

Quantifying Context Mixing in Transformers

Hosein Mohebbi¹ Willem Zuidema² Grzegorz Chrupala¹ Afra Alishahi¹

¹ CSAI, Tilburg University ² ILLC, University of Amsterdam

{h.mohebbi, a.alishahi}@tilburguniversity.edu

w.h.zuidema@uva.nl

grzegorz@chrupala.me

Abstract

Self-attention weights and their transformed variants have been the main source of information for analyzing token-to-token interactions in Transformer-based models. But despite their ease of interpretation, these weights are not faithful to the models' decisions as they are only one part of an encoder, and other components in the encoder layer can have considerable impact on information mixing in the output representations. In this work, by expanding the scope of analysis to the whole encoder block, we propose *Value Zeroing*, a novel context mixing score customized for Transformers that provides us with a deeper understanding of how information is mixed at each encoder layer. We demonstrate the superiority of our context mixing score over other analysis methods through a series of complementary evaluations with different viewpoints based on linguistically informed rationales, probing, and faithfulness analysis.¹

1 Introduction

Transformers (Vaswani et al., 2017), with their impressive empirical success, have become a prime choice of architecture to learn contextualized representations across a wide range of modalities, such as language (Devlin et al., 2019; Brown et al., 2020), vision (Dosovitskiy et al., 2021), vision-language (Radford et al., 2021; Rombach et al., 2022), and speech (Baevski et al., 2020), mainly due to their ability to utilize pairwise interactions between input tokens at every timestep.

To better understand the inner dynamics of Transformers, we need to trace the information flow from the input embeddings up to the output representation (including quantifying the degree of *context mixing*, which we will define below). The attention weights from the multi-head attention mechanisms offer a straightforward starting

point for understanding this flow, and these weights ('raw attention') have been used in many studies (Clark et al., 2019; Kovaleva et al., 2019; Reif et al., 2019; Htut et al., 2019a, *inter alia*). However, the reliability and usefulness of raw attention weights has also been questioned (Jain and Wallace, 2019; Bibal et al., 2022). In particular, attention weights tend to concentrate on uninformative tokens in the context (Voita et al., 2018; Clark et al., 2019), and removing or altering them may lead to the same and sometimes even better model performance on downstream tasks (Jain and Wallace, 2019; Toneva and Wehbe, 2019; Hassid et al., 2022). These findings suggest that Transformers do not solely rely on self-attention, and other components in the encoder block play an essential role in information mixing.

A number of methods have been proposed to compute some form of 'effective attention weights', with the goal of more faithfully tracing the relative contributions of different input tokens at various layers of the Transformer (as we will discuss in Section 2). These methods show an improvement over raw attention, but still ignore key components of the Transformer encoder block. This is a particularly crucial shortcoming, given that most of the parameter budget in a Transformer encoder is spent on position-wise feed-forward networks outside of the self-attention component, which can have a considerable impact on the degree of information mixing in the output representations.

In this paper, we focus on *context mixing*: the property of Transformers that in each node, at each layer, information from the context can be incorporated into the representation of the target token. We propose *Value Zeroing*, a novel approach to quantify the contribution each context token has in determining the final representation of a target token, at each layer of a Transformer. *Value Zeroing* is based on the Explaining-by-Removing intuition (Covert et al., 2021) shared by many posthoc interpretability methods, but it takes advantage of

¹Code is freely available at <https://github.com/hmohebbi/ValueZeroing>

a specific feature of Transformers: it zeroes only the value vector of a token t when computing its importance, but leaves the key and query vectors (and thus the pattern of information flow) intact.

Based on extensive experiments and three complementary approaches to evaluation, we demonstrate that importance scores we can obtain with *Value Zeroing* provide better interpretations than other analysis methods.

Firstly, we use a set of grammatical agreement tasks from the BLiMP corpus (Warstadt et al., 2020) as a case study. Transformer-based models do extremely well on the task of distinguishing grammatical from ungrammatical sentences, and the BLiMP corpus provides information on the *cue* words that determine the difference. We find that *Value Zeroing*, unlike earlier approaches, indeed reveals that Transformers make use of relevant cues.

Secondly, we use information-theoretic probing (Voita and Titov, 2020; Pimentel et al., 2020) as an independent approach to track information flow in Transformer networks. The scores we obtain with *Value Zeroing* turn out to be highly correlated with layer-wise probing performance; that is, probing accuracy is higher in layers where relevant tokens are more effectively utilized by the model.

Thirdly, we assess the faithfulness of our method (Jacovi and Goldberg, 2020); compared to alternative analysis methods, we show that *Value Zeroing* is not only more plausible and human-interpretable, but also more faithful to models' decisions.

2 Related Work

While numerous studies have leveraged the weights assigned by the self-attention mechanism to gain intuition about the information mixing process in Transformers (Clark et al., 2019; Kovaleva et al., 2019; Reif et al., 2019; Lin et al., 2019; Htut et al., 2019b; Raganato and Tiedemann, 2018), it is still a matter of debate whether attention weights are suitable for interpreting the model (see Bibal et al. (2022)'s study for a full discussion). Thus several post-processing interpretability techniques have been proposed to convert these weights into scores that provide a more detailed interpretation of the inner workings of Transformers. We review the main approaches below.

Abnar and Zuidema (2020) propose the *attention-flow* and *attention-rollout* methods to approximate information flow in Transformers based on raw attention weights. The former treats raw

attention weight matrices as a flow network and returns the maximum flow through each input token. The latter recursively multiplies the attention weight matrix at each layer by the preceding ones. There is, however, an unjustified assumption in the formulation of these methods that both multi-head attention and residual connections contribute equally to the computation of the output.

Kobayashi et al. (2020) propose a method that incorporates the norm of the transformed value vectors and report a negative correlation between these norms and raw attention weights on frequent tokens, which partially explains the insufficiency of raw attention weights for context mixing estimation. Kobayashi et al. (2021) extend this method to the whole self-attention block by incorporating Residual connections (RES) and Layer Normalization (LN) (two components with significant impact on both model performance and training convergence (Parisotto et al., 2020; Liu et al., 2020)), but demonstrate that RES and LN components largely cancel out the mixing process. Kobayashi et al. (2021)'s method, however, ignores the effect of the second sublayer in a Transformer's encoder.

Brunner et al. (2020) and Pascual et al. (2021) employ a gradient-based approach for analyzing the interaction of input representations, but the gradient measures the sensitivity between two vectors and ignores the impact of the input vector. In our experiments we show that despite their relative success in explaining model decisions, gradient-based approaches are not suitable for layer-wise analysis.

More recently, the effectiveness of combining these approaches has also been investigated. Modarressi et al. (2022) propose a method that uses Kobayashi et al. (2021)'s scores and incorporates the effects of the second layer normalization; they aggregate those scores using *rollout* (Abnar and Zuidema, 2020) to provide global token attributions. In the same vein, Ferrando et al. (2022) use rollout to aggregate a variant of Kobayashi et al. (2021)'s scores: instead of relying on the Euclidean norm of the transformed vector, they measure Manhattan distance of each transformed vector to the context vector outputted from self-attention block. In both studies, however, the fact that these context vectors might undergo significant changes after passing through the second sublayer in the encoder layer is not taken into account. We will show (in Section 7) that even at a global level, our scores provide better interpretation than prior methods.

3 Our Proposed Method

To remedy for the limited scope of the existing methods, we introduce a new context mixing score that takes into account all components in a Transformer encoder block.

3.1 Background and Notation

In this section, we set up the notation and briefly review the internal structure of an encoder layer in the Transformer architecture.

