

IsoBN: Fine-Tuning BERT with Isotropic Batch Normalization

Wenxuan Zhou, Bill Yuchen Lin, Xiang Ren

Department of Computer Science, University of Southern California, Los Angeles, CA
 {zhouwenx, yuchen.lin, xiangren}@usc.edu

Abstract

Fine-tuning pre-trained language models (PTLMs), such as BERT and its better variant RoBERTa, has been a common practice for advancing performance in natural language understanding (NLU) tasks. Recent advance in representation learning shows that *isotropic* (i.e., unit-variance and uncorrelated) embeddings can significantly improve performance on downstream tasks with faster convergence and better generalization. The isotropy of the pre-trained embeddings in PTLMs, however, is relatively under-explored. In this paper, we analyze the isotropy of the pre-trained [CLS] embeddings of PTLMs with straightforward visualization, and point out two major issues: high variance in their standard deviation, and high correlation between different dimensions. We also propose a new network regularization method, isotropic batch normalization (IsoBN) to address the issues, towards learning more isotropic representations in fine-tuning by dynamically penalizing dominating principal components. This simple yet effective fine-tuning method yields about 1.0 absolute increment on the average of seven NLU tasks.

Introduction

Pre-trained language models (PTLMs), such as BERT (Devlin et al. 2019) and RoBERTa (Liu et al. 2019b), have revolutionized the area of natural language understanding (NLU). Fine-tuning PTLMs has advanced performance on many benchmark NLU datasets such as GLUE (Wang et al. 2019a). The most common fine-tuning method is to continue training pre-trained model parameters together with a few additional task-specific layers. The PTLMs and task-specific layers are usually connected by the embeddings of [CLS] tokens, which are regarded as sentence representations.

Recent works on text representation (Arora, Liang, and Ma 2017; Mu and Viswanath 2018; Gao et al. 2019; Wang et al. 2020) have shown that regularizing word embeddings to be more *isotropic* (i.e., rotational invariant) can significantly improve their performance on downstream tasks. An ideally isotropic embedding space has two major merits: a) all dimensions have the *same variance* and b) all dimensions are *uncorrelated* with each other. These findings align with conventional *feature normalization* techniques (Cogswell et al. 2016; Ioffe and Szegedy 2015; Huang et al. 2018), which aim

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Illustration of the isotropic batch normalization (IsoBN). The [CLS] embedding is normalized by standard deviation and pairwise correlation coefficient to get a more isotropic representation.

to transform input features into *normalized¹, uncorrelated* representations for faster convergence and better generalization ability.

It, however, remains an open question that how isotropic the representations of PTLMs are. Particularly, we want to understand the isotropy of pre-trained [CLS] embeddings in PTLMs, and how we can improve it towards better fine-tuning for downstream tasks. In this paper, we first argue the reason why we want more isotropic embeddings for the [CLS] tokens (Section). Our analysis reveals that the dominating principle components largely hinder the fine-tuning process to use knowledge in other components, due to the lack of isotropy. Then, we analyze the isotropy of the pre-trained [CLS] embeddings. There are two essential aspects of an isotropic embedding space: *unit-variance* and *uncorrelatedness*. Thus, we start our analysis by visualizing the *standard deviation* and *Pearson correlation coefficient* of pre-trained [CLS] embeddings in BERT and RoBERTa on several NLU datasets.

Our visualization and quantitative analysis in Section finds that: 1) the [CLS] embeddings have very *different variance* (Sec.); 2) the [CLS] embeddings construct a few large clusters of dimensions that are *highly correlated* with each other (Sec.). Both findings indicate that pre-trained contextualized word embeddings are far from being isotropic, i.e., normalized and uncorrelated. Therefore, these undesired prior bias from PTLMs may result in sub-optimal performance in fine-tuning for target tasks.

¹We use ‘normalized’ to refer unit-variance in this paper.

Given that pre-trained [CLS] embeddings are very anisotropic, a natural research question is then: *how can we regularize the fine-tuning process towards more isotropic embeddings?* There are two common methods for improving the isotropy of feature representations: whitening transformation and batch normalization (Ioffe and Szegedy 2015). However, both are not practically suitable in the scenario of fine-tuning PTLMs. Whitening transformation requires calculating the inverse of the covariance matrix, which are ill-conditioned in PTLMs’ embeddings. Unfortunately, calculating the inverse is thus numerically unstable, computationally expensive, and incompatible in half-precision training. Batch normalization is proposed to alleviate the inverse-computation issue by assuming that the covariance matrix is *diagonal*, which in turn completely ignores the influence of correlation between dimensions.

Motivated by the research question and limitations of existing works, we propose a new network regularization method, isotropic batch normalization (IsoBN) in Section . As shown in Figure 1, the proposed method is based on our observation that the embedding dimensions can be seen as several groups of highly-correlated dimensions. Our intuition is thus to assume that the absolute correlation coefficient matrix is a *block-diagonal binary matrix*, instead of only a diagonal matrix. The dimensions of the same group have an absolute correlation coefficient of 1 (duplicate of each other), and dimensions in different group of 0 (uncorrelated).

