual Semantic Textual Similarity (Cer et al., 2017, STS) that involves instances from Arabic, English, Spanish, and Turkish (Appendix D). We run our experiments in **Individual** and **Zero-shot** settings. In the former, we perform experiments individually on each language by clustering the corresponding space and applying the isotropy enhancement approach. The goal is to see whether increasing isotropy leads to performance improvement in the multilingual space and to compare the extent of improvements in cross- and multilingual tracks. In the zero-shot scenario, we are interested in evaluating the shared structural properties among languages, specifically, the similarity of the encoded linguistic knowledge in the dominant directions of different languages. To this end, we obtain clusters, their means, and dominant directions on the English dataset and leverage these for isotropy enhancement in other languages.

The reported results in Table 3 show that increasing the isotropy in the multilingual embedding space can enhance the performance in all tracks (multi- and cross-lingual). The improvement could be attributed to the potential of the applied method in discarding frequency bias from the embedding space. The visualization of the embedding space after isotropy enhancement, Figure 1 (right), clearly reveals that the frequency bias is faded after this process. Moreover, the results of the zero-shot setting suggest that the encoded information in dominant directions is similar across the languages because the improvement is compatible with the setting in which the dominant directions are obtained in each track individually.

## 4 Conclusion

In this paper, we provide several analyses on the geometry of multilingual embedding space from the viewpoint of isotropy. We show that, similarly

to its monolingual English counterpart, the multilingual BERT has a highly anisotropic embedding space. However, interestingly, the two spaces differ in their distribution of dimensions. The English BERT has a few high-contribution outlier dimensions, whereas the multilingual space does not possess any such disruptive rouge features. We also investigate the structure of multilingual embeddings from the perspective of frequency-based distribution and show that they have a biased structure towards word frequency and that the distribution is similar across different languages. As the last step, we evaluate the impact of isotropy improvement on the quality of multilingual embeddings. We observe that increasing isotropy can improve multilingual CWRs’ performance on STS and address their frequency bias.

## 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, Jiaji 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.
- Ethan A. Chi, John Hewitt, and Christopher D. Manning. 2020. [Finding universal grammatical relations in multilingual BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5564–5577, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco

- Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. 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, Minneapolis, Minnesota. 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*.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. [Cross-lingual ability of multilingual bert: An empirical study](#). In *International Conference on Learning Representations*.
- Olga Kovaleva, Saurabh Kulshreshtha, Anna Rogers, and Anna Rumshisky. 2021. [BERT busters: Outlier dimensions that disrupt transformers](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3392–3405, Online. Association for Computational Linguistics.
- 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.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Ziyang Luo, Artur Kulmizev, and Xiaoxi Mao. 2021. [Positional artefacts propagate through masked language model embeddings](#). 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 5312–5327, 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*.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- 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.
- William Timkey and Marten van Schijndel. 2021. [All bark and no bite: Rogue dimensions in transformer language models obscure representational quality](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4527–4546, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ivan Vulić, Sebastian Ruder, and Anders Søgaard. 2020. [Are all good word vector spaces isomorphic?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3178–3192, Online. Association for Computational Linguistics.
- 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*.
- Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness 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 833–844, Hong Kong, China. Association for Computational Linguistics.
- Haoran Xu and Philipp Koehn. 2021. [Cross-lingual BERT contextual embedding space mapping with isotropic and isometric conditions](#). *CoRR*, abs/2107.09186.

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.

## A XLM-R

XLM-R is a Transformer-based language model trained on 100 languages (Conneau et al., 2020). In comparison to mBERT, it has seen much more data during its pre-training, leading to higher performance on several downstream tasks. For our analysis, we follow the settings used for mBERT and take RoBERTa (Liu et al., 2019) as its monolingual counterpart.

### A.1 Results

**Probing Isotropy.** Table 4 summarizes the isotropy evaluation for RoBERTa and XLM-R models. As we expected, both models have an extremely anisotropic distribution. However, the degree of anisotropy is much higher in the multilingual model.

**Sensitivity to Rouge Dimensions.** To find out the contribution of individual dimensions to the high cosine similarity between random embeddings and anisotropic distribution, we report the contribution of top-three dimensions in Table 5.

In contrast to mBERT, a few dimensions have a significant role in the anisotropic distribution in the XLM-R model and RoBERTa. Furthermore, the results demonstrate that the contribution of rouge dimensions notably increases in the multilingual model.

**Outliers.** Following our settings in the BERT analysis (Section 3.3), we plot the average representations of the RoBERTa and XLM-R models to investigate the existence of outliers. The visualization of the mean representations across different languages and the monolingual model can be found in Figure 3. As we expected from the results of the previous part, outliers can easily be detected in the mono and multilingual models. The number of outliers is the same in both models. However, the absolute value of the outliers increases significantly in the multilingual models.

**Conclusion.** Our investigations on the geometry of embedding space in the RoBERTa and XLM-R models demonstrate a clear distinction to the BERT

counterpart. We show that although all the mentioned models have anisotropic embedding spaces, they possess different geometrical structures from the perspective of isotropy and rouge dimensions. Working on the reasons behind such discrepancies is an interesting future direction that can enhance our knowledge of how language models shape their underlying representation space.

## B Wordfreq

We have employed the Wordfreq library to investigate word frequency bias in our experiments. This library obtains word frequency from the corpus containing eight different domains in 36 languages. Our target languages are in the *large* category, which means their word lists cover rare words appearing at least once per 100 million words. As a result, the wordfreq could be a suitable tool for our purpose.