Each Transformer encoder layer is composed of two sublayers: a multi-head self-attention mechanism (**MHA**) and a position-wise fully connected feed-forward network (**FFN**), followed by a Residual connection (**RES**) and Layer Normalization (**LN**) around each of these two sublayers. This encoder layer produces the next contextualized representations $(\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n)$ for each token in the context, using the output representations from the previous layer $(\mathbf{x}_1, \dots, \mathbf{x}_n)$.

MHA. For each head $h \in \{1, \dots, H\}$ in the self-attention module, each input vector \mathbf{x}_i is transformed into a query \mathbf{q}_i^h , a key \mathbf{k}_i^h , and a value \mathbf{v}_i^h vector via separate trainable linear transformations:

$$\mathbf{q}_i^h = \mathbf{x}_i \mathbf{W}_Q^h + \mathbf{b}_Q^h \quad (1)$$

$$\mathbf{k}_i^h = \mathbf{x}_i \mathbf{W}_K^h + \mathbf{b}_K^h \quad (2)$$

$$\mathbf{v}_i^h = \mathbf{x}_i \mathbf{W}_V^h + \mathbf{b}_V^h \quad (3)$$

The context vector \mathbf{z}_i^h for the i^{th} token of each attention head is then generated as a weighted sum over the transformed value vectors:

$$\mathbf{z}_i^h = \sum_{j=1}^n \alpha_{i,j}^h \mathbf{v}_j^h \quad (4)$$

where $\alpha_{i,j}^h$ is the raw attention weight assigned to the j^{th} token, and computed as a softmax-normalized dot product between the corresponding query and key vectors:

$$\alpha_{i,j} = \operatorname{softmax}_{\mathbf{x}_j \in \mathcal{X}} \left(\frac{\mathbf{q}_i \mathbf{k}_j^\top}{\sqrt{d}} \right) \in \mathbb{R} \quad (5)$$

Next, the context vector $(\mathbf{z}_i \in \mathbb{R}^d)$ for the i^{th} token is computed by concatenating all the heads' outputs followed by a head-mixing \mathbf{W}_O projection and layer normalization:

$$\mathbf{z}_i = \operatorname{CONCAT}(\mathbf{z}_i^1, \dots, \mathbf{z}_i^H) \mathbf{W}_O \quad (6)$$

$$\mathbf{z}_i = \operatorname{LN}_{\text{MHA}}(\mathbf{z}_i + \mathbf{x}_i) \quad (7)$$

FFN. Each encoder layer also includes two linear transformations with a ReLU activation in between, which is applied to every \mathbf{z}_i separately and identically to produce output token representations $\tilde{\mathbf{x}}_i$:

$$\tilde{\mathbf{x}}_i = \max(0, \mathbf{z}_i \mathbf{W}_1 + \mathbf{b}_1) \mathbf{W}_2 + \mathbf{b}_2 \quad (8)$$

$$\tilde{\mathbf{x}}_i = \operatorname{LN}_{\text{FFN}}(\tilde{\mathbf{x}}_i + \mathbf{z}_i) \quad (9)$$

3.2 Value Zeroing

We aim to measure how much a token uses other context tokens to build its output representation $\tilde{\mathbf{x}}_i$ at each encoder layer. To this end, we treat the self-attention mechanism as a fuzzy hash-table, where we look up the sum of values weighted by the query-key match in the context. Thus in Eq. 4 we replace a value vector associated with token j with a zero vector $\mathbf{v}_j^h \leftarrow \mathbf{0}, \forall h \in H$, where the context vector for the i^{th} token is being computed. This provides an alternative output representation $\tilde{\mathbf{x}}_i^{-j}$ for the i^{th} token that has excluded token j in the mixing process. By comparing the alternative output representation $\tilde{\mathbf{x}}_i^{-j}$ with the original $\tilde{\mathbf{x}}_i$, we can measure how much the output representation is affected by the exclusion of the j^{th} token:

$$\mathcal{C}_{i,j} = \tilde{\mathbf{x}}_i^{-j} * \tilde{\mathbf{x}}_i \quad (10)$$

where the operation $*$ can be any pairwise distance metric that properly considers the characteristics of the model's representation space. We opted for *cosine* distance throughout our experiments as its superiority over other dissimilarity metrics has been supported for textual deep learning models (Yokoi et al., 2020; Hanawa et al., 2021).² Computing Eq. 10 for all $\tilde{\mathbf{x}}_i$ in a given context provides us with a *Value Zeroing* matrix score \mathcal{C} where the value of the cell $\mathcal{C}_{i,j} \in \mathbb{R}$ (i^{th} row, j^{th} column) in the map denotes the degree to which the i^{th} token depends on the j^{th} token to form its contextualized representation.

Note that unlike generic perturbation approaches, our proposed method *does not* remove the token representations \mathbf{x}_i from the input of an encoder. We argue that ablating input token representations cannot be a reliable basis to understand context mixing process since any changes in the input vectors will lead to changes in the query and key vectors (Eq. 1 and 2), resulting in a change in the attention distribution (cf. Eq. 5). Consequently, there will

²More details on the choice of distance metric is discussed in Appendix A.1.

Phenomenon	UID	Example	Target word	Foil word
Anaphor Number Agreement	ana	<u>Many teenagers</u> were helping [MASK].	themselves	herself
Determiner-Noun Agreement	dna	Jeffrey has not passed [MASK] <u>museums</u> .	these	this
	dnaa	Sara noticed [MASK] <u>white hospitals</u> .	these	this
Subject-Verb Agreement	darn	The <u>pictures</u> of Martha [MASK] not disgust Anne.	do	does
	rpsv	<u>Kristen</u> [MASK] fixed this chair.	has	have

Table 1: Examples of the selected tasks with our annotations from the BLiMP benchmark (UIDs are unique identifiers used in BLiMP). Cue words are underlined.

be a discrepancy between the alternative attention weights and those for the original context. Instead, our method only nullifies the *value vector* of a specific token representation. In this way, the token representation can maintain its identity within the encoder layer, but it does not contribute to forming other token representations. Moreover, since our *Value Zeroing* is computed from the encoder’s layer outputs, it incorporates all the components inside an encoder layer such as multi-head attention, residual connection, layer normalization, and also feed-forward networks, resulting in a more reliable context mixing score than previous methods.

4 Experimental Setup

4.1 Data

We used the BLiMP benchmark (Warstadt et al., 2020) which contains a set of pairs of minimally different sentences that contrast in grammatical acceptability under a specific linguistic phenomenon. The benchmark isolates linguistic phenomena such that only one word determines the true label of each sentence. We refer to this crucial context token as the *cue word*. The nature of this task makes it especially suitable for evaluating context mixing scores, since it gives us a strong hypothesis on which context token is the most relevant for the representation of the masked *target word*.

From this benchmark, we select five datasets with three different linguistic phenomena for which Pre-trained Language Models (PLMs) have shown high accuracy to ensure that the model captures the relevant information. We expand contractions such as *doesn’t* → *does not*) and generate dependency trees using SpaCy (Honnibal and Montani, 2017) to extract and annotate the position of target and cue words in a sentence. In Table 1, we provide an example of each phenomenon in the benchmark together with our automated annotations. We accumulate examples from the five selected tasks as a unified dataset for grammatical agreement, result-

ing in 4,276 data points, and divide them equally into Train and Test sets. The Train set is only used for the fine-tuning phase; the Test set is used for all evaluation experiments.

