

Better Pre-Training by Reducing Representation Confusion

Haojie Zhang^{1,2*}, Mingfei Liang^{1*}, Ruobing Xie^{1*}, Zhenlong Sun¹, Bo Zhang¹, Leyu Lin¹

¹WeChat Search Application Department, Tencent, China

²Peking University, China

¹{coldhjzhang, aesopliang, ruobingxie, richardsun, nevinzhang, goshawklin}@tencent.com

²zhanghaojie@stu.pku.edu.cn

Abstract

In this work, we revisit the Transformer-based pre-trained language models and identify two different types of information confusion in position encoding and model representations, respectively. Firstly, we show that in the relative position encoding, the joint modeling about relative distances and directions brings confusion between two heterogeneous information. It may make the model unable to capture the associative semantics of the same distance and the opposite directions, which in turn affects the performance of downstream tasks. Secondly, we notice the BERT with Mask Language Modeling (MLM) pre-training objective outputs similar token representations (last hidden states of different tokens) and head representations (attention weights¹ of different heads), which may make the diversity of information expressed by different tokens and heads limited. Motivated by the above investigation, we propose two novel techniques to improve pre-trained language models: Decoupled Directional Relative Position (DDRP) encoding and MTH² pre-training objective. DDRP decouples the relative distance features and the directional features in classical relative position encoding. MTH applies two novel auxiliary regularizers besides MLM to enlarge the dissimilarities between (a) last hidden states of different tokens, and (b) attention weights of different heads. These designs allow the model to capture different categories of information more clearly, as a way to alleviate information confusion in representation learning for better optimization. Extensive experiments and ablation studies on GLUE benchmark demonstrate the effectiveness of our proposed methods.

*Equal contribution.

¹"attention weights" mainly refer to the dot product between Key and Query in the self-attention module.

²MTH is the abbreviation of our proposed MLM with Token Cosine Differentiation (TCD) and Head Cosine Differentiation (HCD) pre-training task. TCD and HCD are described in detail in sec. 1(2) and sec.3.2.

1 Introduction

The paradigm of pre-training on large-scale corpus and fine-tuning on specific task datasets has swept the entire field of Natural Language Processing (NLP). BERT (Devlin et al., 2018) is the most prominent pre-trained language model, which stacks the encoder blocks of Transformer (Vaswani et al., 2017) and adopts MLM and Next Sentence Prediction (NSP) pre-training tasks, achieving the SOTA results in 2018. After that, a large number of Pre-trained Language Models (PLMs) (Liu et al., 2019; Lan et al., 2020; Raffel et al., 2019; Clark et al., 2020; He et al., 2021) that optimize the Transformer structure and pre-training objectives have emerged, which further improves the performance of the pre-trained language models on multiple downstream tasks. In this work, we identify two different types of information confusion in language pre-training, and explore two conceptually simple and empirically powerful techniques against them as follows:

(1) **Decoupled Directional Relative Position (DDRP) Encoding.** It is well known that relative position encoding is competitive and has been widely used in real PLMs (Shaw et al., 2018; Yang et al., 2019; Wei et al., 2019; Raffel et al., 2019; Su et al., 2021; He et al., 2021; Ke et al., 2021). Despite its great performance, we still notice relative position encoding methods utilizes completely separate parametric vectors to encode different relative position information, which indicates that every single parametric vector needs to learn both distance and directional features. We consider this paradigm of utilizing a single parametric vector to represent both relative distance and direction as a kind of information confusion, and question its rationality. Since relative distance features and the directional features are apparently heterogeneous information that reflects different aspects of positional information, we argue that existing methods may impose

difficult in establishing connections explicitly between parametric vectors of the same distances and the opposite directions, which in turn result in serious information losses in position encoding. Inspired by this, we propose a novel Decoupled Directional Relative Position (DDRP) encoding. In detail, DDRP decomposes the classical relative position embedding (Shaw et al., 2018) into two embeddings, one storing the relative distance features and the other storing the directional features, and then multiply the two together explicitly to derive the final decoupled relative position embedding, allowing originally confused distance and directional information to be as distinguishable as possible.

(2) **Model Representation Differentiations.** We analyze that there is non-negligible confusion in the representation of pre-trained BERT, as evidenced by the high consistency in last hidden states across different tokens and attention weights across different heads, respectively. Similar last hidden states will introduce the *anisotropic problem* (Mimno and Thompson, 2017), which will bound the token vectors to a narrow representation space and thus make it more difficult for the model to capture deep semantics. Considering attention weights contain rich linguistic knowledge (Clark et al., 2019; Jawahar et al., 2019), we argue that high consistency in attention weights also constrains the ability of the model to capture multi-aspect information. To alleviate the representation confusion between different tokens and heads caused by high information overlap, we propose two novel pre-training approaches to stimulate the potential of the pre-trained model to learn rich linguistic knowledge: Token Cosine Differentiation (TCD) objective and Head Cosine Differentiation (HCD) objective. Specifically, TCD attempts to broaden the dissimilarity between tokens by minimizing the cosine similarities between different last hidden states. In contrast, HCD attempts to broaden the dissimilarity between heads by minimizing the cosine similarities between different attention weights. We apply TCD and HCD as two auxiliary regularizers in MLM pre-training, which in turn guides the model to produce more discriminative token representations and head representations. Formally, we define our enhanced pre-training task as MLM with TCD and HCD (MTH).

