

A Cluster-based Approach for Improving Isotropy in Contextual Embedding Space

Sara Rajaei¹ and Mohammad Taher Pilehvar^{1,2}

¹ Iran University of Science and Technology, Tehran, Iran

² Tehran Institute for Advanced Studies, Tehran, Iran

sara_rajaei@comp.iust.ac.ir

mp792@cam.ac.uk

Abstract

The representation degeneration problem in Contextual Word Representations (CWRs) hurts the expressiveness of the embedding space by forming an anisotropic cone where even unrelated words have excessively positive correlations. Existing techniques for tackling this issue require a learning process to re-train models with additional objectives and mostly employ a global assessment to study isotropy. Our quantitative analysis over isotropy shows that a local assessment could be more accurate due to the clustered structure of CWRs. Based on this observation, we propose a local cluster-based method to address the degeneration issue in contextual embedding spaces. We show that in clusters including punctuations and stop words, local dominant directions encode structural information, removing which can improve CWRs performance on semantic tasks. Moreover, we find that tense information in verb representations dominates sense semantics. We show that removing dominant directions of verb representations can transform the space to better suit semantic applications. Our experiments demonstrate that the proposed cluster-based method can mitigate the degeneration problem on multiple tasks.¹

1 Introduction

Despite their outstanding performance, CWRs are known to suffer from the so-called *representation degeneration problem* that makes the embedding space anisotropic (Gao et al., 2019). In an anisotropic embedding space, word vectors are distributed in a narrow cone, in which even unrelated words are deemed to have high cosine similarities. This undesirable property hampers the representativeness of the embedding space and limits the diversity of encoded knowledge (Ethayarajh, 2019).

¹The code for our experiments is available at https://github.com/Sara-Rajaei/clusterbased_isotropy_enhancement/

To better understand the representation degeneration problem in pre-trained models, we analyzed the embedding space of GPT-2 (Radford et al., 2019), BERT (Devlin et al., 2019), and RoBERTa (Liu et al., 2019). We found that, despite being extremely anisotropic in all non-input layers from a global sight, the embedding space is significantly more isotropic from a local point of view (when embeddings are clustered and each cluster is made zero-mean). Motivated by this observation and based on previous studies that highlight the clustered structure of CWRs (Reif et al., 2019; Michael et al., 2020), we extend the technique of Mu and Viswanath (2018) with a further clustering step. In our proposal, we cluster embeddings and apply PCA on individual clusters to find the corresponding principal components (PCs) which indicate the dominant directions for each specific cluster. Nulling out these PCs for each cluster renders a more isotropic space. We evaluated our cluster-based method on several tasks, including Semantic Textual Similarity (STS) and Word-in-Context (WiC). Experimental results indicate that our cluster-based method is effective in enhancing the isotropy of different CWRs, reflected by the significant performance improvements in multiple evaluation benchmarks.

In addition, we provide an analysis on the reasons behind the effectiveness of our cluster-based technique. The empirical results show that most clusters contain punctuation tokens, such as periods and commas. The PCs of these clusters encode structural information about context, such as sentence style; hence, removing them can improve CWRs performance on semantic tasks. A similar structure exists in other clusters containing stop words. The other important observation is about verb distribution in the contextual embedding space. Our experiments reveal that verb representations are separated across the tense dimension in distinct

sub-spaces. This brings about an unwanted peculiarity in the semantic space: representations for different senses of a verb tend to be closer to each other in the space than the representations for the same sense that are associated with different tenses of the same verb. Indeed, removing such PCs improves model’s ability in downstream tasks with dominant semantic flavor.

2 Isotropy in CWRs

Isotropy is a desirable property of word embedding spaces and arguably any other vector representation of data in general (Huang et al., 2018; Cogswell et al., 2016). From the geometric point of view, a space is called isotropic if the vectors within that space are uniformly distributed in all directions. Lacking isotropy in the embedding space affects not only the optimization procedure (e.g., model’s accuracy and convergence time) but also the expressiveness of the embedding space; hence, improving the isotropy of the embedding space can lead to performance improvements (Wang et al., 2020; Ioffe and Szegedy, 2015).