4.2 Target Model

We conduct our experiments on three Transformer-based language models: BERT (uncased, Devlin et al., 2019), RoBERTa (Liu et al., 2019) and ELECTRA (Clark et al., 2020).³ The results for the latter two are reported in Appendix A.3. By replacing the target words with the [MASK] token, we perform a Masked Language Modeling (MLM) task using the model’s pre-trained MLM head. For instance, in the Subject-Verb Agreement example “*The pictures of Martha do not disgust Anne.*”, we replace the verb ‘do’ with the [MASK] token and feed the example to the model.

We perform our experiments on both pre-trained and fine-tuned versions of each model. Including a fine-tuned model in our analysis study gives us a complementary insight into the importance of the cue words, since fine-tuning allows the model to concentrate on the most helpful words for the downstream task of choice (i.e., agreement) and makes sure that target word representations take the cue word into account. We use prompt fine-tuning (Schick and Schütze, 2021a,b; Karimi Mahabadi et al., 2022) and compute Cross Entropy loss only over the output logits corresponding to the target and foil classes. Accuracy is 0.96 for pre-trained and 0.99 for fine-tuned BERT.

4.3 Baselines

Here we describe the existing context mixing methods which we include in our experiments. For each method, we select the m^{th} row of its context mixing map where m is the position of the [MASK] token, resulting in a 1-D array of scores for each context

³Base, with 12 layers and 12 attention heads, obtained from the Transformers library (Wolf et al., 2020).

token. We normalize the scores for all tokens in a sentence so that they are all positive values and sum to one. For completeness, we also include a few gradient-based attribution methods in our comparisons.

Rand: random scores generated from a uniform distribution for each sentence in the dataset.

Attn: raw attention scores α_m from Eq. 5.

Attn-rollout: An aggregation method for approximating the attention flow based on raw attention weights (Abnar and Zuidema, 2020).

Attn-Norm: norm-based method of Kobayashi et al. (2020) that also incorporates the norm of the transformed input vectors to compute context mixing scores.

Attn-Norm + RES + LN: the extended norm-based method of Kobayashi et al. (2021) in which they also incorporate Residual connection (RES) and Layer Normalization (LN) located only in the first sublayer of a Transformer’s encoder.

GlobEnc & ALTI: Two global token attribution methods proposed by Modarressi et al. (2022) and Ferrando et al. (2022). For a fair comparison, we exclude the aggregation through rollout in our first two evaluations (referred to as GlobEnc \rightarrow and ALTI \rightarrow), because we are interested in assessing context mixing scores at each layer separately. In our third evaluation, we compare these methods with others at a global level and use rollout.

GradXinput, IG & DL: feature attribution scores that also make use of *top-down* information from the classification layer on top of the Transformer. We consider three popular variants of gradient-based attribution scores: Gradient \times Input (**GradXinput**) (Samek et al., 2019; Yuan et al., 2019), Integrated Gradients (**IG**) (Sundararajan et al., 2017), and DeepLift (**DL**) (Shrikumar et al., 2017). For gradient-based methods, we use the pre-trained MLM head which has been trained during the pre-training of the BERT to compute the gradient of the true label with respect to the token representations at each layer.

5 Evaluation 1: Cue Alignment

As cue words are the only indicators of the true labels in our dataset, we expect that when the model performs well, it overwhelmingly depends on these words to form the representation of a [MASK] to-

ken in a given context. To quantify the alignment between a context mixing score and the cue word, we first define a binary cue vector ξ according to the following condition:

$$\xi_i = \begin{cases} 1, & \text{the } i^{\text{th}} \text{ token} \in \text{Cue words} \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

Then we compare the cue vector and the prediction of a context mixing score S in two different ways:

Dot Product. We quantify cue alignment as $S \cdot \xi$, which measures the total score mass the model assigns to cue words to form the representation of the target token.

Average Precision. We quantify cue alignment as the average precision between the two vectors, which is a weighted mean of precision at each recall level:

$$AP = \sum_n (R_n - R_{n-1}) P_n \quad (12)$$

where P_n and R_n are the precision and recall at the n^{th} threshold. This metric relies on the ranking of tokens rather than the magnitude of their weights.⁴

Figures 1 and 2 show the alignment between the cue vector and different analysis methods using dot product and average precision for the pre-trained and fine-tuned model, respectively. In comparison with the other context mixing methods, *Value Zeroing* shows a higher degree of the target model incorporating cue words into the representation of the [MASK] token across all layers.

As can be seen from the first two columns in all graphs, raw self-attention weights (**Attn**) always perform worse than even random scores in highlighting cue words. This is in line with previous studies showing that raw attention weights often pay attention to uninformative tokens (Voita et al., 2018; Clark et al., 2019) and do not reflect the appropriate context (Kim et al., 2019). However, we can see a significant improvement in the results for **Attn-norm** where the norm of transformed value vectors are also taken into account, confirming that value vectors play an essential role in the context mixing process. The method **Attn-norm + RES + LN**, which expands **Attn-norm** to the whole self-attention block by adding Residual

⁴We also employed Probes-needed (Zhong et al., 2019) metric in our evaluation which intuitively counts the number of non-cue tokens we need to probe to find cue words based on a given score. As its motivation is similar to Average Precision and the results show the same pattern, we relegate the results with this metric to the Appendix A.2.




Figure 1: Layer-wise alignment between the cue vector and different analysis methods averaged over Test set examples for the **pre-trained** model. Higher value (darker color) is better.




Figure 2: Layer-wise alignment between the cue vector and different analysis methods averaged over Test set examples for the **fine-tuned** model. Higher value (darker color) is better.

connection and Layer Normalization, would seem to show that the model is incapable of utilizing the cue words. However, incorporating also the second part of the encoder layer via our method shows the model does indeed use the cue words. We also find that GlobEnc and ALTI benefit from the rollout aggregation method to provide a global view, but they do not seem to provide good layerwise scores (without rollout).

The gradient-based scores, in contrast to the other methods, highlight the cue words only in the earlier layers of the model. In the next section, by using a layer-wise probing experiment, we will show that these scores are not reliable for identifying the relevant context in individual layers.

6 Evaluation 2: Context Mixing versus Probing

In this section, we investigate the relationship between cue word alignment and probing performance across layers. We hypothesize that if a layer aligns better with the cue word according to a reliable context mixing score, then the representation of the masked token on that layer can be used more effectively by a probing classifier to decode number agreement with the cue word.

To verify our hypothesis, we obtain the representation of masked tokens in test examples across all layers. Since all examples in our dataset share the same number agreement property, we asso-




Figure 3: Layer-wise compression of probing classifiers using pre-trained and fine-tuned representations.

ciate each masked representation with a *Singular* or *Plural* label. Next, we perform an information-theoretic probing analysis using Minimum Description Length (MDL) to measure the degree to which representations encode number agreement. We chose MDL as our probe since it is theoretically justified and has been shown to provide more reliable results than conventional probes (Voita and Titov, 2020; Fayyaz et al., 2021).

To compute MDL, we employed the *online coding* of Voita and Titov (2020). Since MDL can be affected by the number of data points (N), we measure compression as our evaluation metric which is defined as follows:

$$\text{Compression} = \frac{N \cdot \log_2(K)}{\text{MDL}} \quad (13)$$

where K refers to the number of classes (2 in our case). This metric is equal to 1 (no compression) for a random guessing classifier. A higher value for Compression indicates more accurate label prediction for the probing classifier.

Figure 3 reports compression of probing classifiers based on representations obtained from both pre-trained and fine-tuned models across all layers. We also include the results for the embedding layer of the model (layer 0) which can serve as a non-contextualized baseline. We can see a jump in probing performance at layers 4 (1.52 \rightarrow 1.72) and 9 (2.03 \rightarrow 2.45) in the fine-tuned setup, the same layers for which we found a higher alignment with cue words in Figures 1 and 2.

