# Absolute Position Embedding Learns Sinusoid-like Waves for Attention Based on Relative Position 

Yuji Yamamoto and Takuya Matsuzaki<br>Tokyo University of Science<br>1423531@ed.tus.ac.jp matuzaki@rs.tus.ac.jp


#### Abstract

Attention weight is a clue to interpret how a Transformer-based model makes an inference. In some attention heads, the attention focuses on the neighbors of each token. This allows the output vector of each token to depend on the surrounding tokens and contributes to make the inference context-dependent. We analyze the mechanism behind the concentration of attention on nearby tokens. We show that the phenomenon emerges as follows: (1) learned position embedding has sinusoid-like components, (2) such components are transmitted to the query and the key in the selfattention, (3) the attention head shifts the phases of the sinusoid-like components so that the attention concentrates on nearby tokens at specific relative positions. In other words, a certain type of Transformer-based model acquires the sinusoidal positional encoding to some extent on its own through Masked Language Modeling.


## 1 Introduction

The architecture of Transformer (Vaswani et al., 2017) is symmetric with respect to the token position and it captures word order only through the position embedding included in the input. Thanks to this design, Transformer can flexibly learn relationships between tokens while allowing parallelization. To derive a representation of a context, previous language models have used, for instance, recurrent units to process tokens in the given token order, or convolution to aggregate tokens within a given range. In these models, the property of translation invariance of language has been captured through the architecture of the models. In contrast, Transformer receives all tokens at the same time and does not restrict the positions of the tokens on which each token depends; in exchange for it, the translation invariance has to be learned rather than imposed by the architecture.

Position embedding is thus the key to make inferences in Transformer context-dependent. While Transformer uses deterministic sinusoidal position embedding, BERT (Devlin et al., 2019) uses a learnable absolute position embedding. The latter learns the positional representation only through word cooccurrences. Clark et al. (2019) and Kovaleva et al. (2019) investigated the attention weights of the self-attention in BERT and found that the attention in some heads is largely determined by relative positions (Figure 1). This implies that even when the absolute position embedding is used, self-attention can make an inference depending on relative position.

Wang et al. (2021) compared various position embedding methods based on the performance in downstream tasks. Their results showed that local translation invariance and asymmetric position bias with respect to the direction improved the performance. Ravishankar and Søgaard (2021) observed that some columns of absolute position embedding were periodic. Chang et al. (2022) showed that position representation was periodic even in hidden representations. However, it is not clear how the periodicity is used in the model.

In this work, we analyze how attention depends on relative position. As a result, we show that the relative positional dependence of attention emerges due to the following factors.

- The learnable absolute position embedding has sinusoid-like waves with several limited frequencies (§4.1.1).
- Attention heads extract periodic components derived from position embedding from the hidden states. It is made explicit by applying Singular Value Decomposition to the parameters of the pre-trained language model (§4.2).
- The self-attention shifts the phase of the periodic components in the query and the key to decide the direction of attention ( $\$ 4.3, \S 4.4$ ).

Thus, it becomes partially clear how the selfattention equipped with learnable absolute position embeddings makes inferences based on context. However, it is suggested that, when the attention is strongly concentrated on the adjacent tokens, the word embedding is also a factor that enables inference based on relative position (§5).

## 2 Background

In this section, we review the multi-head selfattention mechanism and position embedding.

### 2.1 Multi-Head Self-Attention

The self-attention mixes the representations of tokens at different positions. The input to the $l$-th attention layer is a matrix $X_{l} \in \mathbb{R}^{T \times d}$ whose $k$ th row corresponds to the hidden representation of the $k$-th token as a vector in $\mathbb{R}^{d}$. The output of the self-attention layer is defined as follows:

$$
\begin{align*}
& A_{l h}=\operatorname{softmax}\left(\frac{X_{l} W_{l h}^{Q}\left(X_{l} W_{l h}^{K}\right)^{T}}{\sqrt{d / n}}\right)  \tag{1}\\
& O_{l h}=A_{l h} X_{l} W_{l h}^{V}  \tag{2}\\
& \text { MultiHead }_{l}=\operatorname{concat}\left(O_{l 1}, \ldots, O_{l n}\right) W_{l}^{O} \tag{3}
\end{align*}
$$

where $W_{l h}^{Q}, W_{l h}^{K}, W_{l h}^{V} \in \mathbb{R}^{d \times(d / n)}$ and $W_{l}^{O} \in$ $\mathbb{R}^{d \times d}$ are the parameters, $n$ is the number of attention heads per layer, and the subscripts $l$ and $h$ are the indices of the layer and the head, respectively.

In this paper, we refer to the $h$-th head in the $l$-th layer as "head $(l, h)$ " for short, and omit the subscripts $l, h$ when the discussions are common to all layers and heads. The matrices $X W^{Q}, X W^{K}$, and $A$ are called query, key, and attention weight, respectively. If the $(i, j)$ element of $A$ is large, it is interpreted that the $i$-th token in the sentence attends to the $j$-th token.

### 2.2 Position Embedding

Transformer's position embedding is defined as follows (Sinusoidal Position Embedding; SPE):

$$
\begin{align*}
S P E_{(p o s, 2 i)} & =\sin \left(p o s / 10000^{2 i / d}\right)  \tag{4}\\
S P E_{(p o s, 2 i+1)} & =\cos \left(p o s / 10000^{2 i / d}\right) \tag{5}
\end{align*}
$$

Vaswani et al. (2017) hypothesized that it allows the model to easily learn to attend by relative positions, since the offset between two tokens can be expressed as a matrix product:

$$
\begin{align*}
& {\left[\begin{array}{ll}
\sin x_{i} & \cos x_{i}
\end{array}\right]\left[\begin{array}{cc}
\cos \theta & -\sin \theta \\
\sin \theta & \cos \theta
\end{array}\right] } \\
= & {\left[\begin{array}{ll}
\sin \left(x_{i}+\theta\right) & \cos \left(x_{i}+\theta\right)
\end{array}\right] . } \tag{6}
\end{align*}
$$

BERT's position embedding is a learnable parameter. It is called the Absolute Position Embedding (APE) because each row represents an absolute position in the input sequence.