Extensive experiments on the GLUE benchmark show that DDRP achieves better results than classi-

cal relative position encoding (Shaw et al., 2018) on almost all tasks without introducing the additional computational overhead and consistently outperforms prior competitive relative position encoding models (He et al., 2021; Ke et al., 2021). Moreover, our proposed MTH outperforms MLM by a 0.96 average GLUE score and achieves nearly 2x pre-training speedup on BERT_{BASE}. Both DDRP and MTH are straightforward, effective, and easy to deploy, which can be easily combined with existing pre-training objectives and various model structures. Our contributions are summarized as follows:

- We propose a novel relative position encoding named DDRP, which decouples the relative distance and directional features, giving the model a stronger prior knowledge, fewer parameters, and better results compared to conventional coupled position encodings.
- We analyze the trend of self-similarity of last hidden states and attention weights during pre-training, and propose two novel Token Cosine Differentiation and Head Cosine Differentiation objectives, motivating pre-trained Transformer to better capture semantics in PLMs.
- We experimentally verified by our proposed techniques (DDRP and MTH) that decomposing heterogeneous information and extending representation diversity can significantly improve pre-trained language models. We also analyze the characteristics of DDRP and MTH in detail.

2 Related Work

In recent years, pre-trained language models have made significant breakthroughs in the field of NLP. BERT (Devlin et al., 2018), which proposes MLM and NSP pre-training objectives, is pre-trained on large-scale unlabeled corpus and has learned bidirectional representations efficiently. After that, many different pre-trained models are produced, which further improve the effectiveness of the pre-trained models. RoBERTa (Liu et al., 2019) proposes to remove the NSP task and verifies through experiments that more training steps and larger batches can effectively improve the performance of the downstream tasks. ALBERT (Lan et al., 2020) proposes a Cross-Layer Parameter Sharing technique to lower memory consumption. XL-Net

(Yang et al., 2019) proposes Permutation Language Modeling to capture the dependencies among predicted tokens. ELECTRA (Clark et al., 2020) adopts Replaced Token Detection (RTD) objective, which considers the loss of all tokens instead of a subset. TUPE (Ke et al., 2021) performs Query-Key dot product with different parameter projections for contextual information and positional information separately and then added them up, they also add relative position biases like T5 (Raffel et al., 2019) on different heads to form the final correlation matrix. DEBERTA (He et al., 2021) separately encodes the context and position information of each token and uses the textual and positional disentangled matrices of the words to calculate the correlation matrix.

3 Method

In this section, we analyze in turn two different types of information confusion that exist in the real PLMs: (i) The paradigm of utilizing a single parametric vector of relative position embedding to represent both relative distance and direction. (ii) The high similarity and overlap in model representations. Based on above two investigations, we propose two techniques, **Decoupled Directional Relative Position (DDRP)** Encoding and **MLM with TCD and HCD (MTH)**, respectively, to help the PLMs alleviate information confusion and enhance representation clarity and diversity.

3.1 Decoupled Directional Relative Position (DDRP) Encoding

We first start to introduce DDRP by formulating multi-head attention module of BERT and BERT-R (Shaw et al., 2018). Specifically, BERT formulates multi-head attention for a specific head as follows:

$$Q = HW^Q, K = HW^K, V = HW^V, \quad (1)$$

$$A = \frac{QK^T}{\sqrt{d}}, \quad (2)$$

$$Z = \text{softmax}(A)V, \quad (3)$$

where $H \in R^{S \times D}$ represents the input hidden states; $W^Q, W^K, W^V \in R^{D \times d}$ represent the projection matrix of Query, Key, and Value respectively; $A \in R^{S \times S}$ represents attention weight; $Z \in R^{S \times d}$ represents the single-head output hidden states of self-attention module; S represents input sequence length; D represents the dimension of input hidden states; d represents the dimension of single-head hidden states. Unlike BERT, which

adds the absolute position embedding to the word embedding as the final input of the model, BERT-R first applies relative position encoding. It adds relative position embedding into K in the self-attention module of each layer to make a more interactive influence. Its formulations are as follows:

$$A_{i,j} = \frac{Q_i (K_j + K_{\sigma(i,j)}^r)^T}{\sqrt{d}}, \quad (4)$$

$$\sigma(i,j) = \text{clip}(i-j) + r_s, \quad (5)$$

where Q_i represents Query vector at the i -th position; K_j represents Key vector at the j -th position; r_s represents maximum relative position distance; $\sigma(i,j)$ represents the index of relative position embedding $K^r \in R^{2r_s \times d}$; relative position embedding for K are shared at all different heads. Note that Shaw et al. (2018) has experimentally demonstrated that adding relative position embedding to the interaction between A and V gives no further improvement in effectiveness, so the relative position embedding in V space is eliminated in all our experiments to reduce the computational overhead.