We measure the isotropy of embedding space using the partition function of Arora et al. (2016):

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

where u is a unit vector, w_i is the corresponding embedding for the i^{th} word in the embedding matrix $W \in \mathbb{R}^{N \times D}$, N is the number of words in the vocabulary, and D is the embedding size. Arora et al. (2016) showed that $F(u)$ can be approximated using a constant for isotropic embedding spaces. Therefore, for the set U , which is the set of eigenvectors of $W^T W$, in the following equation, $I(W)$ would be close to one for a perfectly isotropic space (Mu and Viswanath, 2018).

$$I(W) = \frac{\min_{u \in U} F(u)}{\max_{u \in U} F(u)} \quad (2)$$

2.1 Analyzing Isotropy in pre-trained CWRs

Using the above metric, we analyzed the representation degeneration problem globally and locally.

Global assessment. We quantified isotropy in all layers for GPT-2, BERT, and RoBERTa on the development set of STS-Benchmark (Cer et al., 2017). Figure 1 shows the trend of isotropy in all layers based on $I(W)$. Clearly, all CWRs are extremely anisotropic in all non-input layers. While

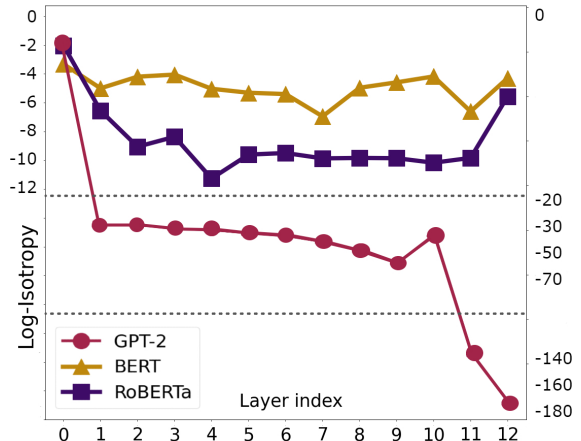


Figure 1: Layer-wise isotropy for different CWRs on the STS-B dev set (\uparrow log-isotropy: \uparrow isotropy). Given the large difference, BERT and RoBERTa are shown on the left axis and GPT-2 on the right.

the isotropy of GPT-2 decreases consistently in upper layers, that for RoBERTa has a semi-convex form in which the last layer (except for the input layer) has the highest isotropy. Also, interestingly, the input layer in GPT-2 is more isotropic than those for the other two models. This observation contradicts with what has been previously reported by Ethayarajh (2019).

Local assessment. In the light of the clustered structure of the embedding space in CWRs (Reif et al., 2019), we carried out a local investigation of isotropy. To this end, we clustered the space using k -means and measured isotropy after making each cluster zero-mean (Mu and Viswanath, 2018). Table 1 shows the results for different number of clusters (each being the average of five runs). When the embedding space is viewed closely, the distribution of CWRs is notably more isotropic. Clustering significantly enhances isotropy for BERT and RoBERTa, making their embedding spaces almost isotropic. However, GPT-2 is still far from being isotropic. This contradicts with the observation of Cai et al. (2021).

A possible explanation for these contradictions is the different metric used by Ethayarajh (2019) and Cai et al. (2021) for measuring isotropy: cosine similarity. Randomly sampled words in an anisotropic embedding space should have high cosine similarities (a near-zero similarity denotes isotropy). However, there are exceptional cases where this might not hold (an anisotropic embedding space where sampled words have near-zero cosine similarities). In Figure 2, we illustrate GPT-

	GPT-2	BERT	RoBERTa
Baseline	5.02E-174	5.05E-05	2.70E-06
$k = 1$	2.49E-220	0.010	0.015
$k = 3$	9.42E-66	0.040	0.290
$k = 6$	1.40E-41	0.125	0.453
$k = 9$	1.18E-41	0.131	0.545
$k = 20$	4.06E-47	0.262	0.603

Table 1: CWRs isotropy after clustering and making each cluster zero-mean separately (results for different number of clusters (k) on STS-B dev set).




Figure 2: GPT-2 embeddings on STS-B dev set before (top) and after (bottom) a local zero-mean operation.

2 embedding space as an example for such an exceptional cases. Making individual clusters zero-mean (bottom) improves isotropy over the baseline (top). However, the embeddings are still far from being uniformly distributed in all directions. Instead, they are distributed around a horizontal line. This leads to a near-zero cosine similarity for randomly sampled words while the embedding space is anisotropic. Hence, cosine similarity might not be a proper metric for computing isotropy.