RoBERTa (Liu et al., 2019) is identical to BERT in architecture. RoBERTa is pre-trained with 512 tokens for all steps whereas BERT is pre-trained with 128 tokens for $90 \%$ of the steps. Hence, RoBERTa's APE recieves the same number of updates regardless of the positions. Wang and Chen (2020) showed that RoBERTa's APE is more orderly than BERT's. We thus use RoBERTa for all experiments to analyze the position dependence of inference.

## 3 Relative Position Dependence of Attention

In this section, we demonstrate that the attention depends on the relative position. Specifically, we analyze in how many heads the attention is dependent on the relative position, and the variation in the direction and concentration of the attention across the heads. The visualizations of the attention weight $A_{l h}$ in Figure 1 shows that, in (a) and (d), each token strongly attends only to the adjacent token, while in (b) and (e), the attention is put on a few nearby tokens on the left or right side. We show that such a pattern is largely invariant with respect to the inputs, by clustering the patterns of the attention on different inputs.

To focus on the relation between the attention and the relative positions, we summarized the attention weight for each input in each head into a vector and applied k-means clustering to them. For that, we defined $t$-offset $\operatorname{trace} \mathrm{tr}_{t}$ as follows:

$$
\operatorname{tr}_{t}(A)= \begin{cases}\sum_{i=1}^{T-t}(A)_{i, i+t} & (t \geq 0)  \tag{7}\\ \sum_{i=1}^{T+t}(A)_{i-t, i} & (t<0)\end{cases}
$$

and transformed $A_{l h}$ to a vector:

$$
\begin{equation*}
\boldsymbol{a}_{l h}=\left[\operatorname{tr}_{-10}\left(A_{l h}\right), \ldots, \operatorname{tr}_{10}\left(A_{l h}\right)\right] \in \mathbb{R}^{21} \tag{8}
\end{equation*}
$$

These vectors represent the trend of attention $A_{l h}$ with respect to relative position.

We input 100 sentences of length 512 created from wikitext-2 (Merity et al., 2017) into RoBERTa and computed vectors $\boldsymbol{a}_{l h}$ for every head. Figure 2 shows the results of applying the k-means to a total of $100 \times 12 \times 12$ vectors $\boldsymbol{a}_{l h}$ when the number of clusters was set to 6 . In the clusters named leftward and next-to-left


Figure 1: Attention weights for the first 15 tokens of an input text.
(resp. rightward and next-to-right), the attention is concentrated on the nearby tokens on the left (resp. right). Particularly, in the clusters named next-to-left and next-to-right, the attention is strongly concentrated on the adjacent tokens. We found that, for each head, the vectors $\boldsymbol{a}_{l h}$ corresponding to the 100 inputs were grouped into one or two clusters. This means that there are some heads that always attended to the same direction for all the input.

## 4 Attention to Nearby Tokens

In this section, we show that the attention depends on relative position due to periodic components in the position embeddings. First, we show that the learnable absolute position embeddings acquire several limited frequency components. Second, we show that some attention heads extract the periodicity derived from position embeddings. Finally, we show that the concentration of attention on the nearby tokens is due to the periodicity.

### 4.1 Learned Representation of Positions

First, we show that APE includes sinusoid-like components using the Discrete Fourier Transform (DFT). Next, we show that the position embeddings are confined to a relatively low-dimensional subspace ( $\sim 15$ dimensions) using Principal Component Analysis (PCA). Finally, we show that the encoded positional information occupies a similar number of dimensions in the hidden states as in the position embedding, and the dimensionality becomes smaller in the higher layers using Canonical Correlation Analysis (CCA).

### 4.1.1 APE Includes Sinusoid-like Waves

We view the RoBERTa's position embedding $E_{p o s} \in \mathbb{R}^{T \times d}$ as a collection of $d$ time series of


Figure 2: For each head, a pie chart illustrates the percentage of labels assigned to 100 vectors $\boldsymbol{a}_{l h}$ (Eq. (7)).
length $T$. We computed the amplitude spectrum of each column of $E_{p o s}$ by applying DFT:

$$
\begin{equation*}
\operatorname{spec}_{i}=\operatorname{abs}\left(\operatorname{DFT}\left(E_{p o s} \boldsymbol{e}_{i}\right)\right) \tag{9}
\end{equation*}
$$

where $e_{i} \in \mathbb{R}^{d}$ is the $i$-th elementary unit vector. The vector $E_{p o s} \boldsymbol{e}_{i}$ is hence the $i$-th column of $E_{\text {pos }}$.

Figure 3a shows the mean and the quartile range of the amplitude spectra of $d=768$ time series. The amplitude spectra peak at some frequencies, indicating that the learnable APE has acquired periodicity through pre-training, even though it is not explicitly defined using sinusoidal waves as in Transformer. In contrast, there are no peaks in the amplitude spectrum of a word embedding sequence of a sample text. ${ }^{1}$ The periodicity is thus an unique property of learned position embeddings.

We investigated whether similar properties are present in pre-trained models other than RoBERTa (Figure 3b and 13). The amplitude spectra of the encoder-only models are similar to RoBERTa regardless of language, but ones of GPT-2 (Radford et al., 2019) are higher at lower frequencies.