Table 2 presents the correlation between layer-wise Compression scores and layer-wise cue alignment scores from Section 5 for different analysis methods. As we can see, alignment according to *Value Zeroing* is highly positively correlated with the probing performance. This suggests that when *Value Zeroing* indicates that the model uses cue words to form representations of the masked tokens in a particular layer, these representations are in fact better at encoding number agreement.

Method	ρ^{PT}	ρ^{FT}
Rand	-0.07	-0.02
Attn	0.10	0.08
Attn-norm	0.52	0.56
Attn-norm + RES	-0.35	-0.24
Attn-norm + RES + LN	0.12	0.17
GlobEnc \neg	-0.01	-0.12
ALTI \neg	-0.01	-0.12
GradXinput	-0.96	-0.99
IG	-0.86	-0.77
DL	-0.97	-1.00
<i>Value Zeroing</i>	0.65	0.64

Table 2: Spearman’s ρ correlation between layer-wise probing performance (Comp.) and layer-wise cue alignments based on Dot Product. PT and FT refer to pre-trained and fine-tuned conditions, respectively.

Recall that based on Figures 1 and 2, according to the gradient-based methods the masked tokens pay more attention to the cue words only in earlier layers. However, we can see a highly negative correlation with probing results for these scores. Due to the nature of the task the probing score goes up monotonically along the layers. At the same time, the gradient attribution score goes up monotonically as you get closer to the bottom embedding layers, suggesting that gradient-based methods are unreliable for layer-wise analysis and identifying important tokens in the context mixing process.

7 Evaluation 3: Faithfulness Analysis

Our experimental results in Sections 5 and 6 show that the *Value Zeroing* score matches our prior linguistically-informed expectations. However, it is not always clear whether a *plausible* context mixing score that matches human expectations is also faithful to the model and reflects its decision making process (Herman, 2017; Wiegrefe and Pinter, 2019; Jacovi and Goldberg, 2020).

In this section we employ the notion of input ablation (Covert et al., 2021) to evaluate the faithfulness of our context mixing score. The influence of a target token on a model’s decision is often estimated as the drop in the model’s predicted probability of the correct class after blanking out the target token from the input. A higher drop for an ablated token indicates that the token is more influential on the model’s decision (DeYoung et al., 2020; Abnar and Zuidema, 2020; Atanasova et al., 2020; Wang et al., 2022). We use this *blank-out* approach as a base for analyzing and comparing the faithfulness of context mixing scores.

Method	ρ^{PT}	ρ^{FT}
Rand	-0.01	0.00
Attn	-0.10	-0.07
Attn-norm	0.19	0.14
Attn-norm + RES	0.03	-0.05
Attn-norm + RES + LN	-0.08	-0.17
GlobEnc	-0.01	-0.09
ALTI	0.17	0.19
GradXinput	0.11	0.16
IG	0.07	0.21
DL	0.20	0.29
<i>Value Zeroing</i>	0.26	0.31

Table 3: Spearman’s ρ correlation between the blank-out scores and different aggregated context mixing and attribution scores. PT and FT refer to pre-trained and fine-tuned conditions, respectively.

To estimate the blank-out scores in BERT, we calculate the probability of its output y using a softmax function normalized over only the corresponding logit values of target t and foil words (cf. Table 1), and compute blank-out scores for a given input token i as $p(y_t|e) - p(y_t|e \setminus e_i)$, where e_i refers to the input embedding of input token i . We compare these blank-out scores with context mixing scores, aggregated across all layers of the model. For gradient-based scores, calculating them with respect to the tokens in the input embedding layer ($\ell = 0$) provides us with aggregated scores since the backpropagation of gradients passes through all layers to the beginning of the model. For other scores, we use the *rollout* (Abnar and Zuidema, 2020) aggregation method.

Table 3 shows Spearman’s rank correlation between the blank-out scores and different aggregated context mixing scores. The highest correlation for our method indicates that *Value Zeroing* is more faithful in explaining the model behaviour compared to other analysis methods.

Qualitative Analysis. We also take a closer look at the aggregated scores for a qualitative comparison. In Figure 4, we illustrate different scores obtained from a fine-tuned BERT model for a correctly classified example, where the model is asked to fill the masked token with one of the verbs *were* or *is* as target and foil classes, respectively. According to *Value Zeroing* scores, the model mainly relies on the main subject (*pictures*) as a cue word to form a contextualized representation of the [MASK] token, while the word *pictures* is also important for the model’s final decision based on



Figure 4: Most influential tokens on the target representation in a fine-tuned BERT model according to different aggregated context mixing scores compared to blank-out scores.

the blank-out scores. In this example, the blank-out score for the cue word is 0.99, meaning the model fully loses its confidence in the target class when the cue word is replaced with an [UNK] token. Surprisingly, gradient-based methods tend to highlight the word *hat* which is an agreement attractor, and attention-based scores tend to focus on the [CLS] token which has been idle during fine-tuning process.

Overall, our faithfulness evaluation and qualitative analysis suggest that *Value Zeroing* can explain model decisions at a global level when it is aggregated across layers. The context mixing maps per layer are provided in Appendix A.4, where some more meaningful patterns can be found in Value Zeroing scores (in both layer-wise and aggregated setups) in contrast to other context mixing scores.

8 Discussion

Although some desiderata such as plausibility (Lei et al., 2016; Strout et al., 2019) and faithfulness (Lakkaraju et al., 2019; Jacovi and Goldberg, 2020) are taken into account when developing explanation and analysis methods, evaluating them is still a challenge due to lack of a standard ground truth. Evaluating context mixing scores, where token-to-token interactions in a context are also considered, is even more challenging. Several studies have used gradient-based scores as an anchor of faithfulness, and measure how strongly context mixing scores correlate with them (Jain and Wallace, 2019; Abnar and Zuidema, 2020; Modarressi et al., 2022). However, the reliability of gradient-based scores can be questioned, especially when different variations of them show considerable disagreement

(Neely et al., 2022; Pruthi et al., 2022; Krishna et al., 2022). Thus, we suggest using controlled tasks for which we have strong prior expectations for evaluating these methods. In our study, we use a set of number agreement tasks to provide such priors, since the cue words are the only sources of information in the context for performing well in the task.

Another point worth discussing is the concern raised by Kobayashi et al. (2021) that BERT tends to preserve token representations rather than mixing them at each layer. We argue that their observation is due to the context-mixing ratio they defined by comparing the norm of residual effects against other token representations. In our view, this ratio is dominated by residuals and neglects the fact that a token representation carried by residual connections is indeed a contextualized representation outputted from previous layers. We keep the residuals intact within the encoder layer by zeroing only the value vectors and focusing on the context mixing performed by all tokens.

9 Conclusion

In this paper, we propose Value Zeroing as a novel approach for quantifying the information mixing process in Transformers to address the shortcomings of previous methods. We performed extensive complementary experiments and showed that our method outperforms others in three different evaluation setups. Since our approach requires no supervision, it could be an interesting option for improving model efficiency by removing token representations across layers.

10 Limitations

As is the case for most attempts at interpreting Deep Learning models, our evaluation of our (and others’) proposed methods are not definite since we have no gold standard of what happens inside a model, although we try to remedy for that by conducting independent and complementary evaluation schemes.

Our proposed method is customized for deep neural models based on the Transformer architecture and cannot be easily generalized to other (mathematically different) modeling architectures. Our evaluations were based on encoder-based models, and focused on the Text modality. In the future, we will extend our experiments to more modalities, such as speech and vision.