3 Cluster-based Isotropy Enhancement

The degeneration problem in the embedding space can be attributed to the training procedure of the underlying models, which are often language models trained through likelihood maximization with the weight tying trick (Gao et al., 2019). Maximizing the likelihood of a specific word embedding (minimizing that for others) requires pushing it towards the direction of the corresponding hidden state, which results in the accumulation of the learnt word embeddings into a narrow cone.

Previous work has shown that nulling out dominant directions of an anisotropic embedding space can make the space isotropic and improve its expressiveness (Mu and Viswanath, 2018). We refer

to this as the *global* approach. This method was proposed for static embeddings. Hence, it might not be optimal for contextual embeddings, especially in the light that the latter tends to have a clustered structure. For instance, recent work suggests that word types (e.g., verbs, nouns, punctuations), entities (e.g., personhood, nationalities, and dates), and even word senses (Michael et al., 2020; Loureiro et al., 2021; Reif et al., 2019) create local distinct clustered areas in the contextual embedding space. Moreover, our local assessment shows that it is not necessarily the case that all clusters share the same dominant directions. Hence, discarding dominant directions that are computed globally is not efficient for removing local degenerated directions. Consequently, it is more logical to have a cluster-specific dropping of dominant directions.

Based on these observations, we propose a cluster-based approach for isotropy enhancement. Specifically, instead of determining dominant directions globally, we obtain them separately for different sub-spaces and discard for each cluster only the corresponding cluster-specific dominant directions. To this end, we employ Principal Component Analysis (PCA) to compute local dominant directions in clusters. Geometrically, principal components (PCs) represent those directions in which embeddings have the most variance (maximum elongation). In our proposed method, we first cluster word embeddings using a simple k -means algorithm. After making each cluster zero-mean, the top PCs of every cluster are removed separately. Adding a clustering step helps us to eliminate the local dominant directions of each cluster. We will show in Section 5 that different linguistic knowledge is encoded in the dominant directions of various clusters. Moreover, numerical results show that in comparison with the global approach, our method can make the embedding space more isotropic, even when the fewer number of PCs are nulled out.

4 Experiments

We carried out experiments on the following benchmarks. As for Semantic Textual Similarity (STS), which is the main benchmark for our experiments, we experimented with STS 2012-2016 datasets (Agirre et al., 2012, 2013, 2014, 2015, 2016), the SICK-Relatedness dataset (SICK-R) (Marelli et al., 2014), and the STS benchmark (STS-B). For the STS task, we report results for GPT-2, BERT, and RoBERTa. We also experimented with a number

	Model	STS 2012	STS 2013	STS 2014	STS 2015	STS 2016	SICK-R	STS-B
Baseline	GPT-2	26.49	30.25	35.74	41.25	46.40	45.05	24.8
	BERT-base	42.87	59.21	59.75	62.85	63.74	58.69	47.4
	RoBERTa-base	33.09	56.44	46.76	55.44	60.88	61.28	56.0
Global approach	GPT-2	51.42	69.71	55.91	60.35	62.12	59.22	55.7
	BERT-base	54.62	70.39	60.34	63.73	69.37	63.68	65.5
	RoBERTa-base	51.59	73.57	60.70	66.72	69.34	65.82	70.1
Cluster-based approach	GPT-2	52.40	72.71	59.23	62.19	64.26	59.51	62.3
	BERT-base	58.34	75.65	63.55	64.37	69.63	63.75	66.0
	RoBERTa-base	54.87	76.70	64.18	67.05	69.28	66.93	71.4

Table 2: Spearman correlation performance of three pre-trained models (baseline) on the Semantic Textual Similarity datasets, before and after isotropy enhancement with the global and cluster-based (our) approach.

	RTE	CoLA	SST-2	MRPC	WiC	BoolQ	Average
Baseline	54.4	38.0	80.1	70.2	60.0	64.7	61.2
Global approach	56.2	38.8	80.2	72.1	60.7	64.9	62.1
Cluster-based approach	56.5	40.7	82.5	72.4	61.0	66.4	63.2

Table 3: Results on the classification tasks (BERT) in terms of accuracy (except for CoLA: Matthew’s correlation).

of classification tasks: Recognizing Textual Entailment from the GLUE benchmark (Wang et al., 2018, RTE), the Corpus of Linguistic Acceptability (Warstadt et al., 2019, CoLA), Stanford Sentiment Treebank (Socher et al., 2013, SST-2), Microsoft Research Paraphrase Corpus (Dolan and Brockett, 2005, MRPC), Word-in-Context (Pilehvar and Camacho-Collados, 2019, WiC), and BoolQ (Clark et al., 2019). For the classification tasks, we limit our experiments to BERT and extract features to train an MLP. Further details on the datasets and system configuration can be found in Appendix B.