The decoder model probably can focus attention on the neighboring tokens of itself without periodicity-based mechanisms (§4.3). For example, if attention is paid more strongly to the backward token, in BERT, attention is focused on the end of the sentence, but in a decoder with causal attention mask, attention is focused on itself. We leave it to future research to clarify whether this

[^0]

Figure 3: Amplitude spectra of position embeddings. Blue line is the mean of $\operatorname{spec}_{i}$ and blue area is the quartile range. Orange line is the spectrum of word embeddings for an input before adding position embeddings.
phenomenon really occurs in the causal model, and in this paper we focus on the case where selfattention is symmetric with respect to position.

### 4.1.2 Dimensionality of Positional Representation

We applied PCA to the position embeddings to see how many dimensions were used to represent the positions. The cumulative principal component contribution rate of the position embedding was $50.51 \%$ for the top 4 dimensions, and $92.23 \%$ for the top 12 dimensions (see Figure 12 in Appendix A for a plot). This means that the positions are mostly encoded in a low dimensional subspace.

We then employed CCA to show how much the input to the self-attention layers included the positional representation. CCA is a method for investigating the relation of two sets of variables by finding the parameters $\boldsymbol{a}$ and $\boldsymbol{b}$ that maximizes the correlation between two synthetic variables $X \boldsymbol{a}$ and $Y \boldsymbol{b}$ given two inputs $X$ and $Y$. Raghu et al. (2017) showed that CCA allows layer-to-layer comparison of deep learning models. We used the representation of neurons and layers proposed by them and computed the correlation between the hidden layer and the position embeddings.

In this study, the $i$-th neuron $\boldsymbol{z}_{i}^{l}$ of layer $l$ and the $l$-th layer $Z^{l}$ are represented as follows:

$$
\begin{align*}
\boldsymbol{z}_{i}^{l} & =\left[z_{i}^{l}\left(\boldsymbol{x}_{1}\right), \ldots, z_{i}^{l}\left(\boldsymbol{x}_{m}\right)\right]^{T}  \tag{10}\\
Z^{l} & =\left[\boldsymbol{z}_{1}^{l}, \ldots, \boldsymbol{z}_{n}^{l}\right] \tag{11}
\end{align*}
$$



Figure 4: CCA coefficient for each layer compared to position embedding.
where $z_{i}^{l}\left(\boldsymbol{x}_{j}\right)$ is the response of the $i$-th neuron to input $\boldsymbol{x}_{j}$. We input 200 sentences of length 512 created from wikitext-2 into RoBERTa and collected the responses of each layer for the input of $m=200 \times 512$ tokens. We then maximized their correlation coefficients $\rho_{k}$ :

$$
\begin{equation*}
\rho_{k}=\max _{\boldsymbol{a}_{k}, \boldsymbol{b}_{k}} \operatorname{Corr}\left(Z^{l_{1}} \boldsymbol{a}_{k}, Z^{l_{2}} \boldsymbol{b}_{k}\right) \tag{12}
\end{equation*}
$$

where $\boldsymbol{a}_{k}$ is chosen such that it is orthogonal to $\boldsymbol{a}_{1}, \ldots, \boldsymbol{a}_{k-1}$ and similarly for $\boldsymbol{b}_{k}$. The CCA coefficients in Figure 4 show that the hidden states of higher layers have lower correlations with the position embeddings. This is in agreement with Lin et al.'s (2019) result that BERT phases out positional information as it passes through the layers.

The CCA result indicates that the number of components that are highly correlated with position embedding is only $5 \sim 20$-dimensions, and the PCA result suggests that it is enough to accommodate the positional representation. This indicates that the hidden states include positional representation in a low-dimensional subspace, similarly to position embedding.

### 4.2 Positional Representation in Self-Attention

It is not at all clear how the positional representation in the hidden states contribute to the inference in self-attention. We thus analyzed how attention weight is calculated, which is one of the most important process in self-attention.

### 4.2.1 Rethinking About Query and Key

The attention weight $A$ is determined by the inner product of the rows of the query and the key matrix. We thus expect that the relative position dependence of the attention can be explained by analyzing the relationship between the query and


Figure 5: Column $i$ of the query and the key in head $(1,1)$.
the key. We begin by rethinking about their definition. The two parameter matrices $W^{Q}$ and $W^{K}$ contribute to the attention only through their product $W^{Q}\left(W^{K}\right)^{T}$ (Eq. (1)). Hence, they can be seen as a parametrization of a rank-restricted $d \times d$ matrix, and we may consider another decomposition of the product $W^{Q}\left(W^{K}\right)^{T}$.

We found that some sinusoid-like components were obtained from the query and the key by applying Singular Value Decomposition (SVD) to $W^{Q}\left(W^{K}\right)^{T}$. Specifically, the query and the key are redefined as the product of the hidden state and the singular vectors:

$$
\begin{gather*}
W^{Q}\left(W^{K}\right)^{T}=U^{Q} S\left(U^{K}\right)^{T}  \tag{13}\\
Q=X U^{Q}, \quad K=X U^{K} \tag{14}
\end{gather*}
$$

where the matrix $S \in \mathbb{R}^{d \times d}$ is a diagonal matrix $\operatorname{diag}\left(s_{1}, \ldots, s_{d / n}, 0, \ldots, 0\right)$ and each element $s_{i}$ is a singular value. In the rest of this paper, we refer to $Q$ and $K$ defined above as the query and the key, respectively. As shown in Figure 5, in the head in which the attention depends on relative position, sinusoid-like waves appear in several columns of the redefined query $Q$ and the key $K$. Moreover, a sine wave is paired with a cosine wave of the same frequency, as the 5 -th and the 6 -th column for head $(1,1)$ shown in Figure 5.