Acknowledgments

This publication is part of the project *InDeep: Interpreting Deep Learning Models for Text and Sound* (with project number NWA.1292.19.399) of National Research Agenda (NWA-ORC) programme. Funding by the Dutch Research Council (NWO) is gratefully acknowledged.

References

- Samira Abnar and Willem Zuidema. 2020. [Quantifying attention flow in transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4197, Online. Association for Computational Linguistics.
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. [A diagnostic study of explainability techniques for text classification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3256–3274, Online. Association for Computational Linguistics.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33:12449–12460.
- Adrien Bibal, Rémi Cardon, David Alfter, Rodrigo Wilkens, Xiaou Wang, Thomas François, and Patrick Watrin. 2022. [Is attention explanation? an introduction to the debate](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3889–3900, Dublin, Ireland. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Gino Brunner, Yang Liu, Damian Pascual, Oliver Richter, Massimiliano Ciaramita, and Roger Wattenhofer. 2020. [On identifiability in transformers](#). In *International Conference on Learning Representations*.
- 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.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. **Electra: Pre-training text encoders as discriminators rather than generators**. In *International Conference on Learning Representations*.
- Ian Covert, Scott M. Lundberg, and Su-In Lee. 2021. Explaining by removing: A unified framework for model explanation. *J. Mach. Learn. Res.*, 22:209:1–209:90.
- 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.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020. **ERASER: A benchmark to evaluate rationalized NLP models**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4443–4458, Online. Association for Computational Linguistics.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. **An image is worth 16x16 words: Transformers for image recognition at scale**. In *International Conference on Learning Representations*.
- Mohsen Fayyaz, Ehsan Aghazadeh, Ali Modarressi, Hossein Mohebbi, and Mohammad Taher Pilehvar. 2021. **Not all models localize linguistic knowledge in the same place: A layer-wise probing on BERToids’ representations**. In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 375–388, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Javier Ferrando, Gerard I. Gállego, and Marta R. Costa-jussà. 2022. **Measuring the mixing of contextual information in the transformer**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8698–8714, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Kazuaki Hanawa, Sho Yokoi, Satoshi Hara, and Kentaro Inui. 2021. **Evaluation of similarity-based explanations**. In *International Conference on Learning Representations*.
- Michael Hassid, Hao Peng, Daniel Rotem, Jungo Kasai, Ivan Montero, Noah A. Smith, and Roy Schwartz. 2022. **How much does attention actually attend? questioning the importance of attention in pretrained transformers**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1403–1416, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Bernease Herman. 2017. The promise and peril of human evaluation for model interpretability. *ArXiv*, abs/1711.07414.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Phu Mon Htut, Jason Phang, Shikha Bordia, and Samuel R. Bowman. 2019a. **Do attention heads in BERT track syntactic dependencies?** *CoRR*, abs/1911.12246.
- Phu Mon Htut, Jason Phang, Shikha Bordia, and Samuel R. Bowman. 2019b. **Do attention heads in bert track syntactic dependencies?** *ArXiv*, abs/1911.12246.
- Alon Jacovi and Yoav Goldberg. 2020. **Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness?** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4198–4205, Online. Association for Computational Linguistics.
- Sarthak Jain and Byron C. Wallace. 2019. **Attention is not Explanation**. 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 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rabeeh Karimi Mahabadi, Luke Zettlemoyer, James Henderson, Lambert Mathias, Marzieh Saeidi, Veselin Stoyanov, and Majid Yazdani. 2022. **Prompt-free and efficient few-shot learning with language models**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3638–3652, Dublin, Ireland. Association for Computational Linguistics.
- Yunsu Kim, Duc Thanh Tran, and Hermann Ney. 2019. **When and why is document-level context useful in neural machine translation?** In *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*, pages 24–34, Hong Kong, China. Association for Computational Linguistics.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. **Attention is not only a weight: Analyzing transformers with vector norms**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7057–7075, Online. Association for Computational Linguistics.

- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2021. [Incorporating Residual and Normalization Layers into Analysis of Masked Language Models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4547–4568, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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.
- Satyapriya Krishna, Tessa Han, Alex Gu, Javin Pombra, Shahin Jabbari, Steven Wu, and Himabindu Lakkaraju. 2022. The disagreement problem in explainable machine learning: A practitioner’s perspective. *ArXiv*, abs/2202.01602.
- Himabindu Lakkaraju, Ece Kamar, Rich Caruana, and Jure Leskovec. 2019. [Faithful and customizable explanations of black box models](#). In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’19, page 131–138, New York, NY, USA. Association for Computing Machinery.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. [Rationalizing neural predictions](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 107–117, Austin, Texas. 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.
- Fenglin Liu, Xuancheng Ren, Zhiyuan Zhang, Xu Sun, and Yuexian Zou. 2020. [Rethinking skip connection with layer normalization](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3586–3598, Barcelona, Spain (Online). International Committee on 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*, abs/1907.11692.
- Ali Modarressi, Mohsen Fayyaz, Yadollah Yaghoobzadeh, and Mohammad Taher Pilehvar. 2022. [GlobEnc: Quantifying global token attribution by incorporating the whole encoder layer in transformers](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 258–271, Seattle, United States. Association for Computational Linguistics.
- Michael Neely, Stefan F. Schouten, Maurits Bleeker, and Ana Lucic. 2022. A song of (dis)agreement: Evaluating the evaluation of explainable artificial intelligence in natural language processing. In *HHAI*.
- Emilio Parisotto, Francis Song, Jack Rae, Razvan Pascanu, Caglar Gulcehre, Siddhant Jayakumar, Max Jaderberg, Raphaël Lopez Kaufman, Aidan Clark, Seb Noury, Matthew Botvinick, Nicolas Heess, and Raia Hadsell. 2020. [Stabilizing transformers for reinforcement learning](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7487–7498. PMLR.
- Damian Pascual, Gino Brunner, and Roger Wattenhofer. 2021. [Telling BERT’s full story: from local attention to global aggregation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 105–124, Online. Association for Computational Linguistics.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4609–4622, Online. Association for Computational Linguistics.
- Danish Pruthi, Bhuwan Dhingra, Livio Baldini Soares, Michael Collins, Zachary Chase Lipton, Graham Neubig, and William W. Cohen. 2022. Evaluating explanations: How much do explanations from the teacher aid students? *Transactions of the Association for Computational Linguistics*, 10:359–375.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Alessandro Raganato and Jörg Tiedemann. 2018. [An analysis of encoder representations in transformer-based machine translation](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 287–297, Brussels, Belgium. Association for Computational Linguistics.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B Viegas, Andy Coenen, Adam Pearce, and Been Kim. 2019. Visualizing and measuring the geometry of bert. In *Advances in Neural Information Processing Systems*, pages 8594–8603.
- Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-

