

# A Simple and Effective Positional Encoding for Transformers

Pu-Chin Chen\*, Henry Tsai\*, Srinadh Bhojanapalli\*,  
Hyung Won Chung, Yin-Wen Chang, Chun-Sung Ferng

Google Research

## Abstract

Transformer models are permutation equivariant. To supply the order and type information of the input tokens, position and segment embeddings are usually added to the input. Recent works proposed variations of positional encodings with relative position encodings achieving better performance. Our analysis shows that the gain actually comes from moving positional information to attention layer from the input. Motivated by this, we introduce **Decoupled positional attention for Transformers (DIET)**, a simple yet effective mechanism to encode position and segment information into the Transformer models. The proposed method has faster training and inference time, while achieving competitive performance on GLUE, XTREME and WMT benchmarks. We further generalize our method to long-range transformers and show performance gain.

## 1 Introduction

Transformers are sequence-to-sequence models that achieve state of the art performance in many Natural Language Processing (NLP) tasks, such as machine translation, language modeling and question answering (Vaswani et al., 2017; Devlin et al., 2018; Yang et al., 2019; Liu et al., 2020). Transformers have two major components: self-attention and a position-wise feed forward layer. Both are permutation equivariant and are not sensitive to the order of input tokens. To make these models position-aware, the position information of the input words is typically added as an additional embedding to the input token embeddings (Vaswani et al., 2017). For example, input embedding ( $W$ ) of a sentence is added to the position embeddings ( $P$ ), resulting in input  $W + P$  to the Transformer. These position embeddings only depend on the location the word appears. For multi-segment tasks,

additional segment embeddings can be added just like the position embeddings (Devlin et al., 2018).

There have been multiple works exploring different ways to include position information in Transformers (Shaw et al., 2018; Yang et al., 2019; Raffel et al., 2020). Many of those note the advantages of using a relative position encoding scheme over absolute position encodings (see also Fig 1). However what causes this difference is not clear. Yun et al. (2020) have shown that Transformers with absolute position encodings are universal approximators of all sequence to sequence functions, proving that absolute position encodings can capture the position information. Hence what causes the superiority of relative position encodings? A systematic study and understanding of the benefits and drawbacks of different position encoding methods is missing. Ke et al. (2020) hypothesised that the cross correlation between word and position embeddings while computing attention could be the cause of poor performance of absolute position encodings. However such cross terms are present in some of the relative position encoding methods (Shaw et al., 2018; Yang et al., 2019), and these methods perform on par or better than the other position encoding schemes (see §4).

In this paper we undertake a systematic study to understand different position encoding methods. We argue that absolute position embeddings mainly suffer from being added at the input. We show, with our experiments on classification, question answering and machine translation tasks, that absolute position encodings added to attention matrices with different parameters for each head improves significantly over absolute position encodings added to the input. This highlights that where the position information is included in the Transformer is important, providing an explanation for the gap in performance between absolute and relative position encodings. We also compare different position encodings and the effect of sharing position encod-

\* The authors contribute equally to this paper. Corresponding author email: puchin@google.com

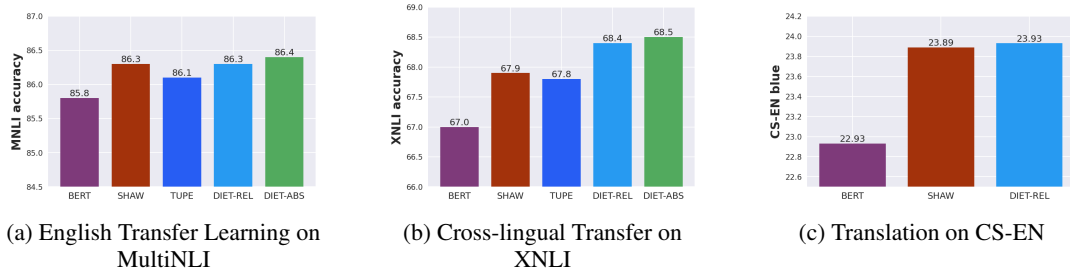


Figure 1: Performance effect of different positional encoding methods for Transformers (see § 2) on two Natural language Inference datasets from GLUE (Wang et al., 2019), XTREME (Hu et al., 2020) and one Neural Machine Translation dataset WMT 18 (Bojar et al., 2018). Absolute positional encoding (DIET-ABS) can achieve better performance than the relative counterpart (DIET-REL), showing the importance of designing the right position encoding method.

ings across different heads and layers of a Transformer. Based on these observations we propose decoupled positional attention and a new segment encoding approach (for tasks with multiple segments), and empirically show its superiority.

We summarize our contributions in this paper below.

- We theoretically and empirically analyze the limitation of the absolute position embeddings added to the input. For both absolute and relative information, we show that encoding position to attention matrix per-head results in superior performance.
- We propose a simple and efficient way to encode position and segment information. The proposed encoding matches the SoTA methods on multiple standard NLP tasks while having a simpler model with lower training/inference costs.
- Our proposed method can be easily applied to long sequence models (DIET-ABS<sup>LIN</sup>) and improve all metrics compared with Linformer (Wang et al., 2020).
- We present ablation studies comparing different position encoding methods and ways of sharing position encoding parameters across heads and layers in Transformer.

## 2 Position Encoding for Transformers

In this section, we briefly review the Transformer models (Vaswani et al., 2017) and discuss previous improvement of position encoding and analyze the limitation of the additive position embedding proposed in the initial and widely-adopted Transformer model.

### 2.1 Transformer

A Transformer block consists of two types of layers: 1) Self-attention layer and 2) Feed forward layers.

**Self-Attention Module** Given input sequence length  $n$ , hidden size  $d$ , multi-head query-key down-projection size  $d_h$ , we define hidden layer input to this attention head as  $\mathbf{X} \in \mathbb{R}^{n \times d}$ , the query projection matrix as  $\mathbf{W}_Q^i \in \mathbb{R}^{d \times d_h}$ , the key projection matrix as  $\mathbf{W}_K^i \in \mathbb{R}^{d \times d_h}$  and the value projection matrix as  $\mathbf{W}_V^i \in \mathbb{R}^{d \times d_h}$ ,  $i \in [h]$ , for  $h$  heads. Usually,  $d_h < d$  as we do multi-head attention with a smaller representation per head ( $d_h = d/h$ ). With that we can write dot-product attention score:

$$\mathbf{A}^i = (\mathbf{X}\mathbf{W}_Q^i)(\mathbf{X}\mathbf{W}_K^i)^\top$$

This attention score is used to compute the output for each head, after scaling and per row normalization using softmax:

$$\text{head}^i = \text{Softmax}(\mathbf{A}^i / \sqrt{d}) \cdot (\mathbf{X}\mathbf{W}_V^i)$$

Output of all attention heads in a layer are concatenated and passed to the next feed-forward layer applied token-wise.

### 2.2 Position Aware Self Attention

Many NLP tasks, such as machine translation, language modeling, are sensitive to the ordering of input words. Since Transformers are permutation equivariant, we usually additionally include the position information in the input. Below we discuss some of the popular position encoding methods.