We benchmark our cluster-based approach with the pre-trained CWRs (baseline) and the global method. As it was mentioned before, this method is similar to ours in its elimination of a few top dominant directions but with the difference that these directions are computed globally (in contrast to our local cluster-based computation). The best setting for each model is selected based on performance on the STS-B dev set. The reported results are the average of five runs.

4.1 Results

Tables 2 and 3 report experimental results. As can be seen, globally increasing isotropy can make a significant improvement for all the three pre-trained models. However, our cluster-based approach can achieve notably higher performance compared to the global approach. We attribute this improvement to our cluster-specific discarding of dominant directions. Both global and cluster-based methods null

out the optimal number of top dominant directions (tuned separately, cf. Appendix B), but the latter identifies them based on the specific structure of a sub-region in the embedding space (which might not be similar to other sub-regions).

5 Discussion

In this section, we provide a brief explanation for reasons behind the effectiveness of the cluster-based approach through investigating the linguistic knowledge encoded in the dominant local directions. We also show that enhancing isotropy reduces convergence time.

5.1 Linguistic knowledge

Punctuations and stop words. We observed that local dominant directions for the clusters of punctuations and stop words carry structural and syntactic information about the sentences in which they appear. For example, the two sentences “A man is crying.” and “A woman is dancing.” from STS-B do not have much in common in terms of semantics but are highly similar with respect to their style. To quantitatively analyze the distribution of this type of tokens in CWRs, we designed an experiment based on the dataset created by Ravfogel et al. (2020). The dataset consists of groups in which sentences are structurally and syntactically similar but have no semantic similarity. We picked 200 different structural groups in which each group has six semantically different sentences. Then, using the k -NN algorithm, we calculated the percentage of

Model	Baseline				Removed PCs			
	ST-SM	ST-DM	DT-SM	Isotropy	ST-SM	ST-DM	DT-SM	Isotropy
GPT-2	48.82	48.19	50.86	2.26E-05	9.32	9.53	9.49	0.17
BERT	13.44	14.24	14.87	2.24E-05	10.31	10.50	10.32	0.32
RoBERTa	5.89	6.31	6.86	1.22E-06	4.78	5.00	4.89	0.73

Table 4: The mean Euclidean distance of a sample occurrence of a verb to all other occurrences of the same verb with the Same-Tense and the Same-Meaning (ST-SM), the Same-Tense but Different-Meaning (ST-DM), and a Different-Tense but the Same-Meaning (DT-SM). Semantically, it is desirable for DT-SM to be lower than ST-DM.




Figure 3: The percentage of nearest neighbours that share similar structural and syntactic knowledge, before (lighter, pattern-filled) and after removing dominant directions in pre-trained CWRs.

nearest neighbours which are in the same group before and after removing local dominant directions. We evaluated this for period and comma, which are the most frequent punctuations, and “the” and “of” as the most contextualized stop words (Ethayarajh, 2019). The reported results in Figure 3 show that the representations for punctuations and stop words are biased toward structural and syntactic information of sentences; hence, removing their dominant directions reduces the number of same-group nearest neighbours. The improvement from our local isotropy enhancement can be partially attributed to attenuating this type of bias.

Verb Tense. Our experiments show that tense is more dominant in verb representations than sense-level semantic information. To have a precise examination of this hypothesis, we used SemCor (Miller et al., 1993), a dataset comprising around 37K sense-annotated sentences. We collected representations for polysemous verbs that had at least two senses occurring a minimum of 10 times. Then, for each individual verb, we calculated Euclidean distance to the contextual representation of the same verb: (1) with the same tense and the same meaning, (2) with the same tense but a different meaning, and (3) with a different tense and the same mean-




Figure 4: The impact of our cluster-based isotropy enhancement on per-epoch performance for two tasks.

ing. The experimental results reported in Table 4 confirm the hypothesis and show the effectiveness of the cluster-based approach in bringing together verb representations that correspond to the same sense, even if they have different tense.

5.2 Convergence time

In the previous experiments, we showed that the contextual embeddings are extremely anisotropic and highly correlated. Such embeddings can slow down the learning process of deep neural networks. Figure 4 shows the trend of convergence for the BoolQ and RTE tasks (dev sets). By decreasing the correlation between embeddings, our method can reduce convergence time.