This method greatly reduces the computation efforts in calculating the inverse, and better models the characteristics of the pre-trained [CLS] embeddings. Our experiments (Sec.) show that the IsoBN indeed improves both BERT and RoBERTa in fine-tuning, yielding about 1.0 absolute increment on average of a wide range of 7 GLUE benchmark tasks. We also empirically analyze the isotropy increment brought by IsoBN via explained variance, which clearly shows that IsoBN produces much more isotropic embeddings than conventional batch normalization method.

To the best of our knowledge, this work is the first one in studying the isotropy of the pre-trained [CLS] embeddings. We believe our findings and the proposed IsoBN method will inspire interesting future research directions in improving pre-training language models as well as better fine-tuning towards more isotropy of PTLMs.

Why Isotropic [CLS] Embeddings?

We formally show our analysis on the principle components of the [CLS] embeddings. Our findings reveal that with random weight initialization, the first few principle components usually take the majority of contributions for the prediction results.

Background knowledge. For text classification, the input text x is first encoded by the PTLM to feature \mathbf{h} (i.e., their [CLS] embeddings), and then classified by a random-initialized softmax layer:

$$\mathbf{h} = \text{PTLM}(x); p_i = \frac{\exp(\mathbf{W}_i^T \mathbf{h})}{\sum_{j=1}^c \exp(\mathbf{W}_j^T \mathbf{h})},$$

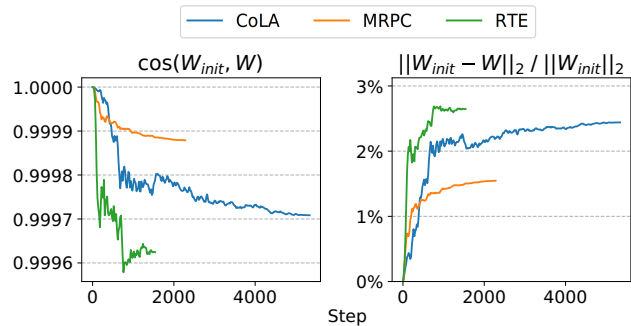


Figure 2: Average cosine similarity and L2 distance of fine-tuned weight \mathbf{W} to the initialized weight \mathbf{W}_{init} during the entire training process. Both measures suggest that the change of weight is very subtle.

where $\mathbf{W} \in \mathcal{R}^{d \times c}$ is a random-initialized learnable parameter. It learns towards mapping the underlying features extracted by PTLMs into target classes for input examples.

Previous work (Dodge et al. 2020) has shown that initialized weight \mathbf{W}_{init} of classifier has a large impact on the model performance. We further find that the final converged weights nearly remain the same to be the initialization after fine-tuning. We visualize this surprising phenomenon in Figure 2. We first project both \mathbf{W}_{init} and \mathbf{W} to the subspace spanned by the top 10 eigenvectors of $\text{Cov}(\mathbf{h})$ to remove the unimportant components, then use two similarity metrics (cosine similarity and L2 distance) to measure the difference between the initialized weight \mathbf{W}_{init} and the fine-tuned weight \mathbf{W} . We observe that the cosine similarity between \mathbf{W}_{init} and \mathbf{W} is extremely close to 1, and their L2 distance is close to 0. It suggests that the weight \mathbf{W} of the classifier in the fine-tuned model is almost determined in initialization. For example, there is a 0.9997 cosine similarity between initialized and fine-tuned weights on CoLA with RoBERTa.

Dominating Principle Components. Knowing that the weight of classifier is nearly fixed, we infer that the classifier may not capture the discriminative information for classification during fine-tuning. We measure the informativeness of each principal component, by comparing the variance of logits produced by it (between the fine-tuned classifier and the optimal classifier). Specifically, we fit a logistic regression model on the entire training data using the scikit-learn (Pedregosa et al. 2011) framework to get the optimal classifier. The variance Var_i of logits by the i^{th} principal component (w_i, \mathbf{v}_i) is calculated by: $\text{Var}_i = w_i \cdot (\mathbf{W}^T \mathbf{v}_i)^2$, where w_i, \mathbf{v}_i are the i^{th} eigenvalue and eigenvector. Intuitively, a decent classifier should maximize the variance along informative principal components and minimize irrelevant ones. We show the average proportion of variance in Figure 3.

As shown in our experiments (Sec.), the top principal components are constantly exaggerated by the fine-tuned weights, especially at the beginning of training process. The first principal component accounts for over 90% variance of logits throughout the training, which thus hinder learning from other useful components. This motivates us to penalize the

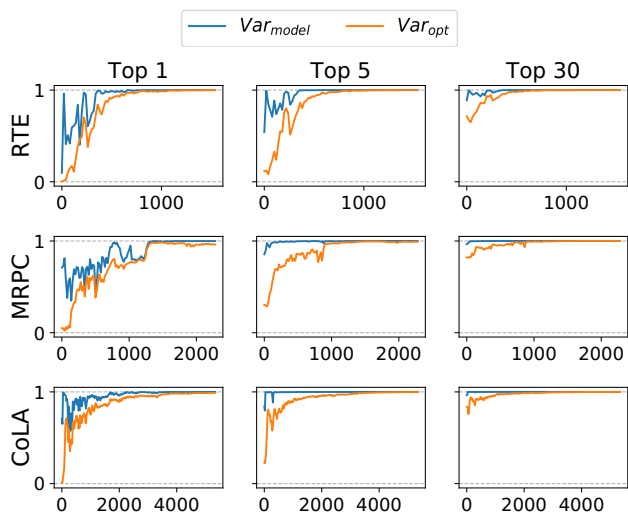


Figure 3: The average percentage of variance of logit produced by the top 1, 5, and 30 principal components. The top principal components are consistently exaggerated.

top principal components dynamically during the fine-tuning for avoid losing knowledge of PTLMs due to dominating components.