#### 2.2.1 Absolute Position Encodings

Absolute position encodings are computed in the input layer and are summed with the input token embeddings. Vaswani et al. (2017) proposed this

for Transformers and it has been a popular choice in the followup works (Radford et al., 2018; Devlin et al., 2018). There are two common variations of the absolute position encodings - fixed and learned.

### 2.2.2 Relative Position Encodings

One drawback of absolute position encoding is that it requires fixed length of input sequence and does not directly capture relative positions to each word. To solve these problems several relative positions schemes have been proposed.

Shaw et al. (2018) proposed using relative position encoding instead of absolute position encoding, and add position embeddings to the key and optionally value projections instead of the input. They show that this new way of encoding position information leads to better performance on machine translation tasks. Yang et al. (2019) simplified this by removing the position embeddings in value projections and showed better performance on the language modeling tasks. Both these approaches use a vector representation to encode position information.

Raffel et al. (2020) use scalars to encode relative position between query and key indices and add directly to the attention scores matrix. They further use logarithmic binning of position information into a fixed number of buckets. All these relative position methods further share the position encoding parameters across layers.

Recently Ke et al. (2020) hypothesised that the cross correlation between position and token embeddings can result in weaker performance of additive absolute position embeddings and instead proposed to add both absolute and relative positional information based attention directly in each head. However such cross terms are present in the method proposed by Shaw et al. (2018), which does competitively with other approaches. We instead hypothesise that position encodings at input limit the rank of the position attention matrix leading to its poor performance.

### 2.3 Limitations of the Input Additive Position Embedding

In this section we discuss some limitations of the de facto way of adding absolute position encodings to the input token embeddings.

We first compare the representation power in terms of the rank of attention matrices achievable with different position encodings.

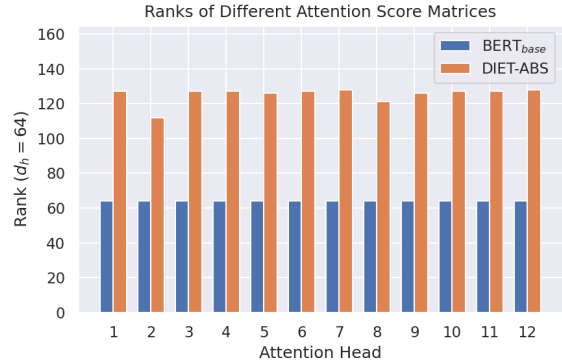


Figure 2: Rank of attention matrices: We present a comparison of the rank of the attention score matrices of a BERT<sub>BASE</sub> model with absolute position embeddings at input v.s. absolute position embeddings per-head (DIET-ABS (1)). With additive positional embedding at input, the attention matrices have much lower rank, limiting the representative power. This is alleviated by DIET-ABS.

**Theorem 1.** Let  $\mathbf{P} \in \mathbb{R}^{n \times d}$  be the input position embedding and  $\hat{\mathbf{P}} \in \mathbb{R}^{n \times d_p}$  be the layer-wise position embeddings. Let  $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{d \times d_h}$  be the query and key projection matrices with head projection size  $d_h$ , and  $d_h < d_p, d$  and  $n \geq d_h + d_p$ . Let  $\mathbf{A}_a = (\mathbf{X} + \mathbf{P})\mathbf{W}_Q\mathbf{W}_K^\top(\mathbf{X} + \mathbf{P})^\top$  and  $\mathbf{A}_r = \mathbf{X}\mathbf{W}_Q\mathbf{W}_K^\top\mathbf{X}^\top + \hat{\mathbf{P}}\hat{\mathbf{P}}^\top$  be the attention matrices computed using input and layer-wise position embeddings respectively. Then for any  $\mathbf{X}, \mathbf{P}, \mathbf{W}_Q, \mathbf{W}_K$

$$\text{rank}(\mathbf{A}_a) \leq d_h.$$

There exists a choice of  $\mathbf{X}, \hat{\mathbf{P}}, \mathbf{W}_Q, \mathbf{W}_K$  such that

$$\text{rank}(\mathbf{A}_r) = d_p + d_h > d_h.$$

**Remarks.** This theorem shows us that the rank of attention matrices is constrained with the absolute position encodings at the input and using per-head position encodings by adding position information to attention matrix directly results in allowing for higher rank attention. See § B for the proof.

Adding the position encodings directly to the input further places a constraint on training dynamics by forcing gradients to be same for both the input token and position embeddings (see § B). Relative position encodings discussed earlier, while addressing some of these concerns, suffer from slower training/inference times (see Table 1) with complex implementations (Shaw et al. (2018); Ke et al. (2020)). In the next section, we present simple position encoding methods that avoid these limitations.

### 3 Proposed Position and Segment Encodings

In the previous section, we learned about the limitations of input additive positional embeddings and existing works. Based on these observations, we propose two minimal/efficient ways to incorporate (absolute/relative) positional encodings along with a novel absolute segment encoding approach. By decoupling position and segment from token embeddings we match the SoTA performance while improving training/inference time (see §3.3).

#### 3.1 Decoupled Absolute Positional Attention

We propose the following simple absolute position encoding method that adds position information to the token attention matrix directly in each attention head. We further also add segment information to the token attention instead of the input embeddings. This way we can set the rank of position encodings independently resulting in higher rank attention matrix, addressing the limitations discussed earlier.

##### DIET-ABS

$$\mathbf{A}_{i,j}^{\text{ABS}} = (\mathbf{X}_i \mathbf{W}_Q)(\mathbf{X}_j \mathbf{W}_K)^\top / \sqrt{d} + (\mathbf{P}_Q \mathbf{P}_K^\top)_{i,j} + E_S(S(i), S(j)), \quad (1)$$

where  $\mathbf{P}_Q, \mathbf{P}_K \in \mathbb{R}^{n \times d_p}$  are low-rank position embedding matrices and  $E_S$  is the absolute segment attention to model interactions between segments defined as

$$E_S(S(i), S(j)) = \mathbf{S}_{i,\hat{j}} \quad (2)$$

where  $S(i) = \hat{i}$  if index  $i$  is in segment  $\hat{i}$ .

Please note that we use the following notation in the above equation.  $\mathbf{A}_{i,j}$  denotes the  $(i, j)$  entry of matrix  $\mathbf{A}$ .  $\mathbf{X}_i$  and  $\mathbf{X}_j$  denote the  $i$ th row and  $j$ th column of  $\mathbf{X}$  respectively. We will follow this notation in the remainder of the paper.

By default, we set  $d_p$  same as  $d_h$ . This already results in potentially a rank  $d_p + d_h$  attention matrix as shown in Theorem 1. To illustrate this, we compare the rank of the attention matrices in the first layer of a baseline BERT model and a DIET-ABS model for a sampled batch in Figure 2. The figure shows that attention matrices of DIET-ABS have higher ranks than the baseline BERT. Our detailed experiment results in § 4 also show that DIET-ABS performs noticeably better. This confirms our earlier observation in Theorem 1 that additive position embeddings at input can constrain the model and

adding the position embeddings per-head removes this constraint and results in better performance.