Compared with BERT, BERT-R models the correlation between words and positions more explicitly, and thus further expands the expression diversity between words. However, we notice that in BERT-R, the vectors from the same distance on both left and right sides are encoded in isolation (as shown in Figure 1(a)), which indicates that every single parametric vector from K^r is forced to maintain distance and direction, two different types of information. Since it is confirmed that directional information is crucial in language modeling (Vu et al., 2016; Fuller, 2002; Shen et al., 2018), we argue that such an approach causes unnecessary information confusion and faces several constraints: (i) Mixing relative distance and directional information for modeling makes information originally in different spaces entangled, which in turn makes the learning of parametric vectors more difficult. (ii) Dot products between word vectors and directionally confused positional vectors bring unnecessary randomness in deep bidirectional representation models.

To alleviate the confusion of distance and direction that exists in BERT-R and allow the model to perceive distances and directions more clearly, we propose a novel Decoupled Directional Relative Position (DDRP) encoding. Specifically, DDRP decouples the relative distance and directional information and maintains them with two different

embeddings. Its formula is as follows:

$$A_{i,j} = \frac{Q_i (K_j + K_{\delta(i,j)}^d)^T}{\sqrt{d}}, \quad (9)$$

$$K_{\delta(i,j)}^d = D_{\rho(i,j)} \odot K_{\delta(i,j)}^{rd}, \quad (10)$$

$$\rho(i,j) = \begin{cases} 1, & \text{if } i-j < 0 \\ 0, & \text{if } i-j = 0 \\ 2, & \text{if } i-j > 0 \end{cases}, \quad (11)$$

$$\delta(i,j) = \text{abs}(\text{clip}(i-j)), \quad (12)$$

where $\rho(i,j)$ represents the index of directional embedding $D \in R^{3 \times d}$; $\delta(i,j)$ represents the index of relative distance embedding $K^{rd} \in R^{r_s \times d}$; K^d represents the relative position matrix. Note that in terms of implementation details, the only difference between DDRP and BERT-R is that DDRP decouples K^r in BERT-R into the element-wise multiplication of D and K^{rd} . We also provide a specific comparison example in Figure 1.

Compared to previous relative position encodings, we summarize the advantage of DDRP as follows: (i) DDRP explicitly extracts the commonalities (relative distances) and differences (directions) in the positional information, leading the model to produce attention that better match the real semantic distributions, which reduces the difficulty of model learning and unlocks the potential of the model. (ii) DDRP compresses the total number of parametric vectors from $2r_s$ to $r_s + 3$.




Figure 1: Fig.1(a) represents the classical relative position matrix; Fig.1(b) represents the decoupled relative position matrices we proposed. Note that the parametric vectors of the same color have the same values.

3.2 Model Representation Differentiations

Token representations. Isotropic distributions have been proved theoretically to be beneficial to token representations, which ensures that the different token vectors are directional uniform, thus maximizing the diversity of token representations. (Mimno and Thompson, 2017). In practice, Mu et al. (2017) have also empirically confirmed the effectiveness of isotropic distributions on static token representations, such as WORD2VEC (Mikolov et al., 2013) and GLOVE (Pennington et al., 2014).

Inspired by the above studies, we also wonder whether contextualized token representations (e.g., last hidden states of BERT) are isotropic. Following Mimno and Thompson (2017), we utilize the cosine similarity to evaluate the degree of isotropy in token representations. The higher similarity, the smaller isotropy; the lower similarity, the greater isotropy. For an input sequence $S = [x_1, \dots, x_n]$, we formulate the last hidden states' average self-similarity as follows:

$$f(S) = \frac{2}{n(n-1)} \sum_{i=1}^{n-1} \sum_{j=i+1}^n \cos(h_i, h_j), \quad (15)$$

where h_i and h_j are the last hidden states of x_i and x_j ; \cos represents cosine similarity.

Head representations. Multi-head attention, is aimed at capturing information in different heterogeneous subspaces and has been experimentally verified different heads correspond well to different linguistic notions (Clark et al., 2019). However, some studies point out that some heads contribute almost nothing to downstream tasks (Kovaleva et al., 2019; Michel et al., 2019; Voita et al., 2019; Correia et al., 2019). We are surprised by this, and speculate that the above problem may be caused by the heavy overlap of information that some heads are concerned about. To verify our point, we utilize cosine similarity to evaluate the degree of overlap in head representations, following token representations. The higher similarity, the higher overlap; the lower similarity, the lower overlap. For multiple heads $H = [H_1^1, \dots, H_m^1, \dots, H_1^L, \dots, H_m^L]$, we formulate the attention weights' average self-similarity as follows:

$$f(H) = \frac{2}{Lm(m-1)} \sum_{l=1}^L \sum_{i=1}^{m-1} \sum_{j=i+1}^m \cos(a_i^l, a_j^l), \quad (16)$$

where L represents the number of Transformer layers; a_i^l and a_j^l are the attention weights of the i -th head and j -th head of the l -th layer.