Furthermore, the introduction of SVD provides a new theoretical interpretation of self-attention. Let $R$ be the orthogonal matrix $R=\left(U^{Q}\right)^{T} U^{K}$. Then $U^{K}$ can be written as $U^{K}=U^{Q} R$ due to the orthogonality of the singular vectors. Thus, the key $K$ can be written as:

$$
\begin{equation*}
K=X U^{K}=X U^{Q} R=Q R \tag{15}
\end{equation*}
$$

That is to say, the rows of the key are the result of an orthogonal transformation of the rows of the


Figure 6: The frequency-wise maxima of the amplitude spectra. The $l$-th row of the heat map corresponds to max-spec ${ }_{l}$. The top figure is a reiteration of Figure 3a.
query. The relation between the query and the key is thus summarized in the matrix $R$.

In addition, since the singular value matrix $S$ is diagonal, the product of the query and the key can be written as follows:

$$
\begin{equation*}
X W^{Q}\left(X W^{K}\right)^{T}=Q S K^{T}=\sum_{i=1}^{d / n} s_{i} \boldsymbol{q}_{i} \boldsymbol{k}_{i}^{T} \tag{16}
\end{equation*}
$$

where $\boldsymbol{q}_{i}$ and $\boldsymbol{k}_{i}$ are the $i$-th columns of $Q$ and $K$, respectively. Eq. (16) says that the subsets of queries $\left[\boldsymbol{q}_{i}\right]$ and keys $\left[\boldsymbol{k}_{i}\right]$ corresponding to the top singular values $\left[s_{i}\right]$ are more important in the calculation of the attention distribution. We hereafter call the matrix $Q S K^{T}$ attention score (i.e., attention before applying softmax).

### 4.2.2 Spectral Analysis of Query and Key

We computed the amplitude spectrum of each column of the query $Q_{l h}$ in each head, using the same procedure for the position embeddings in §4.1.1:

$$
\begin{align*}
\operatorname{spec}_{l h i} & =\operatorname{abs}\left(\operatorname{DFT}\left(Q_{l h} \boldsymbol{e}_{i}\right)\right)  \tag{17}\\
\max -\operatorname{spec}_{l} & =\left[\max _{h, i}\left(\operatorname{spec}_{l h i}\right)_{f}\right]_{f=0, \ldots, T / 2} \tag{18}
\end{align*}
$$

By taking the maximum of $\operatorname{spec}_{l h i}$ among the heads and the columns of $Q_{l h}$, we check if there are high peaks in the spectra of the queries in the $l$-th layer (similarly for the keys and hidden states).

Figure 6 show that the query and key spectra peak at the same frequency bands as the position embeddings, indicating that attention heads extract periodic representations derived from the position embeddings. Furthermore, some of the peaks disappear in the higher layers (downward in the Figure 6), indicating that the periodic components derived from the position embeddings are not dominant in the higher layers. This is consistent with the result that the correlation between the position embeddings and the hiddem states of each layer gradually decreases (§4.1.2). It also agrees with the result shown by Lin et al. (2019) that position information is discarded between the 3rd and 4th layers.

### 4.3 Attention Based on Relative Position is due to the Phase Shift

As shown in Figure 5, there are phase shifts in the sinusoid-like components of the query and the key. In this subsection, we focus on the phase shifts, and clarify that the direction and the width of the phase shifts determine where the attention is concentrated. We measure the direction and the width of the phase shift by cross-covariance and crosscorrelation, defined as follows:

$$
\begin{align*}
\operatorname{xcov}_{j}(t) & = \begin{cases}\sum_{i=1}^{T-t} q_{i, j} k_{i+t, j} & (t \geq 0) \\
\sum_{i=1}^{T+t} q_{i-t, j} k_{i, j} & (t<0)\end{cases}  \tag{19}\\
\operatorname{xcorr}_{j}(t) & =\frac{\left.\operatorname{xcov}_{j}(t)-\mathbb{E}_{t} \operatorname{xcov}_{j}(t)\right]}{\left\|\boldsymbol{q}_{j}\right\|\left\|\boldsymbol{k}_{j}\right\|} \tag{20}
\end{align*}
$$

For example, Figure 7 shows the cross-correlation between the queries and keys in Figure 5. Both $\operatorname{xcorr}_{5}(t)$ and $\operatorname{xcorr}_{6}(t)$ attain a maximal value at $t=-2$ and $t=-3$. It means that the phase shift between the queries and the keys are approximately 2.5 tokens.

It can be explained theoretically that the direction of phase shift coincides with the direction of attention. The cross-covariance is related to the product of query and key (Eq. (16)):

$$
\begin{align*}
\operatorname{tr}_{t}\left(Q S K^{T}\right) & =\sum_{i=1}^{T-t}\left(Q S K^{T}\right)_{i, i+t}  \tag{21}\\
& =\sum_{i=1}^{T-t} \sum_{j=1}^{d / n} s_{j} q_{i, j} k_{i+t, j}  \tag{22}\\
& =\sum_{j=1}^{d / n} s_{j} \mathrm{xcov}_{j}(t) \tag{23}
\end{align*}
$$



Figure 7: The cross-correlation $\operatorname{xcorr}_{i}(t)$ between the queries and the keys in Figure 5.


Figure 8: The sum of cross-covariances weighted by singular values (Eq. (23))



Figure 9: The cross-correlations xcorr $_{j}(t)$ corresponding to the top 10 singular values.
where $\mathrm{tr}_{t}$ is the $t$-offset trace defined in Eq. (7) and $s_{1}, \ldots, s_{d / n}$ are the singular values of $W^{Q}\left(W^{K}\right)^{T}$. According to Eq. (21-23), the sum of the attention scores at relative position $t$ is equal to the weighted sum of the inner product of $\boldsymbol{q}_{j}$ and $\boldsymbol{k}_{j}$ shifted by $t$. Therefore if the $\operatorname{xcov}_{j}(t) \mathrm{s}$ corresponding to the top singular values attain maximal values at the same $t$, the attention is likely to be concentrated around relative position $t$.