- resolution image synthesis with latent diffusion models. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685.
- Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller. 2019. *Explainable AI: interpreting, explaining and visualizing deep learning*, volume 11700. Springer Nature.
- Timo Schick and Hinrich Schütze. 2021a. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021b. [It’s not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. Association for Computational Linguistics.
- Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. 2017. Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, page 3145–3153. JMLR.org.
- Julia Strout, Ye Zhang, and Raymond Mooney. 2019. [Do human rationales improve machine explanations?](#) In *Proceedings of the 2019 ACL Workshop Black-boxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 56–62, Florence, Italy. Association for Computational Linguistics.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, page 3319–3328. JMLR.org.
- William Timkey and Marten van Schijndel. 2021. [All bark and no bite: Rogue dimensions in transformer language models obscure representational quality](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4527–4546, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mariya Toneva and Leila Wehbe. 2019. [Interpreting and improving natural-language processing \(in machines\) with natural language-processing \(in the brain\)](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- 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*, volume 30. Curran Associates, Inc.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. [Context-aware neural machine translation learns anaphora resolution](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics.
- Elena Voita and Ivan Titov. 2020. [Information-theoretic probing with minimum description length](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 183–196, Online. Association for Computational Linguistics.
- Lijie Wang, Yaozong Shen, Shu ping Peng, Shuai Zhang, Xinyan Xiao, Hao Liu, Hongxuan Tang, Ying Chen, Hua Wu, and Haifeng Wang. 2022. A fine-grained interpretability evaluation benchmark for neural nlp. *ArXiv*, abs/2205.11097.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: A benchmark of linguistic minimal pairs for English](#). In *Proceedings of the Society for Computation in Linguistics 2020*, pages 409–410, New York, New York. Association for Computational Linguistics.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation](#). 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 11–20, Hong Kong, China. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Sho Yokoi, Ryo Takahashi, Reina Akama, Jun Suzuki, and Kentaro Inui. 2020. [Word rotator’s distance](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2944–2960, Online. Association for Computational Linguistics.
- Hao Yuan, Yongjun Chen, Xia Hu, and Shuiwang Ji. 2019. [Interpreting deep models for text analysis via optimization and regularization methods](#). In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI’19/IAAI’19/EAAI’19*. AAAI Press.

Ruiqi Zhong, Steven Shao, and Kathleen McKeown. 2019. Fine-grained sentiment analysis with faithful attention. *ArXiv*, abs/1908.06870.

A Appendices

A.1 On the choice of distance function

The purpose of this section is to inspect the impact of selecting different distance metrics when computing Value Zeroing in Eq. 10. Timkey and van Schijndel (2021) questioned the informativity of standard representational distance measures such as cosine and Euclidean by observing that only a small subset of rogue dimensions contribute to the anisotropy of a contextualized representation space. They proposed using simple post-processing techniques to correct for such these rough dimensions. We followed their suggestion and normalized the representations before computing distances, but we did not observe any noticeable difference in our scores compared to using non-normalized representations (Fig. A.1). We also repeated our experiment with Spearman’s and Euclidean distance metrics and observed the same pattern in the results (Fig. A.1).

We believe that in anisotropy studies that use clustering methods, the choice of distance metrics is crucial. However, we compute each token’s distance from itself (not from other tokens) and compare them relatively. This might explain why we observe the same pattern for different distance metrics.

A.2 More metrics

Figure A.2 reports the cue alignment evaluation for BERT model based on Probes-needed (Zhong et al., 2019) metric.

A.3 More PLMs

We replicated our experiment for the cue alignment for two more PLMs; RoBERTa (Liu et al., 2019) and ELECTRA (generator, Clark et al., 2020). As we can see in Figures A.3 and A.4, our method consistently outperforms other methods on all models in both pre-trained and fine-tuned setups. Due to the fact that our scores are based on zeroing value vectors, our method can be easily applied to any Transformer-based models even with different modalities.

Layer	Dot Product					Average Precision					
	1	2	3	4	5	1	2	3	4	5	
12	0.21	0.21	0.21	0.21	0.17	0.17	0.44	0.44	0.44	0.44	0.44
11	0.26	0.26	0.26	0.26	0.19	0.19	0.54	0.54	0.54	0.54	0.54
10	0.21	0.21	0.21	0.21	0.18	0.18	0.45	0.45	0.46	0.46	0.45
9	0.28	0.28	0.28	0.28	0.22	0.22	0.57	0.57	0.57	0.57	0.57
8	0.21	0.21	0.21	0.21	0.19	0.19	0.47	0.47	0.47	0.47	0.47
7	0.17	0.17	0.17	0.17	0.17	0.17	0.43	0.42	0.42	0.42	0.43
6	0.17	0.17	0.17	0.17	0.16	0.16	0.41	0.41	0.41	0.41	0.41
5	0.24	0.24	0.23	0.23	0.19	0.19	0.54	0.54	0.54	0.54	0.54
4	0.21	0.21	0.21	0.21	0.18	0.18	0.48	0.48	0.48	0.48	0.48
3	0.19	0.19	0.19	0.19	0.15	0.15	0.37	0.37	0.38	0.37	0.37
2	0.14	0.15	0.14	0.14	0.13	0.13	0.36	0.37	0.36	0.37	0.36
1	0.10	0.10	0.10	0.10	0.10	0.10	0.32	0.33	0.33	0.32	0.33

(a) Pre-trained BERT

Layer	Dot Product					Average Precision					
	1	2	3	4	5	1	2	3	4	5	
12	0.22	0.22	0.22	0.21	0.17	0.17	0.46	0.45	0.45	0.45	0.45
11	0.25	0.25	0.24	0.24	0.19	0.19	0.53	0.53	0.53	0.53	0.53
10	0.21	0.21	0.20	0.21	0.18	0.18	0.45	0.45	0.45	0.45	0.45
9	0.35	0.35	0.34	0.34	0.25	0.25	0.68	0.68	0.68	0.68	0.68
8	0.24	0.24	0.24	0.24	0.21	0.21	0.53	0.53	0.53	0.53	0.53
7	0.19	0.19	0.19	0.19	0.18	0.18	0.46	0.46	0.46	0.46	0.46
6	0.18	0.18	0.17	0.17	0.17	0.17	0.43	0.43	0.43	0.43	0.43
5	0.25	0.25	0.25	0.25	0.20	0.20	0.58	0.58	0.57	0.57	0.58
4	0.21	0.21	0.21	0.21	0.18	0.18	0.48	0.48	0.48	0.47	0.48
3	0.19	0.19	0.19	0.19	0.15	0.15	0.37	0.37	0.37	0.37	0.37
2	0.14	0.15	0.14	0.14	0.13	0.13	0.36	0.37	0.36	0.37	0.36
1	0.10	0.10	0.10	0.10	0.11	0.11	0.33	0.33	0.33	0.33	0.33

(b) Fine-tuned BERT

Figure A.1: Layer-wise alignment between the cue vector and different Value Zeroing (VZ) scores computed based on 1) different distance metric and 2) whether representations are normalized.

A.4 Qualitative Analysis: Layer-wise Context Mixing Maps

This section illustrates different context mixing maps obtained from a fine-tuned BERT model for the correctly classified example of “The pictures of some hat [MASK] scaring Marcus.”






Figure A.2: The layer-wise alignment based on Probes-needed metric between the cue vector and different analysis methods averaged over Test set examples. Lower value (darker blue) is better.




(a) Pre-trained RoBERTa




(b) Fine-tuned RoBERTa

Figure A.3: Layer-wise alignment between the cue vector and different analysis methods averaged over Test set examples for RoBERTa. Higher value (darker color) is better.



(a) Pre-trained ELECTRA



(b) Fine-tuned ELECTRA

Figure A.4: Layer-wise alignment between the cue vector and different analysis methods averaged over Test set examples for ELECTRA. Higher value (darker color) is better.