With the decoupled positional embedding, we can increase  $d_p$  to any width  $k$  to break the low-rank bottleneck shown in Theorem 1. We call such model DIET-ABS-Rank- $k$ . We also address the efficiency issue introduced by one additional matrix multiplication ( $\mathbf{P}_Q \mathbf{P}_K^\top$ ). As the positional embeddings are independent of the input, we only need to compute the matrix multiplication once for each training batch, and we can cache the computed matrix before running inference. As a result, we observe neglectable training and inference cost increase in this model variant.

#### 3.2 Decoupled Relative Positional Attention

To incorporate relative position inductive bias, we consider a simplified version of the position encoding proposed in T5 (Raffel et al., 2020) without log-binning and per-layer parameter sharing. We further also incorporate our per-head segment encoding as in DIET-ABS. The model can be written as:

##### DIET-REL

$$\mathbf{A}_{i,j}^{\text{REL}} = (\mathbf{X}_i \mathbf{W}_Q)(\mathbf{X}_j \mathbf{W}_K)^\top / \sqrt{d} + \mathbf{R}_{i-j} + E_S(S(i), S(j)). \quad (3)$$

We show an example of this model with two segments in Figure 3.

#### 3.3 Training and Inference Costs

We next show the proposed models introduce little computational overhead compared to the baseline model, making our model more practical than alternatives. We consider two different models - BERT<sub>BASE</sub> model and a smaller model, BERT<sub>SMALL</sub>, that has hidden size 512, 4 layers and 8 attention heads.

In Table 1 we compare the training and inference costs of position encoding methods of Shaw et al. (2018), Ke et al. (2020), DIET-ABS and DIET-REL. We notice that the simplicity of the proposed methods indeed translates to savings in both training and inference times compared to other position encoding approaches. The savings in step times are even more significant for smaller models (BERT<sub>SMALL</sub>) and during inference.

Note that the discrepancy between training and inference speed is likely because gradient updates dominate the cost at training time (Lan et al., 2020). At inference time, we only measure the time of a

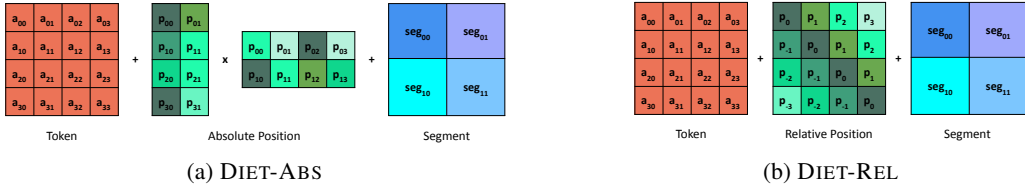


Figure 3: Proposed efficient approach to include position and segment encoding by adding them directly to the token attention matrix per-head. Left figure shows how we encode absolute positional attention. Right figure represents relative positional attention.

	Mode	Shaw et al. (2018)	Ke et al. (2020)	DIET-ABS	DIET-REL
BERT <sub>BASE</sub>	Training	+13%	+1%	+0%	+0%
BERT <sub>BASE</sub>	Inference	+33%	+19%	+0%	+0%
BERT <sub>SMALL</sub>	Training	+24%	+4%	+0%	+0%
BERT <sub>SMALL</sub>	Inference	+65%	+27%	+1%	+0%

Table 1: Pre-training and inference time of Transformers with different position encoding methods in comparison to the baseline BERT model on TPU v2. We observe that simplicity of the DIET-REL and DIET-ABS result in substantial gains in both training and inference time. We notice even more speedup for the smaller BERT<sub>SMALL</sub> model compared to BERT<sub>BASE</sub>.

forward pass which corresponds to costs of using such models in real systems.

### 3.4 Application to Long-range Transformers

Another advantage of our propose approaches is they easily extend to long range Transformer models. For long sequence inputs, Transformers suffer from quadratic dependence of computational complexity with respect to the sequence length. A class of methods reduce this complexity by using a low rank projection of the input sequence for attention computation (Wang et al., 2020; Choromanski et al., 2021; Dai et al., 2020). However, such methods use the default input position encodings, and there has not been much work in incorporating position information per-head without introducing the quadratic computation complexity on the input sequence length. We illustrate the applicability of our methods to such settings by applying DIET-ABS to Linformer (Wang et al., 2020), which projects the attention key and value matrices to a lower dimension  $k$  during attention computation.

**DIET-ABS<sup>LIN</sup>** The proposed method can be written as:

$$\mathbf{A}_{i,j}^{\text{LIN}} = (\mathbf{X}_i \mathbf{W}_Q)((\mathbf{E}\mathbf{X})_j \mathbf{W}_K)^\top / \sqrt{d} + (\mathbf{P}_Q \mathbf{P}_K^\top)_{i,j}, \quad (4)$$

where  $\mathbf{E} \in \mathbb{R}^{k \times n}$ ,  $\mathbf{P}_Q \in \mathbb{R}^{n \times d}$ ,  $\mathbf{P}_K \in \mathbb{R}^{k \times d}$ .

## 4 Experiments

In this section, we present our experimental results comparing different position and segment encoding approaches discussed in earlier sections. We

conduct experiments in three different settings to cover a wide range of use cases. First, we examine the results of a popular transfer learning approach from masked-LM pretraining to the end tasks in GLUE (Devlin et al., 2018). Second, we study zero-shot cross-lingual transferability of the multilingual pretrained models (Hu et al., 2020) to classification and question answering tasks in the XTREME benchmark (Hu et al., 2020). Lastly, we consider training Transformer models from scratch for machine translation.

We compare the following positional encoding approaches - absolute positional embedding (Devlin et al., 2018), relative positional embedding (Shaw et al., 2018), combined absolute and relative positional encoding (Ke et al., 2020), relative scalar approach (Raffel et al., 2020), our proposed DIET-ABS and DIET-REL per-head positional encoding approaches. We denote the methods that add position/segment information directly to input token embeddings with *input*, and methods that add position/segment information directly in attention layer with *per-head*. For complete experimental setup, see Appendix A.

### 4.1 English Transfer Learning Results

**Datasets and Model** For pre-training, we use English Wikipedia and Books datasets (Devlin et al., 2018). For Finetuning tasks we use the datasets from the GLUE benchmark (Wang et al., 2019). We apply sub-word tokenization on raw text data using WordPiece (Wu et al., 2016) with a 30,000 token vocabulary.