Analysis on the similarities between different tokens and heads. With curiosity about the similarity of token representations and head representations, we analyze the self-similarity trends of tokens and heads during the original MLM BERT pre-training. Specifically, we sample 5,000 sentences from the validation set and evaluate the average self-similarity of last hidden states and attention weights under multiple checkpoints during the pre-training stage as shown in Figure 2. We




Figure 2: Average self-similarity of last hidden states and attention weights during MLM pre-training.

can notice that although $f(S)$ and $f(H)$ decrease at the beginning of pre-training, soon they start to rise gradually until the end of the training, and the similarities are always high throughout the training process. These fully demonstrates that in the MLM-based BERT, the overlap between different tokens and heads is strong and information confusion in model representations has become a problem worth to be solved.

Training objectives. To guide the model to produce more discriminative token representations and head representations, we propose a novel MTH pre-training objective, which combines original MLM with two novel Token Cosine Differentiation (TCD) objective and Head Cosine Differentiation (HCD) objective. Specifically, MTH applies the average cosine self-similarity of sampled last hidden states and attention weights as two auxiliary pre-training regularizers besides MLM³. For an input sequence $S = [x_1, \dots, x_n]$, TCD samples n' ($n' \leq n$) tokens uniformly in sequence order to obtain a subsequence $\tilde{S} = [\tilde{x}_1, \dots, \tilde{x}_{n'}]$ and calculates the average cosine self-similarity of the subsequence’s last hidden states as follows:

$$\mathcal{L}_{TCD} = \frac{2}{n'(n'-1)} \sum_{i=1}^{n'-1} \sum_{j=i+1}^{n'} \cos(\tilde{h}_i, \tilde{h}_j), \quad (17)$$

where \tilde{h}_i and \tilde{h}_j are the last hidden states of \tilde{x}_i and \tilde{x}_j . For multiple heads $H^l = [H_1^l, \dots, H_m^l]$ of a specific layer l , HCD randomly

³We have empirically verified that this below sampling strategy can greatly reduce the computational overhead with only a slight performance drop, comparing with regularizing all tokens and heads. In practice, we notice that setting $n' = 50$, $m' = 2$ is fine on both BERT and DDRP, which will compress the additional computational overhead from about 30% to 4%.

samples m' ($m' \leq m$) different heads $\tilde{H}^l = [\tilde{H}_1^l, \dots, \tilde{H}_{m'}^l]$ (Note that HCD samples by layers, so sampled heads may be different across different layers.) and then calculate the average cosine self-similarity of attention weights of sampled heads as follows:

$$\mathcal{L}_{HCD} = \frac{2}{Lm'(m'-1)} \sum_{l=1}^L \sum_{i=1}^{m'-1} \sum_{j=i+1}^{m'} \cos(\tilde{a}_i^l, \tilde{a}_j^l), \quad (18)$$

\tilde{a}_i^l and \tilde{a}_j^l are the attention weights of the i -th head and j -th head in the sampled headset of the l -th layer. Ultimately, we define the global pre-training objective MTH as follows:

$$\mathcal{L}_{MTH} = \mathcal{L}_{MLM} + \alpha_1 \mathcal{L}_{TCD} + \alpha_2 \mathcal{L}_{HCD}, \quad (19)$$

where α_1 and α_2 are hyperparameters.

4 Experiments

4.1 Pre-training Text Corpora

Follow Devlin et al. (2018), we use Wikipedia and BooksCorpus (Zhu et al., 2015), a roughly 16G uncompressed text corpus for pre-training.

4.2 Baselines

We compare DDRP with competitive pre-trained models. BERT (Devlin et al., 2018) equips Transformer (Vaswani et al., 2017) with parametric absolute position encoding. BERT-R uses the relative position encoding proposed by Shaw et al. (2018), which couples relative distance information and directional information for modeling. TUPE (Ke et al., 2021) performs Query-Key dot product with different parameter projections for contextual information and positional information separately and then adds them up, plus the relative position biases like T5 (Raffel et al., 2019). DEBERTA (He et al., 2021) uses two vectors to encode content and position and uses disentangled matrices on their contents and relative positions respectively to compute the attention weights among words.