How Isotropic Are [CLS] Embeddings?

The [CLS] embeddings of PTLMs, regarded as sentence representations, are directly used for fine-tuning (e.g., BERT and RoBERTa) towards a wide range of downstream tasks. Given its impact in fine-tuning, we want to understand their isotropy. As we know unit-variance and uncorrelatedness are two essential aspects of an isotropy space, we start our investigation by analyzing the *standard deviation* and *Pearson correlation coefficient*.

Specifically, we take the corpus of four popular NLU tasks (MRPC, RTE, CoLA, and STS-b) from the GLUE benchmark datasets (Wang et al. 2019b) and then analyze their pre-trained [CLS] embeddings in terms of standard deviation (Section) and correlation (Section) respectively.

Analysis of Standard Deviation

To visualize the standard deviation of embeddings, we take the input sentences of the whole training dataset to PTLMs, then calculate the standard deviation on their [CLS] embeddings, and finally obtain the distribution.

The standard deviation of a nearly isotropic embedding space should concentrate on a very small range of values. Simply put, an ideally isotropic embedding space should have a small variance of the distribution of their standard deviation, i.e., all dimensions of [CLS] embeddings should have almost the same standard deviation.

As shown in Figure 4, we can see that both BERT and RoBERTa do not have such desired property for pre-trained [CLS] embeddings. The standard deviations of the embeddings vary in a very wide range of values (e.g., $[10^{-10}, 1]$ in BERT for MRPC). Interestingly, we can see that RoBERTa is

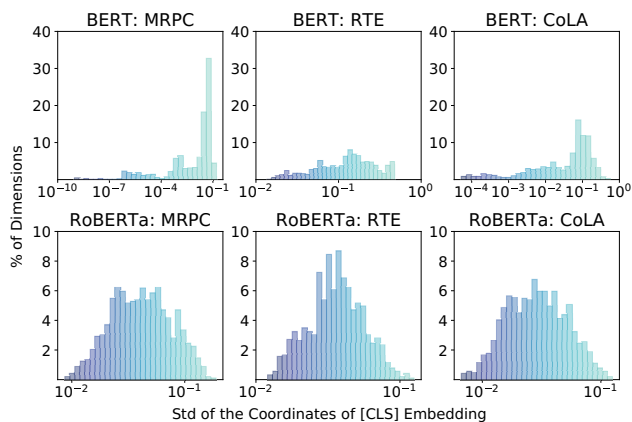


Figure 4: The distribution of the standard deviation (std) of pre-trained [CLS] embeddings. We show the results of BERT-base and RoBERTa-Large on 4 NLU datasets. Note that an (nearly) isotropic embedding space should have (almost) zero variance on the std (i.e., 100% dimensions have the same std).

evidently better than BERT from this perspective (e.g., usually ranging in $[0.01, 1]$). However, the [CLS] embeddings of RoBERTa are still far from being isotropic, as there is no significantly dominant centering standard deviation value.

Analysis of Correlation Coefficient

Correlation between different dimensions of [CLS] embeddings is an essential aspect of isotropy. Embeddings with low correlation between dimensions usually show better generalization on downstream applications (Cogswell et al. 2016). It, however, is relatively ignored by many neural network regularization methods, such as batch normalization (Ioffe and Szegedy 2015).

In order to better visualize the Pearson correlation coefficient of [CLS] embeddings, we cluster the dimensions by their pairwise coefficient, and then re-arrange the dimension index, such that highly correlated dimensions locate near each other. The absolute value of correlations are shown in Figure 5, where darker cells means higher correlation.

We can see that both BERT and RoBERTa usually have very high correlations between different dimensions (i.e., most cells are in dark blue), although the situation is less severe in a few cases such as BERT on CoLA and RoBERTa on RTE. We find that BERT’s embeddings have several large clusters of correlated features, while RoBERTa tends to have a single extreme large cluster.

In either case, such high correlation between embedding dimensions is harmful to future fine-tuning. Recall that the [CLS] embeddings are usually connected to a linear classifier which is uniformly initialized. In the beginning of the fine-tuning process, the classifier will be biased to these features since they gain more importance in back-propagation. This undesired prior prevents models to exploit other potentially value features, and thus require more training data or epochs to converge and generalize in downstream tasks.