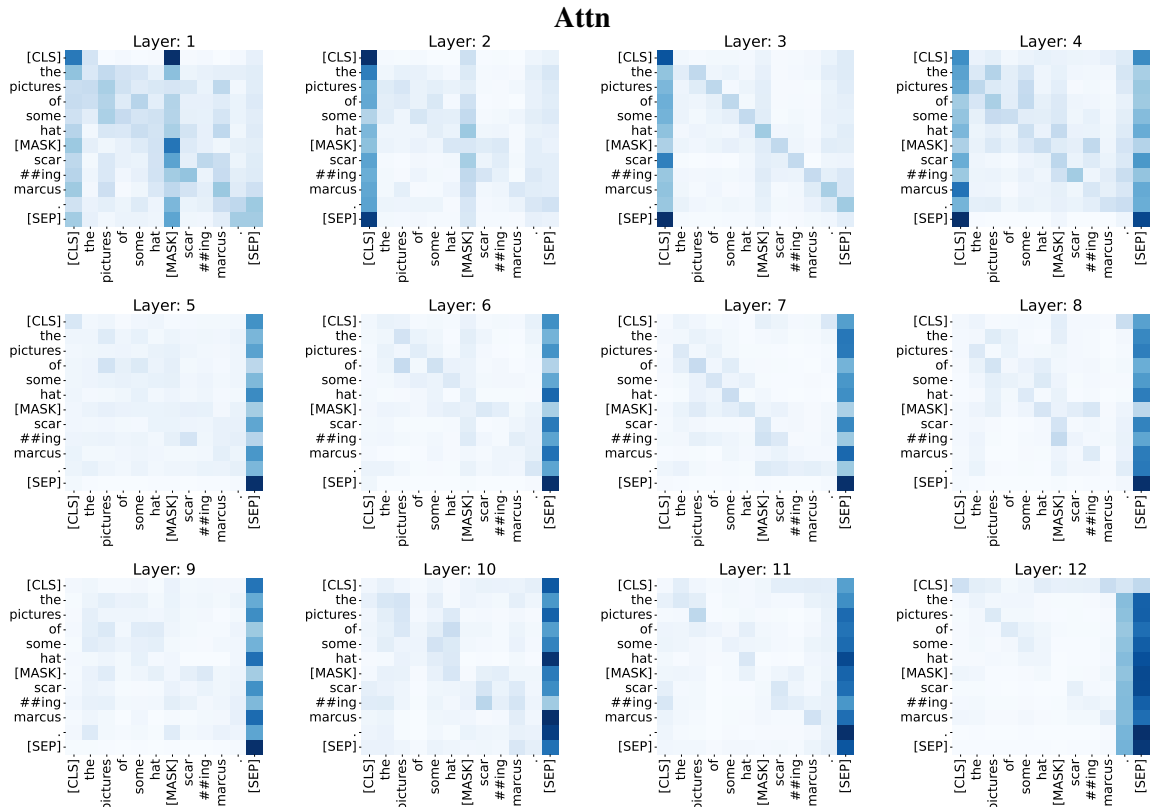


Figure A.5: Raw attention scores (Attn) averaged over all attention heads at each different layer.

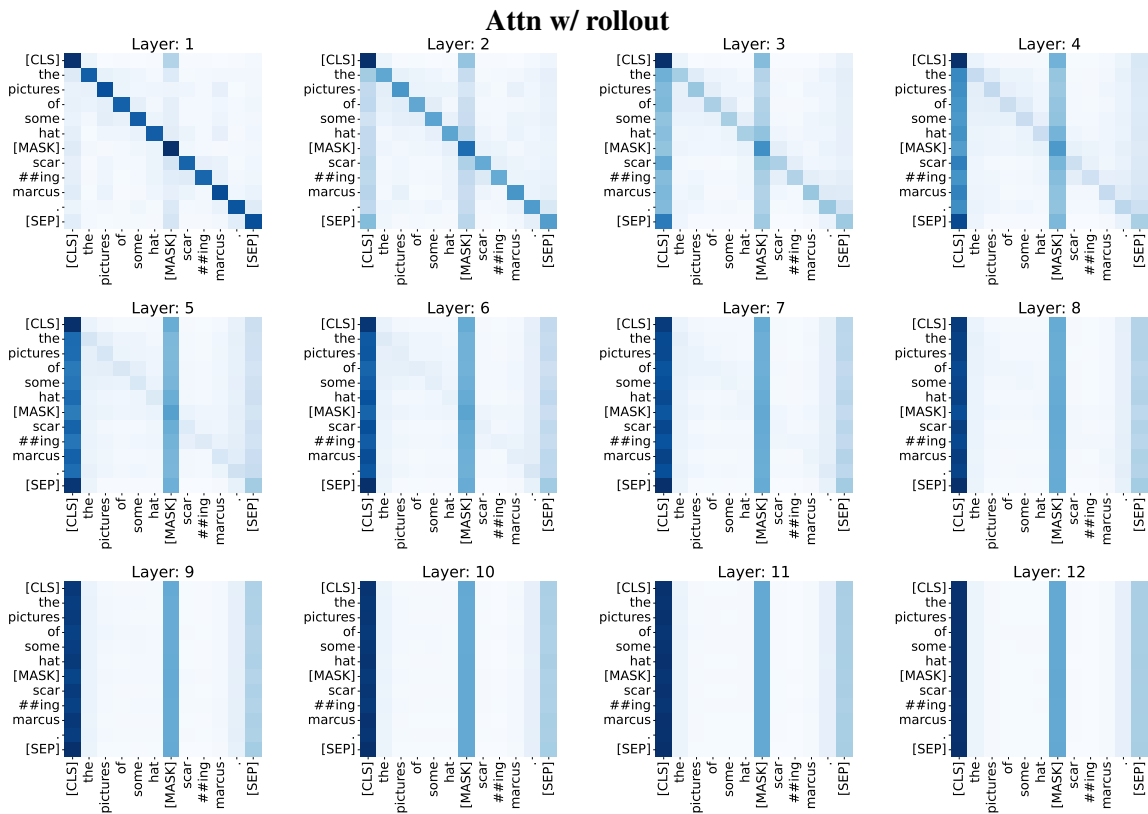


Figure A.6: Raw attention scores (Attn) aggregated by rollout (Abnar and Zuidema (2020)'s method) across layers.