Figure 8 and 9 show that it is actually the case. Figure 9 presents the cross-correlations of the query $\boldsymbol{q}_{j}$ and the key $\boldsymbol{k}_{j}$ for $j$ s corresponding to the top singular values. We can see how the concentration of attention on nearby tokens is formed by the weighted superposition of the crosscorrelations. In head $(1,1)$, the maxima of the cross-correlations near $t=-1$ stack up, while the maxima away from $t=-1$ cancel with the minima of other components. Since there are multiple periodic components with different frequencies (Figure 3a), the attention is not concentrated away from each token. In contrast, in head (8, 9), some cross-correlations have narrow peaks only at $t=-1$, and it makes the attention be very concentrated only on the adjacent tokens. However, it is unlikely that the cause of the narrow peaks
is the sinusoid-like waves, because their period is approximately 8 tokens or more. ${ }^{2}$

### 4.4 Phase Shift Width is the Same even if Frequencies are Different

We saw in the previous subsection that the main factor of the concentration of attention is that the phases of multiple sinusoid-like waves are all shifted by the same number of tokens. In this subsection, we explain it based on the property of the eigenvalues and the eigenvectors of the matrix $R$ that relates the query $Q$ to the key $K$.

Let $\boldsymbol{p}_{i} \in \mathbb{C}^{d}$ and $\lambda_{i} \in \mathbb{C}$ be an eigenvector and an eigenvalue of the matrix $R$, respectively. Since $R$ is an orthogonal matrix, its eigenvalue can be expressed as $\lambda_{i}=e^{j \theta_{i}}$ where $j$ is the imaginary unit. From Eq. (15), the following relation between the query and the key holds:

$$
\begin{equation*}
K \boldsymbol{p}_{i}=Q R \boldsymbol{p}_{i}=Q \lambda_{i} \boldsymbol{p}_{i}=Q \boldsymbol{p}_{i} \cdot e^{j \theta_{i}} \tag{24}
\end{equation*}
$$

Let two conditions be assumed: (1) for each $i=1, \ldots, d$, the time series $Q \boldsymbol{p}_{i}$ is sinusoidal with a single frequency $f_{i}$ and (2) the ratio of the argument $\theta_{i}$ of the eigenvalue $\lambda_{i}$ to the frequency $f_{i}$ is constant regardless of $i$. Then Eq. (24) implies that the phases of $Q \boldsymbol{p}_{i}$ and $K \boldsymbol{p}_{i}$ differ by a constant number of tokens for any eigenvector $\boldsymbol{p}_{i}$ :

$$
\begin{equation*}
\left(Q \boldsymbol{p}_{i}\right)_{t}=\left(K \boldsymbol{p}_{i}\right)_{t+\Delta} \tag{25}
\end{equation*}
$$

where $\Delta=(T / 2 \pi) \cdot\left(\theta_{i} / f_{i}\right)$. Furthermore, $Q_{t}=$ $K_{t+\Delta}$ follows from Eq. (25). We provide the proofs in Appendix D.

We verify whether the two conditions hold by analyzing the periodicity of $Q \boldsymbol{p}_{i}$ and the ratio of the frequency to the argument $\theta_{i}$ of the eigenvalue $\lambda_{i}$. To do so, we define the bivariate functions $g$ for frequency and argument as follows:

$$
\begin{equation*}
g\left(f, \theta_{i}\right)=\operatorname{abs}\left(\operatorname{DFT}\left(Q \boldsymbol{p}_{i}\right)\right)_{f} \tag{26}
\end{equation*}
$$

This function $g$ is shown in Figure 10 as a 2D histogram. Figure 10 shows that the spectrum of $Q \boldsymbol{p}_{i}$ has peaks in different frequency bands according to $\theta_{i}$. It means that the component along each eigenvector $\boldsymbol{p}_{i}$ is fairly band-limited, namely they are sinusoid-like. Furthermore, the spectral peaks

[^1]

Figure 10: 2D histogram of the function $g(f, \theta)$. The section at $\theta=\theta_{i}$ is the amplitude spectrum of $Q \boldsymbol{p}_{i}$.
appear around a straight line. Specifically, the ratio $\Delta(f, \theta)$ of frequency to phase defined as follows is nearly constant:

$$
\begin{equation*}
\Delta(f, \theta)=\frac{T}{2 \pi} \cdot \frac{\theta}{f} \tag{27}
\end{equation*}
$$

Hence, the two conditions are in fact approximately hold. Thus, by multiplying the matrix $R$, the sinusoid-like components in the key are offset by a certain number of tokens relative to the query.

## 5 Attention to the Adjacent Tokens

Figure 9 shows that, in head $(8,9)$, there are not only sinusoid-like waves but also spikes at the relative position $t=-1$. In this section, we dive deeper into the fact that, in some heads, the attention is focused only on the adjacent token. We analyze how the attention changes when we modify either the position embedding or the word embedding component. The following are the results on the modified inputs and insights from them.
Baseline The input is the sum of word embedding and position embedding as usual. This result is shown in Figure 8 and 9.

Only position embedding The word embeddings are replaced with zero vectors. Figure 11a shows that spikes of cross-correlation do not appear for this case. This suggests that position embedding is not the only cause of the spikes.
Only word embedding The position embeddings are replaced with zero vectors. Figure 11 b shows that most of the cross-correlations are flat for this case. It suggests that word embedding contributes less to make the attention dependent on relative position. However,
it is interesting that the cross-covariance at $t=-1$ (i.e., left neighbor) is relatively large even without position information.

Shuffling word embeddings The order of word embeddings is shuffled. Figure 11c shows that spikes appear at the relative position $t=$ -1 even for this case. It suggests that the contribution of position embedding is significant in determining the direction of attention, since the attention is focused on the preceding word regardless of what it actually is.