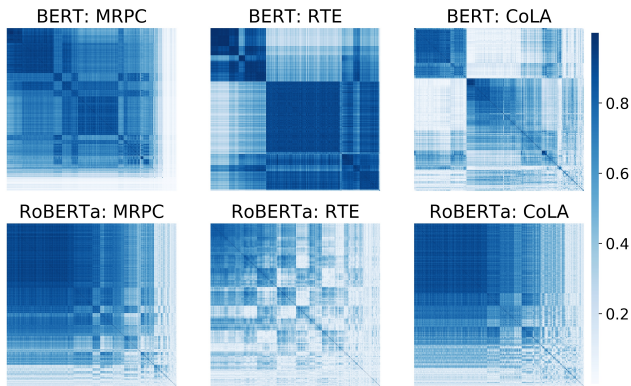


Figure 5: Absolute Pearson correlation coefficients between dimensions of pre-trained [CLS] embeddings. We show the results of BERT-base-based and RoBERTa-Large on four NLU datasets. Note that the dimension indexes of the matrices are rearranged by the clustering results. Ideally, an isotropic embedding space should be 1 (darkest blue) on the diagonal and 0 (white) on other cells. A dark block in a matrix means a cluster of features highly correlated with each other.

We argue that these two findings together indicate that pre-trained language models are far from being isotropic (i.e., normalized and uncorrelated), and thus undesired prior bias may result in sub-optimal model performance for fine-tuning.

Approach

Based on our analysis in Section , we propose a new regularization method, isotropic batch normalization towards learning more isotropic representations of the [CLS] tokens and thus better fine-tuning PTLMs. We first introduce some background knowledge about *whitening* and conventional *batch normalization* methods (Section), then formally introduce the proposed IsoBN (Section), and finally show the implementation details.

Whitening and Batch Normalization

To improve the isotropy of feature representations, there are two widely-used methods: 1) whitening transformation and 2) batch normalization (Ioffe and Szegedy 2015).

Whitening transformation changes the input vector into a white noise vector, and can be defined as a transformation function as follows:

$$\hat{\mathbf{h}} = \Sigma^{-\frac{1}{2}}(\mathbf{h} - \boldsymbol{\mu} \cdot \mathbf{1}^T), \quad (1)$$

where $\Sigma \in \mathcal{R}^{d \times d}$ is the covariance matrix of the input $\mathbf{h} \in \mathcal{R}^{d \times N}$, $\boldsymbol{\mu} \in \mathcal{R}^d$ is the mean of \mathbf{h} . Thus, the transformation is a mapping from $\mathcal{R}^{d \times N} \rightarrow \mathcal{R}^{d \times N}$. This transformation produces a perfectly isotropic embedding space, where the dimensions are uncorrelated and have the same variance. It can be applied in either feature pre-processing (Rosipal et al. 2001) or neural network training (Huang et al. 2018). A similar method is to remove a few top principal components from the embedding space (Arora, Liang, and Ma 2017; Mu and




Figure 6: Illustration of the assumption by batch normalization and our IsoBN with the reference of real correlation. IsoBN assumes that the absolute correlation matrix is *block-diagonal* while batch normalization completely ignores the correlation.

Viswanath 2018). However, these methods are hard to apply in fine-tuning PTLMs, as they require calculating the *inverse of the covariance matrix*. As shown in Section , the embeddings in PTLMs contains groups of highly-correlated dimensions. Therefore, the covariance matrices are ill-conditioned, and calculating the inverse is thus numerically unstable. It is also computationally expensive and incompatible in half-precision training.

Batch normalization (BN) aims to simplify the inverse-computation problem by assuming that the covariance matrix is diagonal, thus the whitening function becomes:

$$\hat{\mathbf{h}} = \Lambda^{-1}(\mathbf{h} - \boldsymbol{\mu} \cdot \mathbf{1}^T), \quad (2)$$

where $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_d)$ is a diagonal matrix consisting the standard deviation of each input dimension. Batch normalization greatly improves the stability and model performance in training deep neural networks.

However, it completely ignores the influence of correlation in the embeddings, and thus not suitable for our interested [CLS] embeddings, where high correlation is a critical issue that needs to be addressed. We seek to design a novel normalization method specially for fine-tuning PTLMs, which can be efficiently computed yet still improve representations towards isotropy property.

Isotropic Batch Normalization

Recall Figure 5, from the correlation matrix of pre-trained embeddings, we observe that on most datasets, the correlation matrix is nearly *block-diagonal*². That is, the embedding dimensions form several clusters of highly-correlated dimensions. Dimensions within the same cluster have an absolute correlation coefficient of nearly 1, while dimensions from different clusters are almost uncorrelated. Inspired by this, we propose an enhanced simplification of the covariance matrix.

We assume that the absolute correlation coefficient matrix is a **block-diagonal binary matrix**. That is, the embedding dimensions can be clustered into m groups $\mathcal{G}_1, \dots, \mathcal{G}_m$, where dimensions of the same group have an absolute correlation coefficient of 1 (duplicate of each other), and dimensions in different group have a correlation coefficient of 0 (uncorrelated). This assumption is illustrate in Figure 6 as a

²A *block diagonal* matrix is a block matrix that is a square matrix such that the main-diagonal blocks are square matrices and all off-diagonal blocks are zero matrices.

conceptual comparison. Comparing with the conventional batch normalization, our assumption takes accounts of correlations and thus is an more accurate approximation of the realistic correlation coefficient matrices. Thereby, instead of whitening the correlation matrix, we want the influence of each group of dimensions similar in the fine-tuning process.