Model	Position	Segment	MNLI	QQP	QNLI	SST2	CoLA	STS-B	Avg
			393k	364k	105k	67k	8.5k	7k	
Devlin et al. (2018)	input	input	85.8 / 85.9	91.1	89.9	93.2	58.7	89.0	84.8
Shaw et al. (2018)	per-head	input	86.3 / 86.0	91.2	90.5	93.2	59.8	89.3	85.2
Raffel et al. (2020)	per-head	input	86.4 / 86.2	91.2	90.1	93.0	59.6	90.1	85.2
Ke et al. (2020)	per-head	input	86.1 / 86.2	91.2	90.3	93.1	59.6	89.6	85.2
DIET-REL	per-head	input	86.0 / 86.1	91.0	89.8	92.8	59.6	89.0	84.9
DIET-REL	per-head	per-head	86.3 / 86.3	91.0	90.5	92.9	60.3	89.3	85.2
DIET-ABS ( $d_p=128$ , share)	per-head	per-head	86.4 / 86.4	90.8	89.5	93.0	59.8	90.2	85.2
Wang et al. (2020) ( $d_p=32$ )	input	input	82.3 / 82.6	90.2	86.3	91.4	53.9	87.6	82.0
DIET-ABS <sup>LIN</sup> ( $d_p=32$ )	per-head	input	83.0 / 83.1	90.6	86.7	92.0	55.7	87.6	82.7

Table 2: GLUE: Results on the GLUE dev set of the finetuned models based on a pre-trained model with 12-layer BERT<sub>BASE</sub> architecture. We report the median of the maximum accuracy over all checkpoints among five runs. We notice that the shared DIET-ABS with rank 128 performs competitively with existing relative positional embedding SoTA models without the inductive bias of the relative positions. The proposed method also improves performance in the low-rank long range transformer setting of (Wang et al., 2020), where relative positional embedding approaches are inefficient to use.

Model	Position	Segment	Classification	Question Answering			Avg
			XNLI 393k	XQuAD 88k	MLQA	TyDiQA 3.7k	
Devlin et al. (2018)	input	input	67.0	66.0 / 49.9	56.2 / 41.0	59.0 / 47.9	55.3
Shaw et al. (2018)	per-head	input	67.9	69.5 / 53.9	58.2 / 43.1	64.8 / 49.9	58.2
Raffel et al. (2020)	per-head	input	68.5	69.9 / 53.5	59.5 / 44.3	63.8 / 50.6	58.6
Ke et al. (2020)	per-head	input	67.8	68.6 / 52.0	58.6 / 43.2	63.9 / 48.7	57.5
DIET-REL	per-head	input	68.0	68.1 / 52.8	57.7 / 42.7	63.3 / 50.9	57.6
DIET-REL	per-head	per-head	68.4	69.4 / 54.4	58.6 / 43.5	62.4 / 49.3	58.0
DIET-ABS ( $d_p=128$ , share)	per-head	per-head	68.5	70.0 / 53.6	59.8 / 44.5	64.6 / 51.5	58.9
Wang et al. (2020) ( $d_p=256$ )	input	input	63.6	59.1 / 43.7	48.9 / 34.0	50.5 / 37.9	48.2
DIET-ABS <sup>LIN</sup> ( $d_p=256$ )	per-head	input	64.4	61.6 / 46.0	52.2 / 37.0	53.6 / 40.9	50.8

Table 3: XTREME: Fine-tune cross-lingual model on English training set (Cross-lingual Transfer). Performance is measured by *accuracy* for classification, and *f1 score / exact match* for question answering. In agreement with results in Table 2 we see in this table that using per-head position encodings is strictly better than absolute position encodings at the input. With layer-wise sharing, DIET-ABS with rank 128 outperforms all SoTA models.

Model	EN-DE	DE-EN	EN-CS	CS-EN
Vaswani et al. (2017)	39.00	38.42	18.55	22.93
Shaw et al. (2018)	40.10	38.90	18.74	23.89
DIET-REL	39.47	38.49	18.68	23.93

Table 4: Machine Translation: We report results comparing different position encoding methods for Transformers on machine translation tasks en-de, de-en, en-cs and cs-en from the Newstest 2018 dataset. We notice that all per-head position encoding schemes (all except the first row) do better than the absolute position embeddings added at the input. Further the proposed simple DIET-REL approach is competitive with other position encoding approaches.

**Results** We examine how different ways of encoding position and segment affect the transfer learning ability of the pre-trained English BERT models by fine-tuning on the GLUE benchmark (Wang et al., 2019), and present the results in Ta-

ble 2. We first notice that all the approaches that encode position features explicitly at per-head level perform better than the baseline additive position encodings at the input (Devlin et al., 2018). All models incorporating relative positions (Shaw et al., 2018; Raffel et al., 2020; Ke et al., 2020), despite their modeling differences, have very similar average score. We show further gains (84.9 to 85.2 for DIET-REL) by moving segment features to per-head.

Interestingly we notice that the proposed absolute position encoding method DIET-ABS, with layer-wise sharing, is on par with all previous SoTA relative positional encodings. This shows that even absolute position encodings can perform better when included per-head instead at the input. We present a detailed ablation study varying the rank and sharing methods of absolute positional attention (DIET-ABS) in Table 8 and Tables 9 in

Appendix C.

For long range input, we consider Linformer (Wang et al., 2020) with a projection dimension of 32. Due to down-projection, we see non-trivial performance drop, when compared to a Transformer. Even for this setting we see that our absolute positional attention DIET-ABS can be used to improve the model’s performance.

## 4.2 Cross-lingual Model Results

**Datasets and Model** For our multilingual experiments, we pre-train the models on Wikipedia corpus in 100 languages similar to (Lample and Conneau, 2019) for 125K steps with a sequence length of 512, and then fine-tune on downstream XTREME tasks (Hu et al., 2020). We use language-independent tokenizer, Sentence Piece (Kudo and Richardson, 2018) model, with 120,000 token vocabulary to encode input text.

**Classification** We conduct 5 trials of fine-tuning for each model on the MultiNLI (Williams et al., 2018) training data, then perform zero-shot predictions on XNLI (Conneau et al., 2018), choosing median accuracy to report.

**Question Answering** We conduct 5 trials of fine-tuning for each model on SQuAD V1.1 dataset, following by zero-shot predictions on XQuAD (11 languages), MLQA (7 languages) and TyDiQA-GoldP (9 languages), choosing median F1 / EM scores to report.

**Results** We present our results on the classification and question answering finetuning tasks in XTREME for different position and segment encoding methods in Table 3. Again all per-head position encoding methods outperform input additive position encodings. Interestingly, our simple DIET-ABS turns out to be the best model, better than other models using relative position features. Layer-wise sharing and per-head segment attention allows DIET-ABS to outperform DIET-REL. We present a detailed ablation study in Table 5 to understand effect of decoupled positional attention variants. Finally, we notice similar advantages in using DIET-ABS with the Linformer (Wang et al., 2020) model in the long range setting.

## 4.3 Translation Results

**Datasets and Model** For the machine translation task we consider two language pairs (both directions) for training - WMT 2018 English-to-German

(en-de), German-to-English (de-en), English-to-Czech (en-cs) and Czech-to-English (cs-en) (Bojar et al., 2018). We test the corresponding models on Newstest 2018 datasets respectively and report the BLEU score output by SacreBLEU (Post, 2018) with default setting. Our setup follows Vaswani et al. (2017) closely and use their Tensor2Tensor framework (Vaswani et al., 2018). Following Vaswani et al. (2017) we use a 6 layer Transformer with encoder-decoder architecture. For more details of our experimental setup please see Appendix A

**Results** We report the BLEU scores of the models in Table 4. We observe that moving positional information from input to per-head attention layer improves BLEU scores. Different variations of per-head positional attention do not make much difference with DIET-REL being competitive with Shaw et al. (2018).