As mentioned at the end of Section 4.3 and also suggested by Figure 11a, it is unlikely that position embedding is the only cause of the strong concentration of attention on the adjacent tokens. However, if the word embeddings identified the adjacent tokens (i.e., if the narrow peak of cross-correlation appeared due to an interaction of the word embeddings of a frequently occurring bigram), the attention would have been moved to non-neighboring positions by the shuffling of word embeddings, but this was not the case. It is thus suggested that the concept of adjacency in RoBERTa is built upon both word meaning and positional information.

## 6 Remark on the Learning Process of Position Embeddings

The training data of the masked language modeling is only word sequences without positional information. It suggests that the relative position dependence of attention is acquired by the combination of two factors: (1) the linguistic property that related words tend to appear nearby due to grammatical rules and collocations, and (2) the property that attention is focused on the word that is syntactically or semantically related and hence gives clue to fill the masked token. In appendix E, we demonstrate that the position representation can be acquired with a much smaller amount of training data than pre-training.

## 7 Related Works

Ravishankar and Søgaard (2021) observed that some columns of absolute position embedding were periodic and sinusoidal position embedding was more effective in multilingual models than other embedding methods. Chang et al. (2022) showed that position representation was periodic even in hidden representations by using Linear


Figure 11: The cross-correlations (left) and the weighted sum of cross-covariances (right) in head $(8,9)$ when we modify either the position embedding or the word embedding component of the input.

Discriminant Analysis, i.e., by identifying the axes separating the different position representations. We showed that sinusoid-like components could be extracted from the hidden states by applying SVD to the pre-trained parameters of the selfattention even though SVD does not have the objective of separating positional representations.

Vaswani et al. (2017) stated that the sinusoidal position embedding allowed position offsets to be represented by a rotation transformation (Eq. (6)), and this could prompt learning of attention that depends on relative position. We showed that the query and the key in the self-attention included sinusoid-like components and the learned parameters in the self-attention shift the phase of the query and the key relatively. This means that the mechanism hypothesized by Vaswani et al. is in fact acquired through pre-training by masked language modeling. However, the frequencies acquired by the position embedding of RoBERTa are only in the specific frequency bands, whereas sinusoidal position embedding has $d / 2$ frequencies (Eq. (4)-(5)). RoBERTa thus seems to have acquired a more economical positon embedding than sinusoidal position embedding.

Recently, there are many variants of position
embedding and Dufter et al. (2022) surveyed them exhaustively. In particular, Rotary Position Embedding (RoPE) (Su et al., 2022), a type of Relative Position Embedding (RPE) (Shaw et al., 2018), relates to the property that the self attetnion acquires the rotation matrix through pre-training. To acquire relative positional dependence, RoPE widens the angle between query and key proportionally to the relative position, while pre-trained self-attention rotates the hidden states containing absolute positional bias by the same angle regardless of position. In other words, APE and selfattention, which are distant components, must acquire frequency and angle of rotation, respectively, to satisfy the relation Eq. (27). If translation invariance is essential to language understanding, the cost of discovering this relation is a possible reason why APEs are inefficient to learn compared to RPEs.

## 8 Conclusion

We analyzed the concentration of attention based on relative position in the self-attention using the learnable absolute position embedding. As a result, we showed that it is due to relative phase shifts of the periodic components in the query and the key derived from position embeddings. Our results explain in part that absolute position embedding allows inference based on relative position.

## 9 Limitations

As mentioned in §4.1.1, the positional representation of the GPT-2 differs significantly from ones of the encoder models. Thus, it is impossible to interpret the inferences of the decoder-only model in the same way as in this study.

In the tasks of predicting sentence structure (e.g., dependency parsing), the relative order of two tokens is important regardless of the distance between them. However, we analyzed the dependence of the output of each token only on the nearby tokens. Thus, it is unclear whether position embedding provides relative position information that helps determine which of the distant tokens precedes the other.

We obtained an interpretable representation by decomposing the attention scores before applying softmax function (Eq. (16)). When analyzing the contribution of the decomposed representations to downstream components (e.g., Value in
self-attention), the non-linearity of softmax function should be taken into account.

## References

Tyler Chang, Zhuowen Tu, and Benjamin Bergen. 2022. The geometry of multilingual language model representations. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, pages 119-136, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. What does BERT look at? an analysis of BERT's attention. In Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP, pages 276-286, Florence, Italy. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171-4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Philipp Dufter, Martin Schmitt, and Hinrich Schütze. 2022. Position information in transformers: An overview. Computational Linguistics, 48(3):733763.

Kazuki Irie, Albert Zeyer, Ralf Schlüter, and Hermann Ney. 2019. Language Modeling with Deep Transformers. In Proc. Interspeech 2019, pages 39053909.

Amirhossein Kazemnejad, Inkit Padhi, Karthikeyan Natesan Ramamurthy, Payel Das, and Siva Reddy. 2023. The impact of positional encoding on length generalization in transformers. arXiv preprint arXiv:2305.19466.

Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of BERT. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 4365-4374, Hong Kong, China. Association for Computational Linguistics.
Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pretraining for natural language generation, translation, and comprehension. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 7871-7880, Online. Association for Computational Linguistics.

Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. Open sesame: Getting inside BERT's linguistic knowledge. In Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP , pages 241-253, Florence, Italy. Association for Computational Linguistics.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692.

Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. Pointer sentinel mixture models. In International Conference on Learning Representations.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. 2017. SVCCA: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc.

Vinit Ravishankar and Anders Søgaard. 2021. The impact of positional encodings on multilingual compression. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, pages 763-777, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

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, New Orleans, Louisiana. Association for Computational Linguistics.

Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. 2022. Roformer: Enhanced transformer with rotary position embedding.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc.

Benyou Wang, Lifeng Shang, Christina Lioma, Xin Jiang, Hao Yang, Qun Liu, and Jakob Grue Simonsen. 2021. On position embeddings in BERT. In International Conference on Learning Representations.