6 Conclusions

In this paper, we proposed a cluster-based method to address the representation degeneration problem in CWRs. We empirically analyzed the effect of clustering and showed that, from a local sight, most clusters are biased toward structural information. Moreover, we found that verb representations are distributed based on their tense in distinct sub-spaces. We evaluated our method on different semantic tasks, demonstrating its effectiveness in removing local dominant directions and improving performance. As future work, we plan to study the effect of fine-tuning on isotropy and on the encoded linguistic knowledge in local regions.

References

- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Iñigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, German Rigau, Larraitz Uria, and Janyce Wiebe. 2015. [SemEval-2015 task 2: Semantic textual similarity, English, Spanish and pilot on interpretability](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 252–263, Denver, Colorado. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2014. [SemEval-2014 task 10: Multilingual semantic textual similarity](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 81–91, Dublin, Ireland. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2016. [SemEval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 497–511, San Diego, California. Association for Computational Linguistics.
- Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. [SemEval-2012 task 6: A pilot on semantic textual similarity](#). In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 385–393, Montréal, Canada. Association for Computational Linguistics.
- Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013. [*SEM 2013 shared task: Semantic textual similarity](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 32–43, Atlanta, Georgia, USA. Association for Computational Linguistics.
- 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.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). 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 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Michael Cogswell, Faruk Ahmed, Ross B. Girshick, Larry Zitnick, and Dhruv Batra. 2016. [Reducing overfitting in deep networks by decorrelating representations](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- 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.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- 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 Tie-Yan Liu. 2019. [Representation degeneration problem in training natural language generation models](#). *CoRR*, abs/1907.12009.
- Chengyue Gong, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2018. [FragE: Frequency-agnostic word representation](#). In *Advances in Neural Information Processing Systems*, volume 31, pages 1334–1345. Curran Associates, Inc.
- Lei Huang, Dawei Yang, Bo Lang, and Jia Deng. 2018. [Decorrelated batch normalization](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML'15, page 448–456. JMLR.org.
- 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](#). *CoRR*, abs/1907.11692.
- Daniel Loureiro, Kiamehr Rezaee, Mohammad Taher Pilehvar, and Jose Camacho-Collados. 2021. [Analysis and Evaluation of Language Models for Word Sense Disambiguation](#). *Computational Linguistics*, pages 1–57.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. [A SICK cure for the evaluation of compositional distributional semantic models](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- 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.
- George A. Miller, Claudia Leacock, Randee Teng, and Ross T. Bunker. 1993. [A semantic concordance](#). In *Human Language Technology: Proceedings of a Workshop Held at Plainsboro, New Jersey, March 21-24, 1993*.
- Jiaqi Mu and Pramod Viswanath. 2018. [All-but-the-top: Simple and effective postprocessing for word representations](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. [WiC: the word-in-context dataset for evaluating context-sensitive meaning 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 1267–1273, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Shauli Ravfogel, Yanai Elazar, Jacob Goldberger, and Yoav Goldberg. 2020. [Unsupervised distillation of syntactic information from contextualized word representations](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 91–106, 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, pages 8594–8603. Curran Associates, Inc.
- 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 treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. 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*.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.

A Isotropy statistics

Table 5 shows isotropy statistics for GPT-2, BERT, and RoBERTa. GPT-2’s embedding space is extremely anisotropic in upper layers. Hence, more PCs are required to be eliminated to make this embedding space isotropic in comparison to BERT and RoBERTa, both in the cluster-based approach and the global one (Mu and Viswanath, 2018). Also, in almost all layers, BERT has higher a isotropy than RoBERTa.

Model	GPT-2	BERT	RoBERTa
layer 0	1.5E-02	4.6E-04	9.1E-03
layer 1	9.9E-24	9.9E-06	2.7E-07
layer 2	2.8E-23	6.3E-05	8.7E-10
layer 3	6.1E-26	8.8E-05	4.2E-09
layer 4	1.6E-27	9.2E-06	5.4E-12
layer 5	3.0E-30	4.8E-06	2.4E-10
layer 6	1.6E-32	3.9E-06	3.1E-10
layer 7	1.3E-37	1.1E-07	1.3E-10
layer 8	3.4E-45	1.0E-05	1.4E-10
layer 9	6.4E-55	2.5E-05	1.3E-10
layer 10	4.1E-32	6.9E-05	6.7E-11
layer 11	1.8E-132	2.4E-07	1.4E-10
layer 12	5.0E-174	5.0E-05	2.7E-06

Table 5: Per-layer isotropy on the STS-B dev set. Numbers have been calculated based on $I(W)$.