## 4.4 Ablation Study

In this section, we share our findings of key factors that affect performance of decoupled positional attention.

**Sharing the Positional Encoding** Previous works (Raffel et al., 2020; Ke et al., 2020; Shaw et al., 2018) used different sharing methods for the positional encodings to reduce the model parameters. We present a detailed study on different forms of sharing positional encodings and its effect on performance. In particular, we compare the following variations in sharing the position encoding parameters across different heads and the layers in the Transformer.

- head-wise - Same parameters are used for all heads in a layer, with different layers using different parameters (Shaw et al., 2018; Ke et al., 2020).
- layer-wise - Sharing of position encoding parameters across layers with different parameters for each head (Raffel et al., 2020).
- none - Every layer and head uses different position encoding parameters.

We present results comparing different sharing methods in Table 5 for XTREME tasks. We make the following observations 1) head-wise sharing is consistently worse than layer-wise, 2) sharing hurts the performance of DIET-REL whereas it improves

Model	Sharing	Segment	Classification	Question Answering			Avg
			XNLI	XQuAD	MLQA	TyDiQA-GoldP	
DIET-REL	-	input	68.0	68.1 / 52.8	57.7 / 42.7	63.3 / 50.9	57.6
DIET-REL	head-wise	input	67.7	66.2 / 51.0	56.0 / 41.1	60.1 / 45.9	55.4
DIET-REL	layer-wise	input	68.0	68.6 / 53.3	58.1 / 43.1	61.3 / 48.2	57.2
DIET-REL	-	per-head	68.4	69.4 / 54.4	58.6 / 43.5	62.4 / 49.3	58.0
DIET-REL	head-wise	per-head	67.8	66.0 / 50.5	55.5 / 40.4	59.2 / 44.6	54.7
DIET-REL	layer-wise	per-head	68.1	68.7 / 53.8	58.4 / 43.2	61.0 / 48.4	57.3
DIET-ABS ( $d_p=64$ )	-	input	68.0	67.4 / 50.5	57.8 / 42.3	61.3 / 46.8	56.3
DIET-ABS ( $d_p=64$ )	-	per-head	67.9	67.5 / 52.4	57.3 / 42.3	61.6 / 46.8	56.5
DIET-ABS ( $d_p=128$ )	-	per-head	68.1	68.2 / 52.0	57.9 / 42.6	61.5 / 47.6	56.8
DIET-ABS ( $d_p=512$ )	-	per-head	68.5	68.0 / 52.0	57.7 / 42.4	61.6 / 48.4	56.9
DIET-ABS ( $d_p=64$ )	layer-wise	input	68.0	69.3 / 53.1	59.3 / 43.9	63.2 / 48.6	57.9
DIET-ABS ( $d_p=64$ )	layer-wise	per-head	68.4	69.3 / 53.2	59.4 / 44.1	63.3 / 48.6	58.0
DIET-ABS ( $d_p=128$ )	layer-wise	per-head	68.5	70.0 / 53.6	59.8 / 44.5	64.6 / 51.5	58.9
DIET-ABS ( $d_p=256$ )	layer-wise	per-head	68.4	69.9 / 53.8	59.6 / 44.2	62.8 / 49.1	58.3
DIET-ABS ( $d_p=512$ )	layer-wise	per-head	67.8	69.0 / 53.2	58.4 / 43.0	62.5 / 48.8	57.5

Table 5: Ablation study on XTREME: We run decoupled positional attention ablation study to understand the effect of 1) sharing positional attention parameters across layers and heads 2) segment attention added at per-head 3) performance of relative and absolute 4) absolute positional attention rank  $d_p$  from 64 to 512.

	English			Multilingual		
	Parameters	+ $\Delta$	GLUE	Parameters	+ $\Delta$	XTREME
Devlin et al. (2018)	110.1M	-	84.8	178.9M	-	55.3
Shaw et al. (2018)	112.9M	+2.5%	85.2	181.7M	+1.7%	57.9
DIET-REL	109.9M	+0.0%	85.2	178.7M	+0.0%	58.0
DIET-REL (share)	109.7M	+0.0%	85.0	178.5M	+0.0%	57.3
DIET-ABS ( $d_p=128$ )	128.6M	+16.8%	85.3	197.4M	+10.0%	56.8
DIET-ABS ( $d_p=128$ , share)	111.3M	+1.1%	85.2	180.1M	+0.6%	58.9

Table 6: Model Parameters: We list the number of model parameters and performance for different position encoding approaches. We observe that sharing hurts the performance of DIET-REL with negligible benefit in the number of parameters. On the contrary, the regularization effect of sharing makes DIET-ABS more stable with lesser parameters to achieve competitive performance.

the performance of DIET-ABS. We summarize the key settings along with the number of model parameters in Table 6. For DIET-REL, sharing brings little effect on saving parameters, and hurts the performance. Hence, we recommend no sharing for relative positional encodings (DIET-REL). On the other hand, it is necessary to share parameters for DIET-ABS in order to keep the number of parameters low. Interestingly, sharing has regularization effect on DIET-ABS, making the model perform better. We choose layer-wise sharing over head-wise sharing for its better performance.

**Segment Encoding** Our novel segment encoding design further improves the model performance showed in Table 5. Both relative and absolute decoupled positional attention models benefit from moving the segment encoding from input to per-head: DIET-REL (+0.4%), layer-wise shared DIET-REL (+0.1%), DIET-ABS (+0.2%), layer-wise shared DIET-ABS (+0.1%). See Appendix D for the results of GLUE benchmark and

Appendix C for segment attention visualization.

**Rank of Absolute Positional Attention** The design of DIET-ABS allows to learn higher rank attention matrices as shown in Theorem 1. To understand the effect of absolute positional attention rank ( $d_p$ ) in practice, we conduct experiments varying the rank from  $d_p = 64$  to  $d_p = 512$ . We present the results in Table 5. We notice that the performance improves as we increase the rank from 64 to 128. However there is a performance saturation in further increasing it to 512. We present a visualization of the rank of the positional attention matrix in Appendix B.