## C Frequency-based Distribution

Frequency-based distribution can negatively affect the expressiveness of space. Though it is a well-known bias in pre-trained LMs (e.g., BERT and GPT-2), it is not studied in a multilingual setting. As discussed in Section 3.4, we have studied frequency bias in mBERT and demonstrated that, similarly to its monolingual counterparts, mBERT suffers from frequency-based distribution in its space. The illustration of this bias and the impact of the cluster-based approach on mitigating the issue can be found in Figure 4.

## D Multilingual STS Task

Multi and cross-lingual Semantic Textual Similarity (STS) is the main task in our experiments. STS is a paired sentence task in which samples have been labeled by a score in the continuous range of 0 (irrelevant) to 5 (most semantic similarity). In the multilingual tracks, in a pair, both sentences are in the same language, while sentences have different languages in the cross-lingual tracks. The reason behind choosing STS as the target task for our experiments is that Multilingual BERT has a pretty low performance on it.

In our experiments, we take the average of all tokens in a sentence as the sentence representation and consider the cosine similarity of the sentence representations in a sample as the semantic similarity score.

	RoBERTa	XLM-R					
	En	En	Es	Ar	Tr	Su	Sw
$I_{Cos}(\mathcal{W})$	0.77	0.96	0.96	0.96	0.95	0.95	0.96
$I_{PC}(\mathcal{W})$	2.5E-6	3.8E-11	3.5E-9	3.4E-9	5.3E-9	5.9E-9	5.9E-9

Table 4: The isotropy of RoBERTa and XLM-R on a sub-set of Wikipedia, reporting based on  $I_{Cos}(\mathcal{W})$  and  $I_{PC}(\mathcal{W})$ .

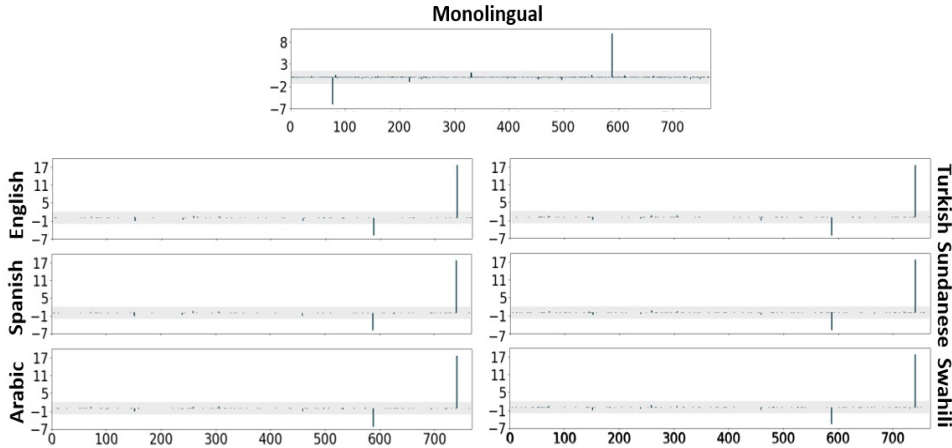


Figure 3: The average representation in RoBERTa (top) and XLM-R (bottom). The shaded area denotes  $3\sigma$ . Both models exhibit extreme outliers across their representations.

	$I_{Cos}(\mathcal{W})$	First	Second	Third
RoBERTa	0.77	0.703	0.251	0.007
English	0.96	0.895	0.098	0.000
Spanish	0.96	0.896	0.097	0.003
Arabic	0.96	0.879	0.108	0.003
Turkish	0.95	0.884	0.111	0.003
Sundanese	0.95	0.884	0.097	0.001
Swahili	0.96	0.897	0.088	0.001

Table 5: The contribution of top-three dimensions to the expected cosine similarity ( $I_{Cos}(\mathcal{W})$ ) in the RoBERTa and XLM-R embedding spaces.

## E Cluster-based Isotropy Enhancement

We pick the cluster-based approach (Rajae and Pilehvar, 2021) to improve the isotropy in multilingual embedding space. In this method, the embeddings are clustered using the k-means clustering algorithm, and then dominant directions of every cluster are nulled out independently. Dominant directions have been calculated using Principal Component Analysis (PCA). The primary key in this method is obtaining dominant principal components (PCs) of clustered areas in the embedding space separately, which makes this approach suitable for exploring the clustered structure of the multilingual CWRs. We apply the cluster-based

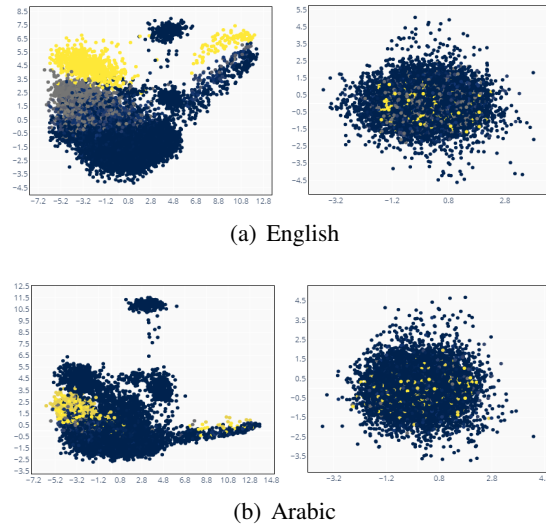


Figure 4: Degenerated (left) and isotropic (right) embedding spaces for the two languages. Frequency-based distribution can be easily detected using two top PCs in the space (lighter colors indicate higher frequency). Eliminating top dominant directions not only makes the embedding space isotropic, but also removes frequency bias in multilingual CWRs.

approach to multi and cross-lingual CWRs with two different settings, Individual and Zero-shot. The number of clusters and discarded dominant directions are 7 and 12, respectively.