4.3 Experimental Settings

Following the previous practice, we use a base-size model for training, which consists of 12 Transformer encoder layers, each containing 12 heads with an input dimension of 768. During pre-training, we directly use the maximum training length of 512 without taking any form of random

Models	Steps	RTE	STS-B	MRPC	CoLA	SST-2	QNLI	QQP	MNLI	Avg.
BERT (MLM)	1M	70.75	89.66	87.50	59.65	92.20	91.23	91.00	84.33	83.29
BERT-R (MLM)	1M	71.84	89.68	87.99	60.82	92.66	91.54	91.13	85.45	83.89
TUPE (MLM)	1M	68.59	89.61	86.02	<u>62.82</u>	92.66	91.26	91.04	84.88	83.36
DEBERTA (MLM)	1M	<u>73.28</u>	89.14	87.99	60.60	92.66	92.14	91.00	85.93	84.09
DDRP (MLM)	1M	72.20	<u>90.01</u>	<u>88.25</u>	<u>62.82</u>	<u>92.41</u>	92.31	91.24	<u>86.02</u>	<u>84.41</u>
DDRP (MTH)	1M	75.09	90.41	88.72	63.36	92.66	<u>92.24</u>	<u>91.22</u>	86.22	85.00

Table 1: Results on the development set of the GLUE benchmark for base-size pre-trained models. The best results are bolded, and the second results are underlined.

Approaches	Steps	RTE	STS-B	MRPC	CoLA	SST-2	QNLI	QQP	MNLI	Avg.
BERT (MLM)	500k	68.23	88.92	86.74	57.05	91.97	90.41	90.74	83.41	82.18
BERT (MTH)	500k	71.84	89.41	86.76	61.40	92.08	90.59	90.76	83.61	83.31
BERT (MLM)	1M	70.75	89.66	87.50	59.65	92.20	91.23	91.00	84.33	83.29
BERT (MTH)	1M	73.64	90.16	88.48	62.31	92.43	91.21	91.12	84.67	84.25
DDRP (MLM)	1M	72.20	90.01	88.25	62.82	92.41	92.31	91.24	86.02	84.41
DDRP (MTH)	1M	75.09	90.41	88.72	63.36	92.66	92.24	91.22	86.22	85.00

Table 2: Development scores on GLUE benchmark. BERT (MLM/MTH) represents pre-trained BERT_{BASE} with MLM/MTH pre-training objective. DDRP (MLM/MTH) represents pre-trained DDRP_{BASE} with MLM/MTH pre-training objective.

injection, and for examples less than 512 in length, we do not use the next document for padding. We remove Next Sentence Prediction (NSP) task and only keep Masked LM (MLM) as our pre-training task for all models unless noted otherwise. Considering that shorter documents may be missing semantics, we discard documents of length less than 8. We adopt the whole word masking strategy and split the whole words longer than 4 into individual subtokens. Following Devlin et al. (2018), we set the batch size to 256 sequences, the peak learning rate to $1e-4$, and the training steps to 1M. We grid search α_1 and α_2 of TCD and HCD in $\{0.01, 0.1, 1.0\}$. Eventually, we set $\alpha_1 = 1.0$ for TCD and set $\alpha_2 = 0.01$ for HCD. All the models are implemented based on the code practice of BERT⁴ in Tensorflow. We conduct all experiments on 16 Tesla-V100 GPUs (32G). All the pre-training hyperparameters are supplemented in Appendix A. To make a fair comparison, we implement BERT, BERT-R, TUPE, DEBERTA, and DDRP⁵ with the same pre-training hyperparameters and model configurations, which are consistent with vanilla BERT.

4.4 Results on GLUE Benchmark

We evaluate models on eight different English understanding tasks from General Language Under-

standing Evaluation (GLUE) benchmark (Wang et al., 2019). The datasets cover four types of tasks: natural language inference (RTE, QNLI, MNLI), paraphrase detection (MRPC, QQP), linguistic acceptability (CoLA), and sentiment classification (SST-2). For all experiments, STS-B and CoLA are reported by Pearson correlation coefficient and Matthews correlation coefficient, and other tasks are reported by Accuracy. All the fine-tuning hyperparameter configurations can be found in Appendix B. Following Ke et al. (2021), we fine-tune with five random seeds and report the median results.

4.4.1 Comparing Prior Competitive Models with DDRP

The overall comparison results are shown in Table 1. Firstly, we can notice that all the various relative position encoding models perform better than BERT, which proves that relative position encoding is a more competitive approach to encode position information. Sencodly, it is easy to find that DDRP outperforms all the strong baselines, which demonstrates modeling relative position encoding by clarifying the originally confused relative distance and directional information more clearly is more effective. Thirdly, DDRP pre-trained with MTH can consistently outperform BERT-R/DEBERTA by a 1.11/0.91 average GLUE score, which indicates that DDRP can be effectively compatible with better pre-training objectives to perform stronger.

⁴<https://github.com/google-research/bert>

⁵Following Shaw et al. (2018) and Raffel et al. (2019), we set $r_s = 64$ for all the relative position encoding models.