#### 4.5 Positional Attention Pattern Visualization

We next visualize the learned positional attention patterns of DIET-ABS in Figure 4. We first note that DIET-ABS has learned to capture the relative positional relations between inputs. Also note that, for the the index zero (the [CLS] token), decoupled absolute positional attention usually learns a spe-



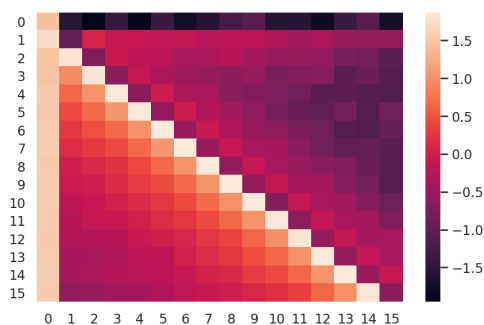


Figure 4: Visualization of learned positional attention patterns of DIET-ABS. Note that in addition to capturing the the relative positional relations, the model also learn to attend to [CLS] at index 0, suggesting the dedicated [CLS] untying design in Ke et al. (2020) is not necessary with DIET-ABS.

cial pattern. This pattern cannot be solely modeled by existing relative positional embedding methods, and some existing works (Ke et al., 2020) handled this case specifically by introducing new parameters. This shows the benefit of DIET-ABS in not requiring any carefully designed inductive biases as in existing approaches (Shaw et al. (2018); Raffel et al. (2020)), which may not generalize across tasks.

## 5 Conclusion

In this paper we theoretically and empirically examined the limitation of additive position embedding at input and showed that having per-head position embeddings results in better performance. We argued that the superior performance of some of the relative position encoding methods come from their per-head addition to attention matrix rather than the position information being relative vs absolute. Indeed we show that using absolute position encodings per-head results in better performance. Motivated by this we propose a simple per-head position and segment attention method that achieves the state-of-the-art performance on multiple NLP tasks and is more computationally efficient than existing approaches.

## References

Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Philipp Koehn, and Christof Monz. 2018. Findings of the 2018 conference on machine translation (WMT18). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 272–303, Belgium, Brussels. Association for Computational Linguistics.

Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarrlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy Colwell, and Adrian Weller. 2021. Rethinking attention with performers.

Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Zihang Dai, Guokun Lai, Yiming Yang, and Quoc V. Le. 2020. Funnel-transformer: Filtering out sequential redundancy for efficient language processing.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional Transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *Proceedings of Machine Learning and Systems 2020*, pages 7449–7459.

Guolin Ke, Di He, and Tie-Yan Liu. 2020. Rethinking the positional encoding in language pre-training. *arXiv preprint arXiv:2006.15595*.

Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing.

Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining.

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.

Xiaodong Liu, Kevin Duh, Liyuan Liu, and Jianfeng Gao. 2020. Very deep transformers for neural machine translation.

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*.

Matt Post. 2018. A call for clarity in reporting bleu scores. In *WMT*.

Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *Technical Report, OpenAI*.

- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468.
- Ashish Vaswani, Samy Bengio, Eugene Brevdo, Francois Chollet, Aidan N Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, et al. 2018. Tensor2tensor for neural machine translation. *arXiv preprint arXiv:1803.07416*.
- 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 *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019*.
- Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. [Linformer: Self-attention with linear complexity](#).
- Yu-An Wang and Yun-Nung Chen. 2020. What do position embeddings learn? an empirical study of pre-trained language model positional encoding. In *EMNLP 2020*.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#).
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#).
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.
- Chulhee Yun, Srinadh Bhojanapalli, Ankit Singh Rawat, Sashank J Reddi, and Sanjiv Kumar. 2020. Are Transformers universal approximators of sequence-to-sequence functions? In *International Conference on Learning Representations*.

## A Experimental setup

In this section we present more details of our experimental setup.

**Pre-training** We pre-train the models using a masked LM task (Devlin et al., 2018) and do not use the Next Sentence Prediction (NSP) loss as suggested in RoBERTa (Liu et al., 2019). Each input is constructed with full sentences from documents, and packed up to the maximum sequence length. We use the same architecture as BERT<sub>BASE</sub> (Devlin et al., 2018) ( $L = 12$ ,  $H = 768$ ,  $A = 12$ ) for our experiments.

**Fine-tuning** Some downstream tasks have different groups of full sentences provided at inputs. For those tasks (e.g. MNLI, CoLA, XNLI, SQuAQ), we fine-tune models with supplemental segment encoding discussed in Section §3. We leave models for other tasks unchanged as their pre-training correspondences.

**Hyper-parameters** Hyper-parameters we use are presented in Table 7.

	English		Multilingual	
	Pretrain	Finetune	Pretrain	Finetune
Max Steps	500K	5 or 10 epochs	125K	3 epochs
Learning Rate	0.0018	{1e-5, 2e-5, 3e-5, 4e-5}	0.0018	{1e-5, 2e-5, 3e-5, 4e-5}
Warmup Proportion	0.025	0.1	0.025	0.1
Sequence Length	128	128	512	512
Batch Size	4096	32	4096	32
Checkpoint Interval	20k	3.5k	20k	3.5k

Table 7: Hyperparameters for all models

**Translate** For our Translate experiments we follow the setup of Vaswani et al. (2017) and use their Tensor2Tensor framework (Vaswani et al., 2018). We train using WMT18 ((Europarl v7, Common Crawl corpus and News Commentary v13) en-de, de-en, en-cs and cs-en datasets. We report BLUE scores provided by SacreBLEU (Post, 2018) on newstest 2018 dataset. We train a 6 layer Transformer model. Any changes to position encoding are applied to all the attention layers both in the encoder and decoder. We use Adam optimizer and train for 250k steps. For decoding we use beam search with beam size 10 and length penalty 0.6.

## B Proofs

*Proof of Theorem 1.* The first claim follows easily by observing that rank of product of an two matrices is upper bounded by the minimum of the individual ranks.

$$\begin{aligned} \text{rank}(\mathbf{A}_a) &= \text{rank}((\mathbf{X} + \mathbf{P})\mathbf{W}_Q\mathbf{W}_K^\top(\mathbf{X} + \mathbf{P})^\top) \\ &\leq \min(\text{rank}(\mathbf{X} + \mathbf{P}), \text{rank}(\mathbf{W}_Q), \text{rank}(\mathbf{X} + \mathbf{P}), \text{rank}(\mathbf{W}_K)) \\ &\leq d_h. \end{aligned}$$

$$\text{rank}((\mathbf{X} + \mathbf{P})\mathbf{W}_Q\mathbf{W}_K^\top(\mathbf{X} + \mathbf{P})^\top) \leq d_h, \text{ where } \mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{d \times d_h}$$

The last inequality follows from  $\text{rank}(\mathbf{W}_Q) \leq d_h$  as  $\mathbf{W}_Q \in \mathbb{R}^{d \times d_h}$ .

To prove the second claim we follow a construction approach. Let us first take  $\mathbf{W}_Q = \mathbf{W}_K$  to be same matrices with first  $d_h$  rows being identity matrix and the remaining  $d - d_h$  rows being all zeros. Then

$$\mathbf{W}_Q\mathbf{W}_K^\top = \begin{pmatrix} I_{d_h, d_h} & 0_{d_h, d-d_h} \\ 0_{d-d_h, d_h} & 0_{d-d_h, d-d_h} \end{pmatrix}.$$

Here  $I_{d_h, d_h}$  denotes the identity matrix in  $\mathbb{R}^{d_h \times d_h}$  and  $0_{d_h, d}$  denotes the all zeros matrix in  $\mathbb{R}^{d_h, d}$ .