We first normalize each dimension to unit-variance, similar to batch normalization, for convenience of further derivation. This makes the dimensions in the same group exactly same to each other. Then, for dimension $i \in \mathcal{G}_{g(i)}$, it is repeated in embeddings by $|\mathcal{G}_{g(i)}|$ times. Therefore, the normalization transformation becomes:

$$\hat{\mathbf{h}}^{(i)} = \frac{1}{\sigma_i \cdot |\mathcal{G}_{g(i)}|} (\mathbf{h}^{(i)} - \mu_i \cdot \mathbf{1}^T). \quad (3)$$

The dimensions of embeddings, however, are not naturally separable into hard group divisions. Thus, we create a soft version of computing the size of a feature-group $|\mathcal{G}_{g(i)}|$ via the correlation coefficient matrix ρ :

$$|\mathcal{G}_{g(i)}| \rightsquigarrow \gamma_i = \sum_{j=1}^d \rho_{ij}^2. \quad (4)$$

This equation produces the same result as $|\mathcal{G}_{g(i)}|$ when our assumption holds in real correlation matrix. Finally, our transformation can be written as:

$$\hat{\mathbf{h}}^{(i)} = \frac{1}{\sigma_i \cdot \gamma_i} (\mathbf{h}^{(i)} - \mu_i \cdot \mathbf{1}^T). \quad (5)$$

The major difference between our method and conventional batch normalization is the introduction of the γ term, as a way to explicitly consider correlation between feature dimensions. As shown in our experiments (Section), γ can greatly improve the isotropy of embedding. We name our proposed normalization method as *isotropic batch normalization* (IsoBN), as it is towards more isotropic representations during fine-tuning.

The IsoBN is applied right before the final classifier. In experiments, we find that a modified version of IsoBN achieves better performance. The mean μ is highly unstable during training and does not affect the principal components of representations, so we remove it in our implementation. The scaling term $(\sigma \cdot \gamma)^{-1}$ has a small magnitude and damages the optimization of training loss, so we re-normalize it to keep the sum of variances in transformed embeddings same as the original one. We further introduce a hyper-parameter β , which controls the normalization strength, since it is shown in Section that the dominating eigenvalue problem varies from datasets. The modified IsoBN is written as:

$$\theta_i = (\sigma_i \cdot \gamma_i + \epsilon)^{-\beta}, \quad (6)$$

$$\bar{\theta} = \frac{\sum_{i=1}^d \sigma_i^2}{\sum_{i=1}^d \sigma_i^2 \theta_i^2} \cdot \theta, \quad (7)$$

$$\hat{\mathbf{h}} = \bar{\theta} \odot \mathbf{h}. \quad (8)$$

In the IsoBN, the calculation of the scaling factor relies on the covariance and standard deviation statistics of the embedding. We keep two moving average caches and update them in training because the estimated statistics from a single batch are not accurate. The whole algorithm of IsoBN is shown in Algorithm 1.

Algorithm 1: IsoBN Transformation

Input: Embedding h over a mini-batch:

$\mathcal{B} = \{h_{1\dots m}\}$; moving covariance Σ ; moving standard deviation σ ; momentum α .

Output: transformed embedding \hat{h} ; updated Σ , σ .

if training then

$$\begin{aligned} \mu_{\mathcal{B}} &= \frac{1}{m} \sum_{i=1}^m h_i \\ \sigma_{\mathcal{B}} &= \sqrt{\frac{1}{m} \sum_{i=1}^m (h_i - \mu_{\mathcal{B}})^2} \\ \Sigma_{\mathcal{B}} &= \frac{1}{m} (h - \mu_{\mathcal{B}})^T (h - \mu_{\mathcal{B}}) \\ \sigma &= \sigma + \alpha(\sigma_{\mathcal{B}} - \sigma) \\ \Sigma &= \Sigma + \alpha(\Sigma_{\mathcal{B}} - \Sigma) \end{aligned}$$

$$\rho = \frac{\Sigma}{\sigma \sigma^T}$$

Compute γ by Eq. 4

Compute scaling factor θ by Eq. 6 and Eq. 7

$$\hat{h} = \theta \odot h$$

Evaluation

In this section, we first present the setup of our experiments (i.e. the datasets, frameworks, and hyper-parameters), then discuss the empirical results, and finally evaluate the isotropy gain through the lens of explained variance.

Experiment Setup

Our implementation of PTLMs is based on HuggingFace Transformer (Wolf et al. 2019). The model is fine-tuned with AdamW (Loshchilov and Hutter 2019) optimizer using a learning rate in the range of $\{1 \times 10^{-5}, 2 \times 10^{-5}, 5 \times 10^{-5}\}$ and batch size in $\{16, 32\}$. The learning rate is scheduled by a linear warm-up (Goyal et al. 2017) for the first 6% of steps followed by a linear decay to 0. The maximum number of training epochs is set to 10. For IsoBN, the momentum α is set to 0.95, the ϵ is set to 0.1, and the normalization strength β is chosen in the range of $\{0.25, 0.5, 1\}$.

We apply early stopping according to task-specific metrics on the dev set. We select the best combination of hyper-parameters on the dev set. We fine-tune the PTLMs with 5 different random seeds and report the median and standard deviation of metrics on the dev set.