B Experimental Setup

B.1 Dataset details

STS. In the Semantic Textual Similarity task, the provided labels are between 0 and 5 for each paired sentence. We first calculate sentence embeddings by averaging all word representations in each sentence and then compute the cosine similarity between two sentence representations as a score of semantic relatedness of the pair.

RTE. The Recognizing Textual Entailment dataset is a classification task from the GLUE benchmark (Wang et al., 2018). Paired sentences are collected from different textual entailment challenges and labeled as *entailment* and *not-entailment*.

CoLA. The Corpus of Linguistic Acceptability (Warstadt et al., 2019) is a binary classification task in which sentences are labeled whether they are grammatically acceptable.

SST-2. The Stanford Sentiment Treebank (Socher et al., 2013) is a binary sentiment classification task.

MRPC. The Microsoft Research Paraphrase Corpus (Dolan and Brockett, 2005) consists of paired sentences, and the goal is determining whether, in a pair, sentences share similar semantics or not.

WiC. Word-in-Context (Pilehvar and Camacho-Collados, 2019) is a binary classification task in which it should be determined if a target word in two different contexts refers to the same meaning.

BoolQ. Boolean Questions (Clark et al., 2019) is a Question Answering classification task. Every sample includes a passage and a yes/no question about the passage.

B.2 Configurations

For the classification tasks, we trained a simple MLP on the features extracted from BERT. The proposed cluster-based approach has two hyperparameters: the number of clusters and the number of PCs to be removed. We selected both of them from range [5, 30] and tuned them on the STS-B dev set. In the cluster-based approach, The optimal number of clusters for GPT-2, BERT, and RoBERTa are respectively 10, 27, and 27. For BERT and RoBERTa, 12 top dominant directions have been removed, while the number is 30 for GPT-2 regarding its extremely anisotropic embedding space. The tuning of the number of PCs to be eliminated in the global method has been done similarly to the cluster-based approach (on the STS-B dev set): 30, 15, and 25 for GPT-2, BERT, and RoBERTa, respectively.

C Isotropy on STS datasets

In Table 6, we present the isotropy of the contextual embedding spaces calculated using $I(W)$ on the STS benchmark. The results reveal the effectiveness of the proposed method in enhancing the isotropy of the embedding space.

D Word frequency bias in CWRs

CWRs are biased towards their frequency information, and words with similar frequency create local regions in the embedding space (Gong et al., 2018; Li et al., 2020). From the semantic point of view, this is certainly undesirable given that words with similar meanings but different frequencies could be

Model	STS 2012	STS 2013	STS 2014	STS 2015	STS 2016	SICK-R	STS-B
<i>Baseline</i>							
GPT-2	1.4E-178	1.0E-170	1.4E-172	2.9E-177	6.0E-174	9.9E-140	2.6E-105
BERT	3.1E-05	1.9E-04	2.6E-04	3.7E-07	2.8E-04	4.2E-05	1.1E-04
RoBERTa	3.1E-06	3.1E-07	3.8E-06	3.8E-06	3.5E-06	3.7E-07	2.9E-06
<i>Global approach</i>							
GPT-2	0.57	0.40	0.05	0.12	0.60	0.57	0.51
BERT	0.48	0.41	0.55	0.72	0.65	0.63	0.58
RoBERTa	0.67	0.87	0.87	0.84	0.85	0.90	0.88
<i>Cluster-based approach</i>							
GPT-2	0.71	0.74	0.47	0.74	0.74	0.78	0.70
BERT	0.68	0.61	0.77	0.81	0.75	0.82	0.73
RoBERTa	0.89	0.91	0.93	0.92	0.89	0.94	0.90

Table 6: Isotropy of CWRs on multiple STS datasets calculated based on $I(W)$; a higher value indicates a more isotropic embedding space. Our cluster-based method significantly increases the isotropy of embedding space on all datasets.

located far from each other in the embedding space. This phenomenon can be seen in Figure 5. The encoded knowledge in the local dominant directions partly correspond to frequency information. The embedding space visualization reveals that our approach performs a decent job in removing frequency bias in pre-trained models.