We let  $\mathbf{X}$  be such that the first  $d$  rows form an identity matrix and rest are zeros -  $\mathbf{X}^\top = [I_{d, d}, 0_{n-d, d}]$ . Hence  $\mathbf{X}\mathbf{W}_Q\mathbf{W}_K^\top\mathbf{X}^\top$  becomes a similar diagonal matrix with

$$\mathbf{X}\mathbf{W}_Q\mathbf{W}_K^\top\mathbf{X}^\top = \begin{pmatrix} I_{d_h, d_h} & 0_{d_h, n-d_h} \\ 0_{n-d_h, d_h} & 0_{n-d_h, n-d_h} \end{pmatrix}.$$

Choose  $d_p = n > d_h$  and let  $\hat{\mathbf{P}} = I$ . Now choosing  $\hat{\mathbf{P}}$  with zeros in the first  $n - d_p$  columns and identity in the last  $d_p$  columns ( $\hat{\mathbf{P}} = [0_{d,n-d_p}, I_{d_p,d_p}]$ ) gives

$$\hat{\mathbf{P}}\hat{\mathbf{P}}^\top = \begin{pmatrix} 0_{n-d_p,n-d_p} & 0_{n-d_p,d_p} \\ 0_{d_p,n-d_p} & I_{d_p,d_p} \end{pmatrix}.$$

Combining these two gives us

$$\begin{aligned} \text{rank}(\mathbf{A}_r) &= \text{rank}(\mathbf{X}\mathbf{W}_Q\mathbf{W}_K^\top\mathbf{X}^\top + \hat{\mathbf{P}}\hat{\mathbf{P}}^\top) \\ &= \min(d_h + d_p, n) > d_h. \end{aligned}$$

□

Let  $\mathbf{X} \in \mathbb{R}^{n \times d}$  be the input word embeddings in dimension  $d$  with sequence length  $n$ . We have trainable position embeddings  $\mathbf{P} \in \mathbb{R}^{n \times d}$ , which are added to the input sequence before feeding into the model  $g$ . For a given input  $\mathbf{X}$  and label  $y$ , the objective for a loss function  $\ell$  is as follows:

$$L = \ell(g(\mathbf{X} + \mathbf{P}), y) \quad (5)$$

**Theorem 2.** *Let  $\mathbf{X}$  and  $\mathbf{P}$  be trainable embedding matrices in  $\mathbb{R}^{n \times d}$ . Then the gradients of the loss function in equation (5), at any point  $(\mathbf{X}, y)$ , and for any differentiable functions  $\ell$  and  $g$ , are same for  $\mathbf{X}$  and  $\mathbf{P}$ .*

**Remarks.** This theorem shows us that the gradients are same for the input token embeddings and position embeddings. While in standard NLP tasks the inputs  $X$  can be different in each step due to different input tokens being present in each mini batch, the result still suggests that additive position embedding can limit the model from learning the relative importance of position encodings with respect to token embeddings based on the training task at hand.

*Proof of Theorem 2.* The above theorem follows by just computing the gradients and showing they are equal for each step.

Gradients of the above objective w.r.t  $\mathbf{X}$  and  $\mathbf{P}$  are as follows.

$$\begin{aligned} \nabla_{\mathbf{X}}L &= \nabla_g L \cdot \nabla_{\mathbf{X}+\mathbf{P}}g \cdot \nabla_{\mathbf{X}}(\mathbf{X} + \mathbf{P}) \\ &= \nabla_g L \cdot \nabla_{\mathbf{X}+\mathbf{P}}g \\ \nabla_{\mathbf{P}}L &= \nabla_g L \cdot \nabla_{\mathbf{X}+\mathbf{P}}g \cdot \nabla_{\mathbf{P}}(\mathbf{X} + \mathbf{P}) \\ &= \nabla_g L \cdot \nabla_{\mathbf{X}+\mathbf{P}}g. \end{aligned}$$

The above computation of gradient follows from chain rule. This shows that the gradients of  $L$  w.r.t.  $\mathbf{X}$  and  $\mathbf{P}$  are the same. □

## C Attention Visualization

In this section, we examine the model internals to understand how the proposed model works. We first visualize the model internals of different modeling alternatives to argue our proposed model is sensible.

**Why We Remove the Input Embedding** To understand if it is sensible to remove the input additive embedding after adding position scalars per-head, we add additive position embedding to our DIET-ABS model. Then, we examine the position embedding of the BERT model and our DIET-ABS variant with additive position embedding. Figure 5 shows that, when the model has both absolute scalar and additive absolute position embedding, the position embedding encodes almost no information — all position embeddings at input are similar.

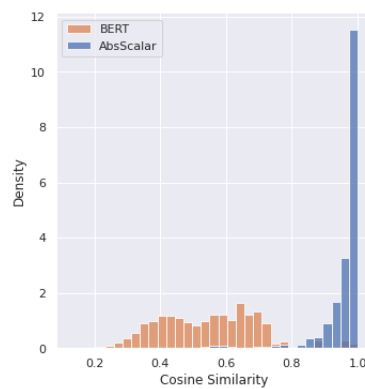


Figure 5: The cosine similarity distribution between all absolute position pairs of the input additive positional embedding for the baseline BERT model and the proposed DIET-ABS. We observed that, after the position features are added to each head as in DIET-ABS, the input position embedding contains almost no information — all input position pairs are similar.

**The Effect of Segment Attention** We also examine the effect of adding segment attention on top of the position attention. Figure 6 shows some representative patterns. We observe that segment attention enables the model to attend more to parts of the sequence that belongs to certain segments.

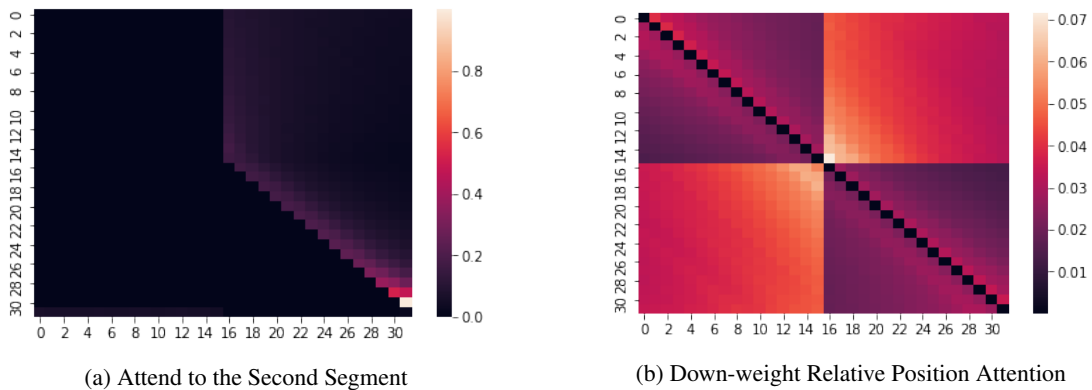


Figure 6: We consider input of length 32 with two segments. The second segment starts at index 16. We observe the attention patterns in the DIET-REL model without token-to-token attention.