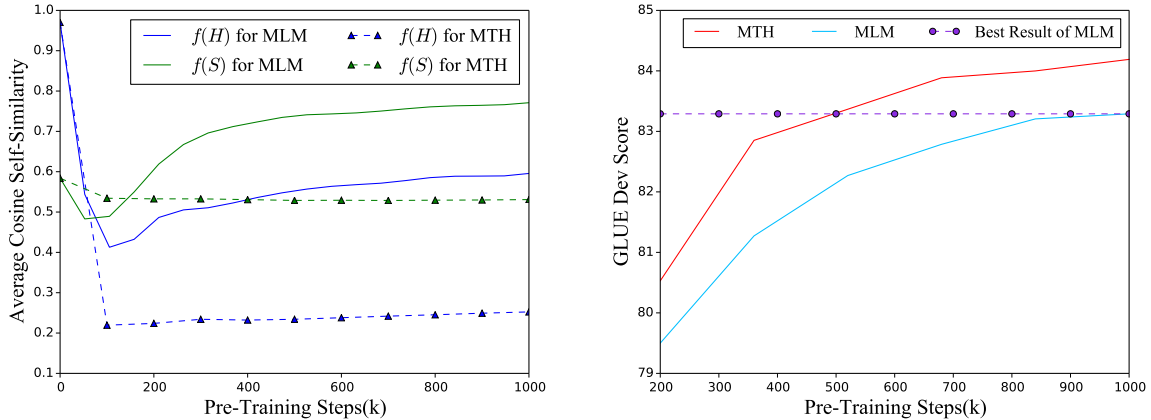


Figure 3: The left figure (a) represents the trend of average cosine self-similarity of token representations and head representations during pre-training. The right figure (b) represents the trend of GLUE average score during pre-training.

Moreover, compared to BERT-R, DDRP introduces nothing in complexity while DEBERTA increases the computational cost about 25%, we also consider DDRP is a more time-efficient alternative than the recent state-of-the-art model DEBERTA (as analyzed in Sec.5.4).

4.4.2 Comparing MTH with MLM

As illustrated in Table 2, BERT (MTH) outperforms BERT (MLM) by a 0.96 average GLUE score and is consistently better on 7 out of 8 tasks. When combining MTH with strong DDRP, it still brings an improvement by 0.59 GLUE average score. Notably, BERT (MTH) can achieve better results compared with well-trained BERT (MLM) while only using 50% training steps. Since MTH utilizes cosine similarity and sampling strategy for the penalty, only a very slight computational cost is introduced. All the above statistics can fully verify that decreasing the similarity in model representations can effectively alleviate information overlap and increase representation diversity, which in turn leads to consistent and stable improvement across different model structures.

5 Analysis and Discussion

5.1 Ablation Studies

Effect of DDRP. As shown in Table 1, BERT-R outperforms BERT by 0.6 points on average. Based on BERT-R, our proposed DDRP outperforms BERT-R by 0.52 points averagely without imposing additional computational costs. It is worth noting that compared to BERT-R, DDRP helps a great deal on low-resource tasks, such as CoLA, while further

improving the performance on high-resource tasks, such as QNLI and MNLI. These fully demonstrate utilizing separate parametric vectors to represent distances and directions, two apparently heterogeneous information, can be beneficial to the model, and further justifies the value of dissociating confusing information that is confounded in similar spaces.

Effect of TCD and HCD. MTH brings in two additional TCD and HCD regularizers besides the original MLM task. To further evaluate the relative contributions of the HCD and TCD, we develop one variation, which is BERT pre-trained with MLM and TCD. Table 3 summarizes the results on the base-size models. Firstly, it shows a 0.42 average GLUE score drop when HCD is removed from MTH, especially on MRPC, CoLA, and MNLI. Secondly, there is a 0.54 average GLUE score drop when TCD is progressively removed, especially on RTE, STS-B, and CoLA. These results indicate that both TCD and HCD regularizers play a crucial role in improving performance.

5.2 Analysis on MTH

To further understand why MTH works, we compare MLM and MTH in terms of average self-similarity of token representations and head representations and performance during pre-training in Figure 3. As shown in Figure 3.(a), it is easy to find that MTH’s average self-similarity is much lower than MLMs’. We can also clearly notice from Figure 3.(b) that the average GLUE score of MTH is always about one point higher than MLM’s during the whole pre-training process. These confirm that

Model	RTE	STS-B	MRPC	CoLA	SST-2	QNLI	QQP	MNLI	Avg.
MTH	73.64	90.16	88.48	62.31	92.43	91.21	91.12	84.67	84.25
w/o HCD	73.28	90.41	87.25	60.85	92.23	91.37	91.00	84.22	83.83
w/o TCD&HCD	70.75	89.66	87.50	59.65	92.20	91.23	91.00	84.33	83.29

Table 3: Ablation study for MTH. Note that MTH (w/o TCD&HCD) equals simply using MLM in pre-training.

(i) the differentiation of tokens and heads is important for model optimization; (ii) MTH can help to produce more discriminative token representations and head representations, extend the representation space of tokens and heads, and thus improve the performance.