Experimental Results

We evaluate IsoBN on two PTLMs (BERT-base-cased and RoBERTa-large) and seven NLU tasks from the GLUE benchmark (Wang et al. 2019b). The experiments results are shown in Table 1. Using IsoBN improves the evaluation metrics on all datasets. The average score increases by 1% for BERT-base and 0.8% for RoBERTa-large. For small datasets (MRPC, RTE, CoLA, and STS-B (Cer et al. 2017)), IsoBN obtains an average performance improvement of 1.6% on BERT and 1.3% on RoBERTa. For large datasets (MNLI (Williams, Nangia, and Bowman 2018), QNLI (Rajpurkar et al. 2016), and SST-2 (Socher et al. 2013)), IsoBN obtains an average performance improvement of 0.15% on BERT and 0.25% on RoBERTa. This experiment shows that by improving the isotropy of embeddings, our IsoBN results in better fine-tuning performance.

Method	Avg	MNLI	QNLI	RTE	SST-2	MRPC	CoLA	STS-B
BERT-base (ReImp)	81.37	83.83 (.07)	90.82 (.1)	67.87 (1.1)	92.43 (.7)	85.29 (.9)	60.72 (1.4)	88.64 (.7)
BERT-base-IsoBN	82.36	83.91 (.1)	91.04 (.1)	70.75 (1.6)	92.54 (.1)	87.50 (.6)	61.59 (1.6)	89.19 (.7)
RoBERTa-L (ReImp)	88.16	90.48 (.07)	94.70 (.1)	84.47 (1.0)	96.33 (.3)	90.68 (.9)	68.25 (1.1)	92.24 (.2)
RoBERTa-L-IsoBN	88.98	90.69 (.05)	94.91 (.1)	87.00 (1.3)	96.67 (.3)	91.42 (.8)	69.70 (.8)	92.51 (.2)

Table 1: Empirical results on the dev sets of seven GLUE tasks. We run 5 times with different random seeds and report median and std. IsoBN outperforms the conventional fine-tuning method around 1.0 absolute increment.

$EV_1/EV_2/EV_3$	MRPC	RTE	CoLA	STS-b
BERT-base	0.76 / 0.87 / 0.89	0.88 / 0.93 / 0.95	0.49 / 0.58 / 0.64	0.89 / 0.94 / 0.96
BERT-base+BN	0.74 / 0.84 / 0.86	0.70 / 0.89 / 0.93	0.37 / 0.59 / 0.63	0.69 / 0.88 / 0.92
BERT-base+IsoBN	0.37 / 0.68 / 0.77	0.49 / 0.72 / 0.85	0.25 / 0.37 / 0.48	0.41 / 0.69 / 0.85
RoBERTa-L	0.86 / 0.90 / 0.91	0.53 / 0.66 / 0.70	0.83 / 0.88 / 0.90	0.87 / 0.90 / 0.92
RoBERTa-L+BN	0.64 / 0.73 / 0.76	0.36 / 0.50 / 0.57	0.61 / 0.70 / 0.75	0.65 / 0.72 / 0.77
RoBERTa-L+IsoBN	0.18 / 0.36 / 0.43	0.15 / 0.29 / 0.37	0.21 / 0.38 / 0.49	0.17 / 0.32 / 0.45

Table 2: The explained variance on BERT-base and RoBERTa-large. Compared to batch normalization, our method can greatly reduce the explained variance and thus improve the isotropy of embeddings.

Experiments of EV_k Metric

To quantitatively measure the isotropy of embeddings, we propose to use *explained variance* (EV) as the metric for isotropy, which is defined as:

$$EV_k(\mathbf{h}) = \frac{\sum_{i=1}^k \lambda_i^2}{\sum_{j=1}^d \lambda_j^2}, \quad (9)$$

where $\mathbf{h} \in \mathcal{R}^{N \times d}$ is the [CLS] embeddings, λ_i is the i^{th} largest singular value of the matrix \mathbf{h} . Note that N is the number of sentences in a certain corpus, and d is the dimension of hidden states in the last layer of a pre-trained language model.

This metric measures the difference of variance in different directions of the embedding space. Intuitively, if the EV_k value is small, the variations of embedding tend to distribute equally in all directions and lead to more angular symmetric representation. If the EV_k value is large, most of the variations will concentrate on the first few directions, and the embedding space will degrade to a narrow cone. Thus, the EV_k is a good metric of the isotropy of embedding space.

We use EV_k as the isotropy metric because it enjoys two beneficial properties:

- It is invariant to the magnitude of the embeddings, and thus comparisons between different models and datasets is more fair.
- It is also invariant to the mean value of the embeddings, aligning with sentence classification/regression tasks of our interest.

We compute the EV_k metric on two PTLMs (BERT-base and RoBERTa-large) and 4 tasks (MRPC, RTE, CoLA, STS-B). For IsoBN, the normalization strength β is 1. We show the first three EV_k value (EV_1 , EV_2 , and EV_3) in Table 2.