**Shifting Pattern Learned from Absolute Positional Attention** Using relative position encoding gives generally better results despite smaller improvement scale compared to moving feature encoding per-head. To understand this, we visualize the attention pattern of the absolute positional attention and found two representative patterns in DIET-ABS in Figure 7. We observe that even given absolute position features, the model learns a “shifting pattern” for the most part. Different from Wang and Chen (2020) which claimed absolute position only learns local patterns, we show the position attention can actually attend to longer context. However, the shifting pattern can be modeled directly by relative position. Thus, DIET-REL can be a better model choice with fewer parameters and more accurate inductive bias in some applications.

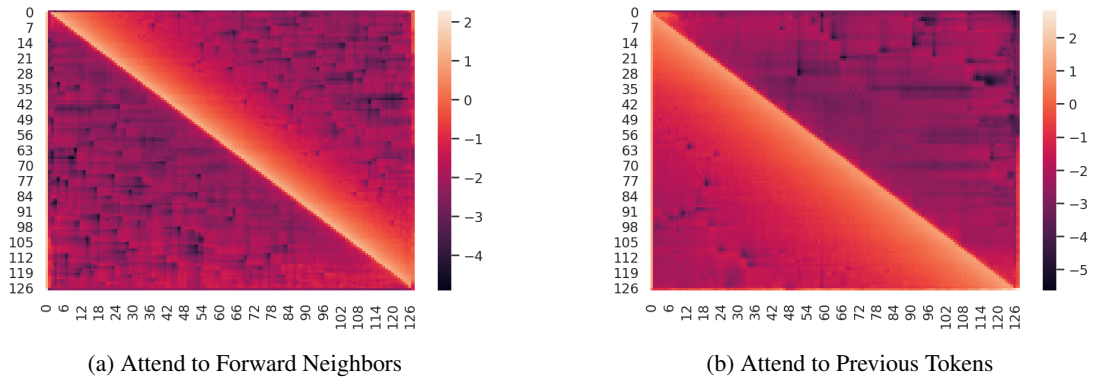


Figure 7: Sampled position attention score patterns for the DIET-ABS model. We can see a clear shifting patterns generated by the model. Such patterns can be modeled better by relative positional scalar encodings.

**Rank of Positional Attention Matrices** In Figure 8, we present a comparison of rank of position attention matrices for a BERT<sub>BASE</sub> model with absolute position embeddings at input ( $\mathbf{P}_Q \mathbf{W}_Q \mathbf{W}_K^T \mathbf{P}_K^T$ ) v.s. absolute position embeddings per-head (DIET-ABS (1),  $(\mathbf{P}_Q \mathbf{P}_K^T)$ , where  $\mathbf{P}_Q, \mathbf{P}_K \in \mathbb{R}^{n \times d_p}$ ). With additive positional embedding at input, position attention matrices have much lower rank, limiting the representative power. This is alleviated by DIET-ABS.

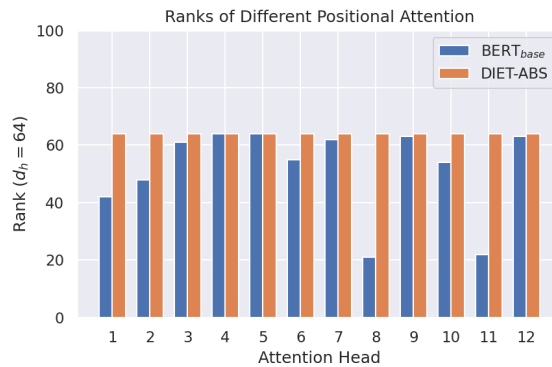


Figure 8: Rank of positional attention matrices

## D Additional Ablation Study on GLUE

Earlier we present an ablation study on XTREME in Table 5 for decoupled positional attention variants. We compare DIET-REL and DIET-ABS against the baseline (Devlin et al., 2018). We now present a similar study on the GLUE benchmark in Table 8 and observe similar results.

**Positional Encoding** In Table 8, moving positional embeddings from input to per-head improves average score for both DIET-REL (+0.1%) and DIET-ABS (+0.2%).

**Segment Encoding** In Table 8, moving segment embeddings from input to per-head improves both DIET-REL (+0.3%) and DIET-ABS (+0.05%).

**Sharing Strategies** Sharing plays an important role for DIET-ABS. In Table 9, we find that sharing will degrade the performance of DIET-REL (-0.2% layer-wise, -0.3% head-wise). For DIET-ABS, sharing makes the model more stable, and able to compete with DIET-REL.

Model	Position	Segment	MNLI 393k	QQP 364k	QNLI 105k	SST2 67k	CoLA 8.5k	STS-B 7k	Avg
Devlin et al. (2018)	input	input	85.8 / 85.9	91.1	89.9	93.2	58.7	89.0	84.8
DIET-REL	per-head	input	86.0 / 86.1	91.0	89.8	92.8	59.6	89.0	84.9
DIET-REL	per-head	per-head	86.3 / 86.3	91.0	90.5	92.9	60.3	89.3	85.2
DIET-ABS ( $d_p=64$ )	per-head	input	86.1 / 85.8	91.2	90.0	93.0	58.9	89.9	85.0
DIET-ABS ( $d_p=64$ )	per-head	per-head	86.1 / 86.1	91.2	90.2	93.0	58.9	89.8	85.0
DIET-ABS ( $d_p=64$ , share)	per-head	per-head	86 / 86.8	91.1	90.4	92.9	59.3	89.8	85.2
DIET-ABS ( $d_p=128$ , share)	per-head	per-head	86.4 / 86.4	90.8	89.5	93.0	59.8	90.2	85.2

Table 8: Position and segment ablation study on GLUE: DIET-REL and DIET-ABS demonstrate the advantages of moving both positional and segment embedding from input to per-head.

Model	Sharing	MNLI 393k	QQP 364k	QNLI 105k	SST2 67k	CoLA 8.5k	STS-B 7k	Avg
DIET-REL	-	86.3 / 86.3	91.0	90.5	92.9	60.3	89.3	85.2
DIET-REL	layer-wise	86.5 / 86.3	91.1	90.0	93.0	58.8	89.6	85.0
DIET-REL	head-wise	85.8 / 85.7	91.2	90.2	92.8	59.8	89.1	84.9
DIET-ABS ( $d_p=64$ )	-	86.1 / 86.1	91.2	90.2	93.0	58.9	89.8	85.0
DIET-ABS ( $d_p=128$ )	-	86.7 / 86.5	91.2	90.6	92.8	60.1	89.4	85.3
DIET-ABS ( $d_p=64$ )	layer-wise	86 / 86.8	91.1	90.4	92.9	59.3	89.8	85.2
DIET-ABS ( $d_p=128$ )	layer-wise	86.4 / 86.4	90.8	89.5	93.0	59.8	90.2	85.2

Table 9: Sharing ablation study on GLUE: We run ablation study to understand the effect of sharing position encoding parameters across layers and heads. We notice that sharing improves the performance of DIET-ABS, but hurts the performance of DIET-REL with both layer-wise or head-wise sharing.