5.3 Analysis on DDRP

In this subsection, we intend to analyze the attention maps of DDRP as a way to investigate why making only slight modifications on BERT-R can bring great gains. To better examine and explain the ability of DDRP to capture information in both left and right directions, we divide multiple heads into two groups evenly, where group 1 consists of the heads that focus most on the right side, and group 2 consists of the heads that focus most on the left side. As shown in Figure 4, it is easy to observe a distinct upper triangle effect in group 1 and a distinct lower triangle effect in group 2, which indicates that DDRP may allow the model to be more precise in the perception of direction, a piece of information that is crucial to understanding semantics. To further confirm our point, we sample 5,000 sentences from the validation set and count the percentage of sentences with upper and lower triangular effects according to Algorithm 1 (more details can be seen in Appendix C). It is observed that 92.11% of sentences have an up-down triangle effect. We also count the percentage for BERT-R with the same process and observe only 78.94% of the sentences have an up-down triangle effect. All the phenomena and statistics fully reveal that DDRP can make different heads focus on the token information interaction in different directions and reduce confusion between heads, thus improving the effectiveness and rationality of the model.

5.4 Complexity Analyses

DDRP. Compared with BERT, DDRP introduces additional parameters: $D \in R^{3 \times d}$ and $K^{rd} \in R^{r_s \times d}$. The total increase in parameters is $3 \times d + r_s \times d$. For base-size model ($D = 768, L = 12, S = 512, N = 12, d = 64$)⁶, the total increase

⁶N is the number of head.




Figure 4: Attention visualization for a sampled batch of sentences. From left to right is the attention visualization for group 1, group 2, and global, respectively.

amounts to 0.0043M, which is negligible. Compared with BERT, the additional computational complexity for both BERT-R and DDRP is $O(SD)$. Since DEBERTA equips different heads with unshared K^r s, the additional computational complexity for DEBERTA is $O(NSD)$. Overall, BERT-R and DDRP increase the computational cost about 5%, and DEBERTA increases the computational cost about 30%. Although DDRP introduces a slight computational cost compared to BERT, it is more time-efficient than DEBERTA and outperforms all the above models.

MTH. Since the two regularizers of TCD and HCD are based on cosine similarity and sampling strategy, they do not introduce too much computational cost. Compared with MLM, MTH only increases a slight computational cost of about 4% while bringing excellent improvement.

6 Conclusion

In this work, we analyze and identify potential information confusion in the relative position encoding and model representations, respectively, and design two novel techniques to address these problems: DDRP (Decoupled Directional Relative Position) encoding and MTH (MLM with TCD and HCD) pre-training objectives. Specifically, DDRP decouples relative distance features and directional features to eliminate unnecessary randomness in the self-attention module. MTH utilize TCD and HCD as two regularizers to supervise the model to always maintain a certain level of critical thinking. The experimental results show that DDRP achieves better performance compared with various relative

position encoding models and MTH outperforms MLM by a large margin. We believe that reducing information confusion in representation learning may have broader application scenarios, and leave this area of exploration for future work.

7 Limitations

Our limitations lie in inducing additional computational costs. Compared with BERT, the additional computational complexity for DDRP is $O(SD)^7$, which is reflected in the 5% increase in computational cost. Compared with MLM, MTH with sampling strategy increases the computational cost by about 4%.

References

- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019. What does bert look at? an analysis of bert’s attention. In *BlackBoxNLP*.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. In *ICLR*.
- Gonçalo M. Correia, Vlad Niculae, and André F. T. Martins. 2019. Adaptively sparse transformers. In *EMNLP-IJCNLP*.
- Graham Neubig Devendra Singh Sachan. 2018. Parameter sharing methods for multilingual self-attentional translation models. In *WMT*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Gillian Fuller. 2002. The arrow-directional semiotics: Wayfinding in transit. *Social semiotics*.
- J Pino X Li H Gong, Y Tang. 2021. Pay better attention to attention: Head selection in multilingual and multi-domain sequence modeling. In *NeurIPS*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. Deberta: Decoding-enhanced bert with disentangled attention. In *ICLR*.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does bert learn about the structure of language? In *ACL*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. Spanbert: Improving pre-training by representing and predicting spans. *TACL*, 8:64–77.
- Guolin Ke, Di He, and Tie-Yan Liu. 2021. Rethinking positional encoding in language pre-training. In *ICLR*.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of bert. In *EMNLP*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations. In *ICLR*.
- 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 preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2018. Fixing weight decay regularization in adam.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? In *NeurIPS*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- David Mimno and Laure Thompson. 2017. The strange geometry of skip-gram with negative sampling. In *EMNLP*.
- Jiaqi Mu, Suma Bhat, and Pramod Viswanath. 2017. All-but-the-top: Simple and effective postprocessing for word representations. *arXiv preprint arXiv:1702.01417*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. In *NAACL*.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, Shirui Pan, and Chengqi Zhang. 2018. Disan: Directional self-attention network for rnn/cnn-free language understanding. In *AAAI*.
- Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2021. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*.