We observe that before normalization, the pre-trained [CLS] embeddings have very high EV_k value. The average

EV_3 value is around 0.86 for both BERT-base and RoBERTa-large. For some datasets (e.g. STS-b), the top three principal components already explain over 90% of the variance. Batch normalization can only reduce the EV_k value by a small margin (0.025 for BERT and 0.148 for RoBERTa on EV_3), because it ignores the correlations among embedding dimensions. Our proposed IsoBN greatly reduces the EV_k value (0.123 for BERT and 0.425 for RoBERTa on EV_3).

We also visualize the distribution of EV_k values in Figure 7. We choose the first 50 EV_k value for BERT-base and first 200 EV_k value for RoBERTa-large. We observe that with IsoBN, we can decrease the EV_k value of pre-trained embeddings. This experiment shows that compared to batch normalization, IsoBN can further improve isotropy of [CLS] embedding of PTLMs.

Related Work

Input Normalization

Normalizing inputs (Montavon and Müller 2012; He et al. 2016; Szegedy et al. 2017) and gradients (Schraudolph 1998; Bjorck et al. 2018) has been known to be beneficial for training deep neural networks. Batch normalization (Ioffe and Szegedy 2015) normalizes neural activations to zero-mean and unit-variance using batch statistics (mean, standard deviation). It is both empirically and theoretically proved that batch normalization can greatly smooth the loss landscape (Santurkar et al. 2018; Bjorck et al. 2018; Ghorbani, Krishnan, and Xiao 2019), which leads to faster convergence and better generalization. One drawback of batch normalization is that it ignores the correlations between input dimensions. Some methods (Huang et al. 2018, 2019) seek to calculate the full whitening transformation, while the IsoBN greatly simplifies the process by using the block-diagonal assumption, and succeeds to reduce the effect of dominating principal components by simple scaling of input dimensions.




Figure 7: The EV_k value on BERT-base and RoBERTa-large with no normalization, batch normalization, and IsoBN. We choose the max K to be 50 for BERT and 200 for RoBERTa. Compared to batch normalization, IsoBN can greatly reduce the EV_k value and thus improve the isotropy of $[CLS]$ embeddings.

Fine-tuning Language Models

Pre-trained language models (Devlin et al. 2019; Liu et al. 2019b) achieve the state-of-the-art performance on various natural language understanding tasks. To adapt the pre-trained language model to target tasks, the common practice, in which a task-specific classifier is added to the network and jointly trained with PTLMs, already produces good results (Peters, Ruder, and Smith 2019). Some works show that with well-designed fine-tuning strategies, the model performance can be further improved, by adversarial training (Zhu et al. 2020; Jiang et al. 2020), gradual unfreezing (Howard and Ruder 2018; Peters, Ruder, and Smith 2019), or multi-tasking (Clark et al. 2019; Liu et al. 2019a). To our best knowledge, this is the first work to study the isotropy of text representation for fine-tuning.

Spectral Control

The spectrum of representation has been studied in many subareas. On pre-trained word embeddings like Glove (Pennington, Socher, and Manning 2014), some work (Mu and Viswanath 2018) shows that removing the top one principal component leads to better performance on text similarity tasks. On contextual embeddings like Elmo (Peters et al. 2018) and GPT (Radford 2018), spectrum is used as a measure of the ability to capture contextual knowledge (Ethayarajh 2019). On text generation, it is shown that large eigenvalues hurt the expressiveness of representation and cause the degradation problem (Gao et al. 2019; Wang et al. 2020). In contrast, we study the spectrum of $[CLS]$ embedding in pre-trained language models, and show that controlling the top principal components can improve fine-tuning performance.

Conclusion

Our major contributions in this paper are two-fold:

- We study the isotropy of the pre-trained $[CLS]$ embeddings. Our analysis is based on straightforward visualization about standard deviation and correlation coefficient.
- The proposed regularization method, IsoBN, stably improves the fine-tuning of BERT and RoBERTa towards more isotropic representations, yielding an absolute increment around 1.0 point on 7 popular NLU tasks.

We hope our work points to interesting future research directions in improving pre-training language models as well as better fine-tuning towards more isotropy of PTLMs.

References

Arora, S.; Liang, Y.; and Ma, T. 2017. A Simple but Tough-to-Beat Baseline for Sentence Embeddings. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

Bjorck, J.; Gomes, C. P.; Selman, B.; and Weinberger, K. Q. 2018. Understanding Batch Normalization. In Bengio, S.; Wallach, H. M.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 7705–7716.