Yu-An Wang and Yun-Nung Chen. 2020. What do position embeddings learn? an empirical study of pre-trained language model positional encoding.

In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 6840-6849, Online. Association for Computational Linguistics.

## A Cumulative Principal Component Contribution Rate



Figure 12: Cumulative principal component contribution rate of RoBERTa position embedding.

## B Amplitude Spectra of Various Models

We computed the amplitude spectra of the position embeddings of bert-base-uncased, cl-tohoku/bert-base-japanese-whole-word -masking, xlm-roberta-base, and gpt2 in the same way as Figure 3a. In Figure 13, the encoder models BERT and RoBERTa both have peaks in their amplitude spectra, and RoBERTa has a higher peaks. On the other hand, the GPT-2 decoder model has only low-frequency components. This suggests that the representation of absolute position embedding is similar between encoder models, regardless of language, but differs significantly between encoder models and decoder models. In fact, Irie et al. (2019) and Kazemnejad et al. (2023) showed that the position embeddings is unnecessary for the decoder models.

## C Comparing Different Architectures

This paper analyzed an encoder-only model (RoBERTa). In this section, we apply our methods to other transformer-based architectures: decoderonly and encoder-decoder. The target pre-trained models are GPT-2 (gpt2) for decoder-only and BART (facebook/bart-base) (Lewis et al., 2020) for encoder-decoder, both of which use absolute position embedding and are available on huggingface. To reduce the gap between architectures, when analyzing the decoders, we use the attention scores, which is the matrix before causal masking and softmax function are applied, instead of the attention weights.

The attention matrices of GPT-2 show two main patterns: one related to position and the other not. Figure 14a shows that the attention head of GPT2 pays stronger attention to the closer tokens by


Figure 13: Amplitude spectra of position embedding of various models.
concentrating attention on the backward tokens, which are masked in the subsequent processing. Figure 15 shows the result of clustering the attention scores of GPT-2 using k-means. The number of clusters was set to 2 and the inputs to k -means were $t$-offset traces (Eq.(8)) from $t=-30$ to 30 . The heads that depend on position are found in the lower layers.

Along §4.4, we investigate the relationship between the frequency of the hidden state and the angle of the rotation matrix inherent in the parameters of attention head. Figure 16a differs from the case of RoBERTa (Figure 10) in that the peaks appear horizontally rather than holding a constant ratio between frequency and angle. Furthermore, according to Figure 13, the dominant components in position embedding of GPT-2 are those with frequencies below 10 , i.e., with periods longer than $51.2(=512 / 10)$ tokens. These differences from RoBERTa imply that even when the attention is focused within a few tokens in a particular head, it cannot be attributed to the mechanism described in §4.3.

In BART, the tendency in the positional dependence of attention differs between encoder and
decoder. In encoder, the patterns of attention weights (Figure 17) and their trend (Figure 18), such as direction and strength, are similar to those of RoBERTa (Figure 1 and 2). On the other hand, comparing Figure 14 and Figure 20, the patterns of decoder is different from those of GPT-2. Such patterns in the BART decoder could not be clustered based on the strength or direction of attention using k-means. Like in GPT-2, Figure 21 shows that the peaks appear horizontally in the decoder.
Figures for GPT-2


Figure 14: Attention scores of GPT-2 for the first 300 tokens of an input text.


Figure 15: Trend of attention scores in GPT-2.


Figure 16: 2D histogram of the function $g(f, \theta)$ in GPT-2.

Figures for BART encoder


Figure 17: Attention weights of BART encoder for the first 15 tokens of an input text.


Figure 18: Trend of attention weights in BART encoder.


Figure 19: 2D histogram of the function $g(f, \theta)$ in BART encoder.

Figures for BART decoder

(a) head $(2,5)$

(b) head $(3,6)$

(c) head( 6,6 )

Figure 20: Attention scores of BART decoder for the first 300 tokens of an input text.


Figure 21: 2D histogram of the function $g(f, \theta)$ in BART decoder.

## D Theorems About the Phase Shift Between the Query and the Key

Theorem 1. The following two assumptions are given:

1. $Q \boldsymbol{p}_{i}$ is sinusoidal with a single frequency $f_{i}$.
2. The ratio of the argument $\theta_{i}$ of an eigenvalue $\lambda_{i}$ to the frequency $f_{i}$ is constant regardless of $i$.

In this case, Eq. (25) is derived from Eq. (24):

$$
\begin{align*}
& K \boldsymbol{p}_{i}=Q \boldsymbol{p}_{i} \cdot e^{j \theta_{i}}  \tag{24}\\
\Longrightarrow & \left(Q \boldsymbol{p}_{i}\right)_{t}=\left(K \boldsymbol{p}_{i}\right)_{t+\Delta} \tag{25}
\end{align*}
$$

Proof. Since $Q \boldsymbol{p}_{i}$ is sinusoidal, the $t$-th element in polar form is as follows:

$$
\begin{equation*}
\left(Q \boldsymbol{p}_{i}\right)_{t}=r_{i} \exp \left[j 2 \pi f_{i} t / T\right] \tag{28}
\end{equation*}
$$

where $r_{i}$ is the absolute value of $\left(Q \boldsymbol{p}_{i}\right)_{t}$. Then the $t$-th element of $K \boldsymbol{p}_{i}$ can be written as follows from Eq. (24):

$$
\begin{align*}
\left(K \boldsymbol{p}_{i}\right)_{t} & =r_{i} \exp \left[j 2 \pi f_{i} t / T\right] \cdot \exp \left[j \theta_{i}\right]  \tag{29}\\
& =r_{i} \exp \left[j\left(2 \pi f_{i} t / T+\theta_{i}\right)\right]  \tag{30}\\
& =r_{i} \exp \left[j 2 \pi f_{i}\left(t+\Delta_{i}\right) / T\right] \tag{31}
\end{align*}
$$

where $\Delta_{i}=(T / 2 \pi) \cdot\left(\theta_{i} / f_{i}\right)$. Now, since the ratio of the argument $\theta_{i}$ to the frequency $f_{i}$ is constant, $\Delta_{i}$ is independent of $i$, namely, $\Delta_{i}$ can be replaced with a constant $\Delta$. Thus, $\left(Q \boldsymbol{p}_{i}\right)_{t}=\left(K \boldsymbol{p}_{i}\right)_{t+\Delta}$ holds from Eq. (28) and Eq. (31).