⁷Here, S is the input sequence length and D is the dimension of token representations.

- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *ACL*.
- Ngoc Thang Vu, Pankaj Gupta, Heike Adel, and Hinrich Schütze. 2016. Bi-directional recurrent neural network with ranking loss for spoken language understanding. In *ICASSP*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *ICLR*.
- Junqiu Wei, Xiaozhe Ren, Xiaoguang Li, Wenyong Huang, Yi Liao, Yasheng Wang, Jiashu Lin, Xin Jiang, Xiao Chen, and Qun Liu. 2019. Nezha: Neural contextualized representation for chinese language understanding. *arXiv preprint arXiv:1909.00204*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *ICCV*.

A Appendix A. Hyperparameters for Pre-Training

As shown in Table 4, we list the pre-training hyperparameter configurations. To make a fair comparison, all models’ pre-training hyperparameter configurations in our experiments are identical to vanilla BERT (Devlin et al., 2018).

Hyperparameter	
Vocab size	3,0522
Hidden size	768
Attention heads	12
Layers	12
Training steps	1M
Warmup ratio	0.01
Batch size	256
Learning rate	1e-4
Adam ϵ	1e-6
Adam (β_1, β_2)	(0.9, 0.999)
Learning rate schedule	linear
Weight decay	0.01
Clip norm	1.0
Dropout	0.1

Table 4: Hyperparameters used for pre-trained models.

B Appendix B. Hyperparameters for Fine-Tuning

As shown in Table 5, we enumerate the hyperparameter configurations to fine-tune the tasks on the GLUE benchmark (Wang et al., 2019). We grid search these fine-tuning hyperparameter configurations for all models. Following the BERT, we do not show results on the WNLI GLUE task for the Dev set results.

Hyperparameter	GLUE
Batch size	{16, 32}
Learning rate	{ $1e-5$, $2e-5$, $3e-5$ }
Epoch	{4, 6}
Adam ϵ	1e-6
Adam (β_1, β_2)	(0.9, 0.999)
Learning rate schedule	linear
Weight decay	0.01
Clip norm	1.0
Dropout	0.1
Warmup ratio	0.1

Table 5: Hyperparameters used for fine-tuning on the GLUE benchmark.

C Appendix C. Details for Up-Down Triangle Effect

Here we provide more details for the up-down triangle effect (in Sec. 5.3). It is rather difficult and non-intuitive to analyze the directional information in 12 different attention heads. Since previous studies have considered to group multiple heads in the self-attention module (Devendra Singh Sachan, 2018; H Gong, 2021), we thereby attempt to divide the heads into two groups evenly. Specifically, we divide the heads that focus more on the right side information into group 1 and the heads that focus more on the left side information into group 2, wishing to reveal the directional information encoded in attention weights more explicitly. We experimentally verified that this grouping approach could provide a better presentation of the directional information. Therefore, we combined this approach to conduct a comparative analysis of DDRP and BERT-R to demonstrate the more powerful directional perception of DDRP.

As illustrated in Figure 4, group 1 is more focused on the right information (greater attention values in the upper triangle region) and group 2 is more focused on the left information (greater attention values in the lower triangle region). To further analyze the universality of this phenomenon, we design the Algorithm 1 to quantitatively analyze the ability of DDRP to capture information on the both left and right sides. To make a fair comparison, we also conduct the same process for BERT-R.

Algorithm 1 Count up-down triangle percentage

Require: N : total number of sentences; n : total number of sentences that have been processed; ms : maximum sentence length; mn : total number of sentences that match the upper and lower triangle; t : the threshold value that satisfies the up-down triangular effects; amp : attention map obtained by averaging attention maps in specific group.

```
1: Initialize  $ms \leftarrow 64, n \leftarrow 1, mn \leftarrow 0, t \leftarrow 0.7$ 
2: while  $n \leq N$  do
3:   Divide heads in equal with greater attention values in
   the (upper/lower) triangle region into (group 1/group 2)
   and obtain ( $amp1/amp2$ ).
4:   // prepare for the upper and lower triangle
5:   Sum the values in upper and lower triangles of  $amp1$ 
   respectively and obtain  $amp1_{up}, amp1_{down}$ .
6:   Sum the values in upper and lower triangles of  $amp2$ 
   respectively and obtain  $amp2_{up}, amp2_{down}$ .
7:    $amp1_{up} \leftarrow amp1_{up}/ms, amp1_{down} \leftarrow$ 
    $amp1_{down}/ms$ 
8:    $amp2_{up} \leftarrow amp2_{up}/ms, amp2_{down} \leftarrow$ 
    $amp2_{down}/ms$ 
9:   // compute for the upper and lower triangle
10:  if  $amp1_{up} \geq t$  and  $amp2_{down} \geq t$  then
11:     $mn \leftarrow mn + 1$ 
12:  else
13:    Continue
14:   $n \leftarrow n + 1$ 
15: return  $float(mn/N)$ 
```