Cer, D.; Diab, M.; Agirre, E.; Lopez-Gazpio, I.; and Specia, L. 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)*, 1–14. Vancouver, Canada: Association for Computational Linguistics. doi:10.18653/v1/S17-2001.
- Clark, K.; Luong, M.-T.; Khandelwal, U.; Manning, C. D.; and Le, Q. V. 2019. BAM! Born-Again Multi-Task Networks for Natural Language Understanding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5931–5937. Florence, Italy: Association for Computational Linguistics. doi:10.18653/v1/P19-1595.
- Cogswell, M.; Ahmed, F.; Girshick, R. B.; Zitnick, L.; and Batra, D. 2016. Reducing Overfitting in Deep Networks by Decorrelating Representations. In Bengio, Y.; and LeCun, Y., eds., *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics. doi:10.18653/v1/N19-1423.
- Dodge, J.; Ilharco, G.; Schwartz, R.; Farhadi, A.; Hajishirzi, H.; and Smith, N. A. 2020. Fine-Tuning Pretrained Language Models: Weight Initializations, Data Orders, and Early Stopping. *ArXiv abs/2002.06305*.
- Ethayarajh, K. 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)*, 55–65. Hong Kong, China: Association for Computational Linguistics. doi:10.18653/v1/D19-1006.
- Gao, J.; He, D.; Tan, X.; Qin, T.; Wang, L.; and Liu, T. 2019. Representation Degeneration Problem in Training Natural Language Generation Models. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Ghorbani, B.; Krishnan, S.; and Xiao, Y. 2019. An Investigation into Neural Net Optimization via Hessian Eigenvalue Density. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 2232–2241. PMLR.
- Goyal, P.; Dollár, P.; Girshick, R. B.; Noordhuis, P.; Wesolowski, L.; Kyrola, A.; Tulloch, A.; Jia, Y.; and He, K. 2017. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. *ArXiv abs/1706.02677*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, 770–778. IEEE Computer Society. doi:10.1109/CVPR.2016.90.
- Howard, J.; and Ruder, S. 2018. Universal Language Model Fine-tuning for Text Classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 328–339. Melbourne, Australia: Association for Computational Linguistics. doi:10.18653/v1/P18-1031.
- Huang, L.; Yang, D.; Lang, B.; and Deng, J. 2018. Decorrelated Batch Normalization. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 791–800. IEEE Computer Society. doi:10.1109/CVPR.2018.00089.
- Huang, L.; Zhou, Y.; Zhu, F.; Liu, L.; and Shao, L. 2019. Iterative Normalization: Beyond Standardization Towards Efficient Whitening. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, 4874–4883. Computer Vision Foundation / IEEE. doi:10.1109/CVPR.2019.00501.
- Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In Bach, F. R.; and Blei, D. M., eds., *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, 448–456. JMLR.org.
- Jiang, H.; He, P.; Chen, W.; Liu, X.; Gao, J.; and Zhao, T. 2020. SMART: Robust and Efficient Fine-Tuning for Pre-trained Natural Language Models through Principled Regularized Optimization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2177–2190. Online: Association for Computational Linguistics. doi:10.18653/v1/2020.acl-main.197.
- Liu, X.; He, P.; Chen, W.; and Gao, J. 2019a. Multi-Task Deep Neural Networks for Natural Language Understanding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4487–4496. Florence, Italy: Association for Computational Linguistics. doi:10.18653/v1/P19-1441.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M. S.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L. S.; and Stoyanov, V. 2019b. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *ArXiv abs/1907.11692*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Montavon, G.; and Müller, K.-R. 2012. Deep Boltzmann Machines and the Centering Trick. In *Neural Networks: Tricks of the Trade*.
- Mu, J.; and Viswanath, P. 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.

- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12: 2825–2830.
- Pennington, J.; Socher, R.; and Manning, C. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. Doha, Qatar: Association for Computational Linguistics. doi:10.3115/v1/D14-1162.
- Peters, M.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 2227–2237. New Orleans, Louisiana: Association for Computational Linguistics. doi:10.18653/v1/N18-1202.
- Peters, M. E.; Ruder, S.; and Smith, N. A. 2019. To Tune or Not to Tune? Adapting Pretrained Representations to Diverse Tasks. In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, 7–14. Florence, Italy: Association for Computational Linguistics. doi:10.18653/v1/W19-4302.
- Radford, A. 2018. Improving Language Understanding by Generative Pre-Training.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2383–2392. Austin, Texas: Association for Computational Linguistics. doi:10.18653/v1/D16-1264.
- Rosipal, R.; Girolami, M.; Trejo, L. J.; and Cichocki, A. 2001. Kernel PCA for feature extraction and de-noising in nonlinear regression. *Neural Computing & Applications* 10(3): 231–243.
- Santurkar, S.; Tsipras, D.; Ilyas, A.; and Madry, A. 2018. How Does Batch Normalization Help Optimization? In Bengio, S.; Wallach, H. M.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 2488–2498.
- Schraudolph, N. N. 1998. Accelerated Gradient Descent by Factor-Centering Decomposition.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A.; and Potts, C. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 1631–1642. Seattle, Washington, USA: Association for Computational Linguistics.
- Szegedy, C.; Ioffe, S.; Vanhoucke, V.; and Alemi, A. A. 2017. Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. In Singh, S. P.; and Markovitch, S., eds., *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, 4278–4284. AAAI Press.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2019a. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2019b. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Wang, L.; Huang, J.; Huang, K.; Hu, Z.; Wang, G.; and Gu, Q. 2020. Improving Neural Language Generation with Spectrum Control. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Williams, A.; Nangia, N.; and Bowman, S. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1112–1122. New Orleans, Louisiana: Association for Computational Linguistics. doi:10.18653/v1/N18-1101.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; and Brew, J. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv abs/1910.03771*.
- Zhu, C.; Cheng, Y.; Gan, Z.; Sun, S.; Goldstein, T.; and Liu, J. 2020. FreeLB: Enhanced Adversarial Training for Natural Language Understanding. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.