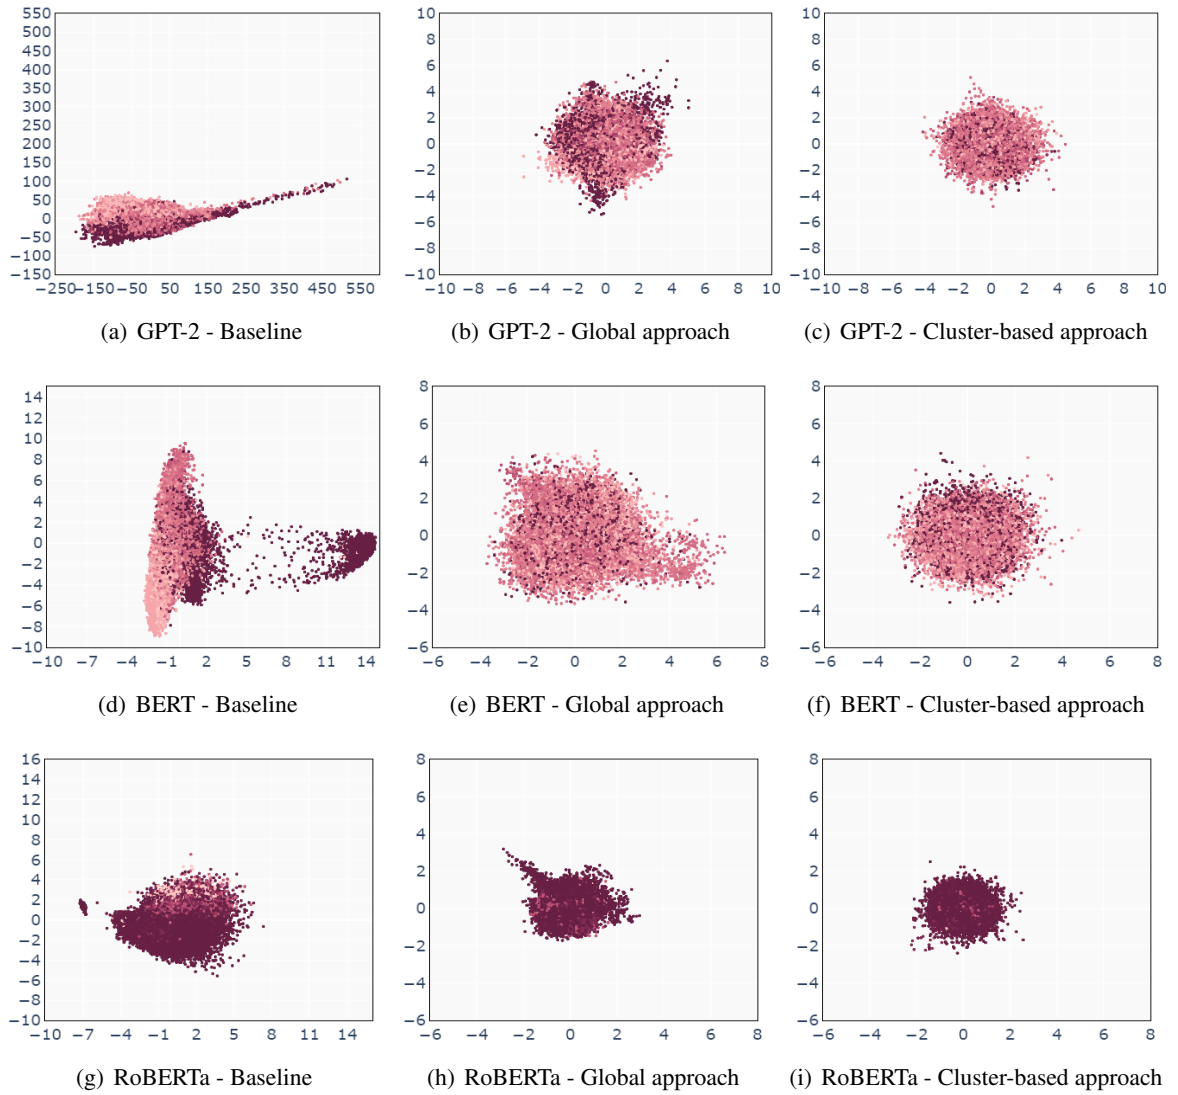


Figure 5: Contextual Word Representations visualization using PCA on STS-B dev set. Colors indicate word frequency in the Wikipedia dump (the lighter point, the more frequent).