Let the matrix $P=\left[\boldsymbol{p}_{1}, \boldsymbol{p}_{2}, \ldots, \boldsymbol{p}_{d}\right]$, then the following equation holds:

$$
\begin{equation*}
(Q P)_{t}=(K P)_{t+\Delta} \tag{32}
\end{equation*}
$$

Theorem 2. When, in a given basis, the position of each token in the key is offset by $\Delta(\in \mathbb{Z})$ tokens relative to the query, then in any basis, there is an
offset of $\Delta$ tokens between the query and the key. That is,

$$
\begin{array}{r}
\exists P \in \mathbb{C}^{d \times d} \text { s.t. }|P| \neq 0 \text { and }(Q P)_{t}=(K P)_{t+\Delta} \\
\Longrightarrow \forall B \in \mathbb{C}^{d \times n},(Q B)_{t}=(K B)_{t+\Delta} \tag{34}
\end{array}
$$

where $Q, K \in \mathbb{R}^{T \times d}$ are the query and the key respectively and $n \in \mathbb{N}$ is less than or equal to $d$. The subscript $t$ indicates the $t$-th row of the matrix.

Proof. Let $B$ be an arbitrary matrix. Then,

$$
\begin{align*}
(Q B)_{t i} & =\left(Q P P^{-1} B\right)_{t i}  \tag{35}\\
& =\sum_{k=1}^{d}(Q P)_{t k}\left(P^{-1} B\right)_{k i}  \tag{36}\\
& =\sum_{k=1}^{d}(K P)_{t+\Delta, k}\left(P^{-1} B\right)_{k i}(\because \text { Eq. }  \tag{33}\\
& =\left(K P P^{-1} B\right)_{t+\Delta, i}  \tag{38}\\
& =(K B)_{t+\Delta, i}
\end{align*}
$$

In particular, when $B=I$, the equation $Q_{t}=$ $K_{t+\Delta}$ is obtained.

## E How the Relative Position Dependence of Attention Emerges

How does a masked language model acquire the concept of "neighborhood" even though absolute position embedding learns without the information about the order of the tokens? In this section, we demonstrate the process of acquiring the ability to focus attention on nearby tokens by relearning the position embedding under the masked language modeling.

Suppose that the following sentence is input into a pre-trained model:

```
"This <unk> is <unk> an <unk> <mask> ."
```

Since this sentence is collapsed with many UNK tokens, the model cannot fill the MASK correctly. If we re-learn the position embedding, will the model be able to fill it?


Figure 22: Attention weight with relearned position embedding

| PE | predicted top-5 tokens |
| :--- | :--- |
| relearned | exercise essay experiment <br> article adventure |
| RoBERTa base | and the, to |

Table 1: Top 5 MASK candidates when 3 UNKs are put between each token in "This is an MASK."

## E. 1 Experiment

We additionally train RoBERTa with the following configurations.

Dataset We used 8645 samples consisting of more than 128 tokens from wikitext-2.

Putting UNK token Three UNK tokens were inserted after each token in all sentences; the token sequence $\left[t_{1}, t_{2}, \ldots, t_{128}\right]$ was expanded to $\left[t_{1}\right.$, unk, unk, unk, $t_{2}$, unk, unk, unk, $\ldots$, $t_{128}$, unk, unk, unk]. ${ }^{3}$

Freezing parameters We freezed all parameters except position embeddings. This encourages the position embeddings to learn to focus attention only on position $t=1,5,9,13, \cdots$ and prevents positionindependent shortcut solutions (e.g., shortening the UNK embedding to ignore it).

Position embedding The position embeddings were initialized with random values following a normal distribution $\mathcal{N}\left(0,0.02^{2}\right)$ before training.

## E. 2 Results

Table 1 shows that, before the additional training, RoBERTa cannot fill MASK with the appropriate tokens (in this case, nouns beginning with a vowel) when UNKs are inserted, but after the additional

[^2]training, the models can predict them. The attention weights during inference are visualized in Figure 22. Since the attentions are concentrated every fourth token, the model with relearned position embeddings recognizes the relative positions of the non-UNK tokens. This result suggests that the attention is not focused on the tokens that cooccur frequently, but on the tokens that are informative to fill the MASK token. Thus, the relative position dependence of attention is simply caused by the fact that related words appear nearby.

In summary, this experiment indicates that the relative position dependence of attention is caused by the combination of two factors: (1) the linguistic property that related words tend to appear nearby due to grammatical rules and collocations, and (2) the property that attention is focused on words that are related in some sense.


[^0]:    ${ }^{1}$ The text taken from abstract and introduction of (Vaswani et al., 2017) was used as input to create the figures in this paper, unless otherwise stated.

[^1]:    ${ }^{2}$ Figure 3a shows that the maximum frequency of the position embeddings is around 60 . The minimum period is hence 512/60 $(>8)$ tokens.

[^2]:    ${ }^{3}$ Since Transformer is symmetric with respect to position, this process is similar to padding. However, the PAD token of RoBERTa published on hugging face is zero-vector, which is inconvenient for analyzing interactions between word embeddings. Thus, we choose UNK tokens.