Attn-norm

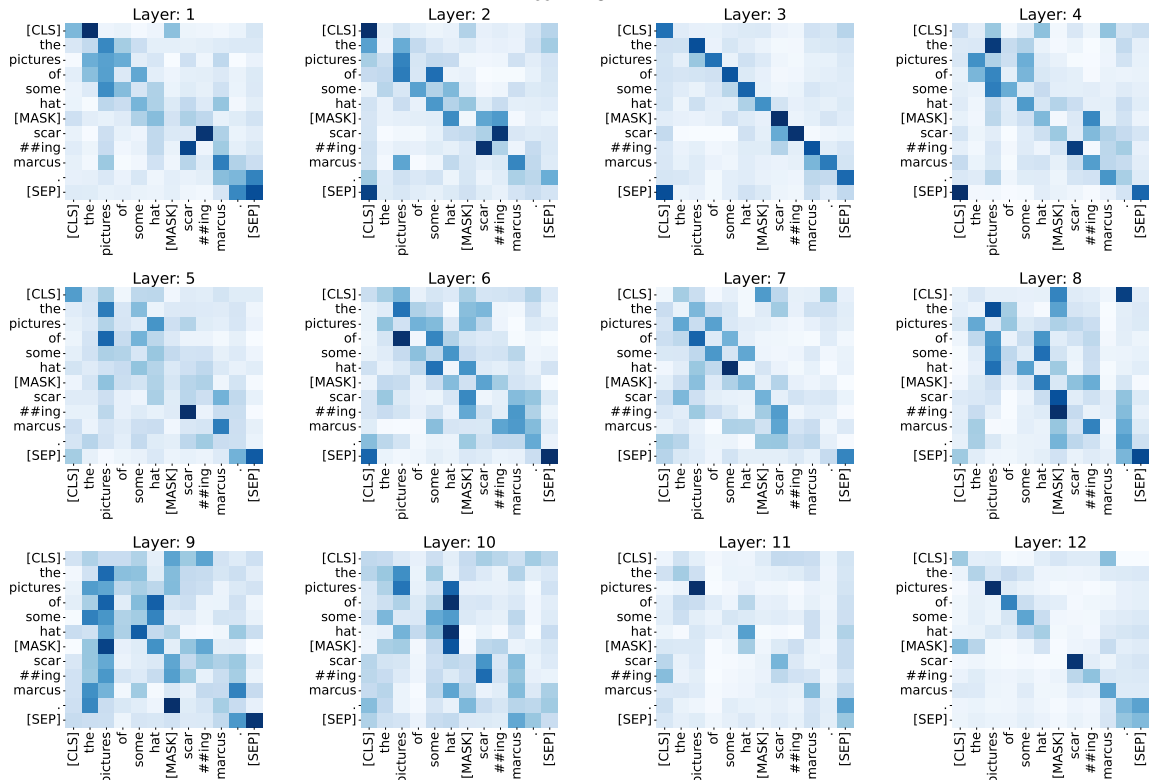


Figure A.7: Kobayashi et al. (2020)'s scores (Attn-norm) across layers.

Attn-norm w/ rollout

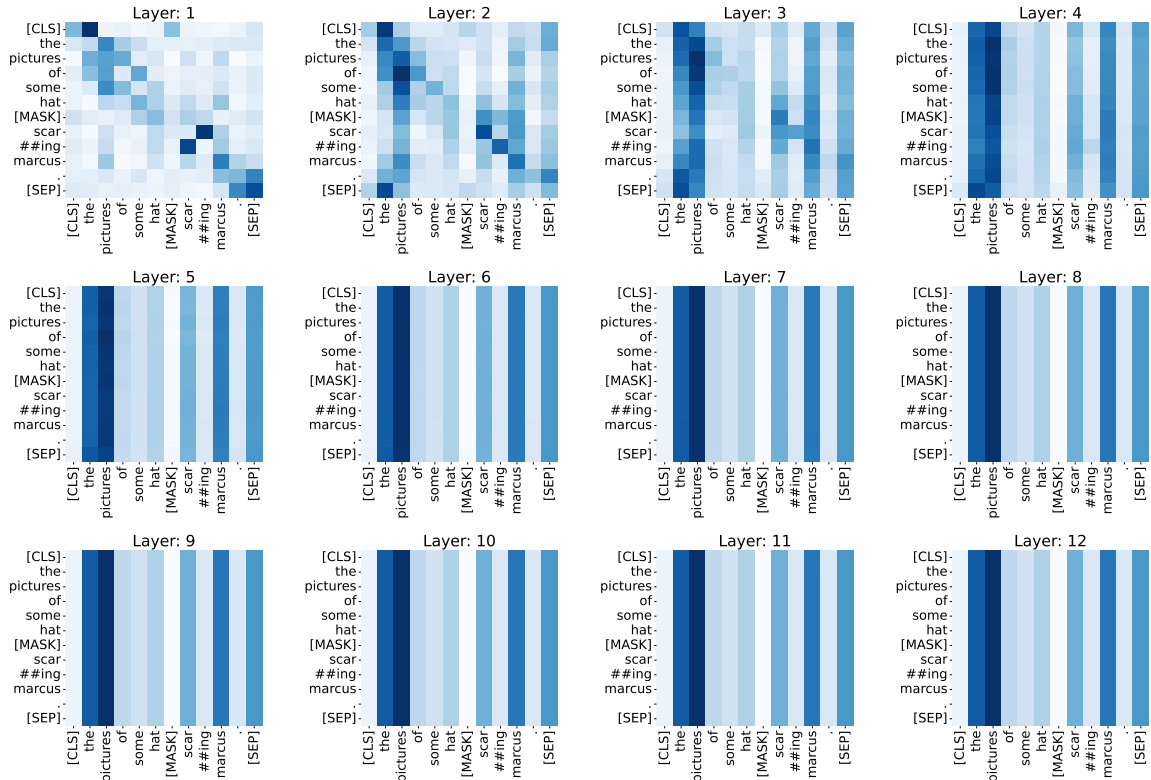


Figure A.8: Kobayashi et al. (2020)'s scores (Attn-norm) aggregated by rollout method across layers.

Attn-norm + RES

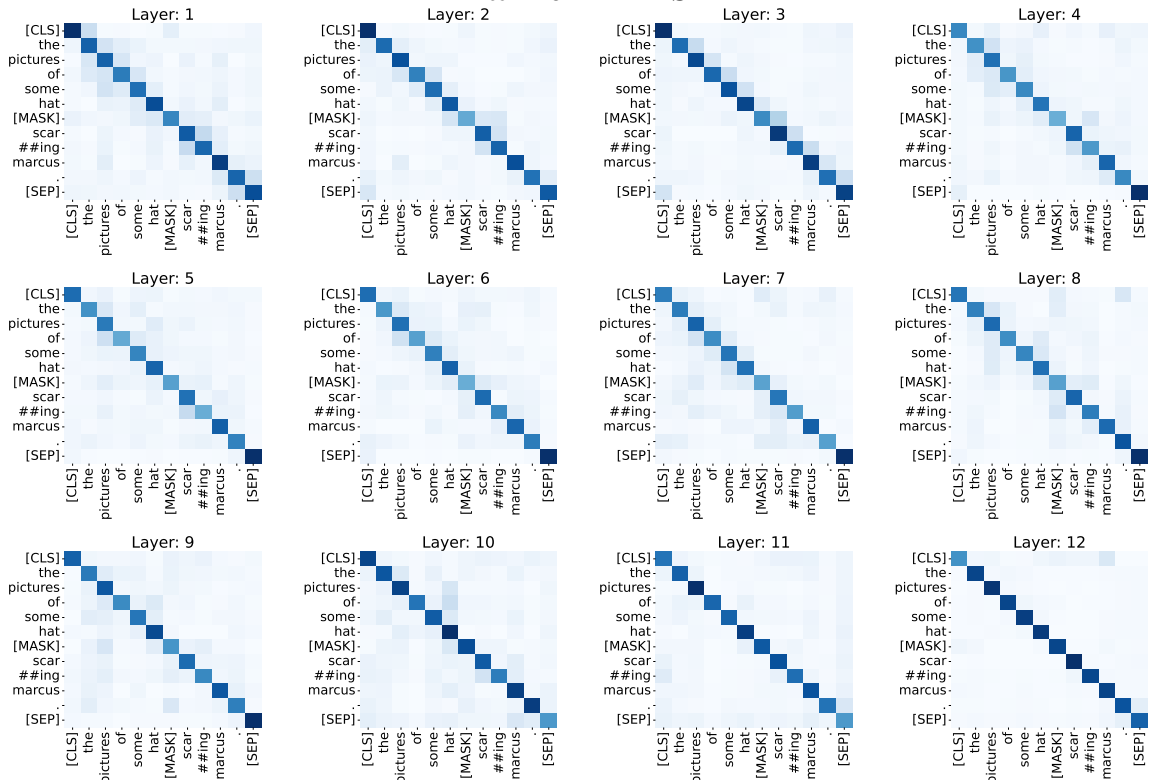


Figure A.9: Kobayashi et al. (2021)'s scores (Attn-norm + RES) across layers.

Attn-norm + RES w/ rollout

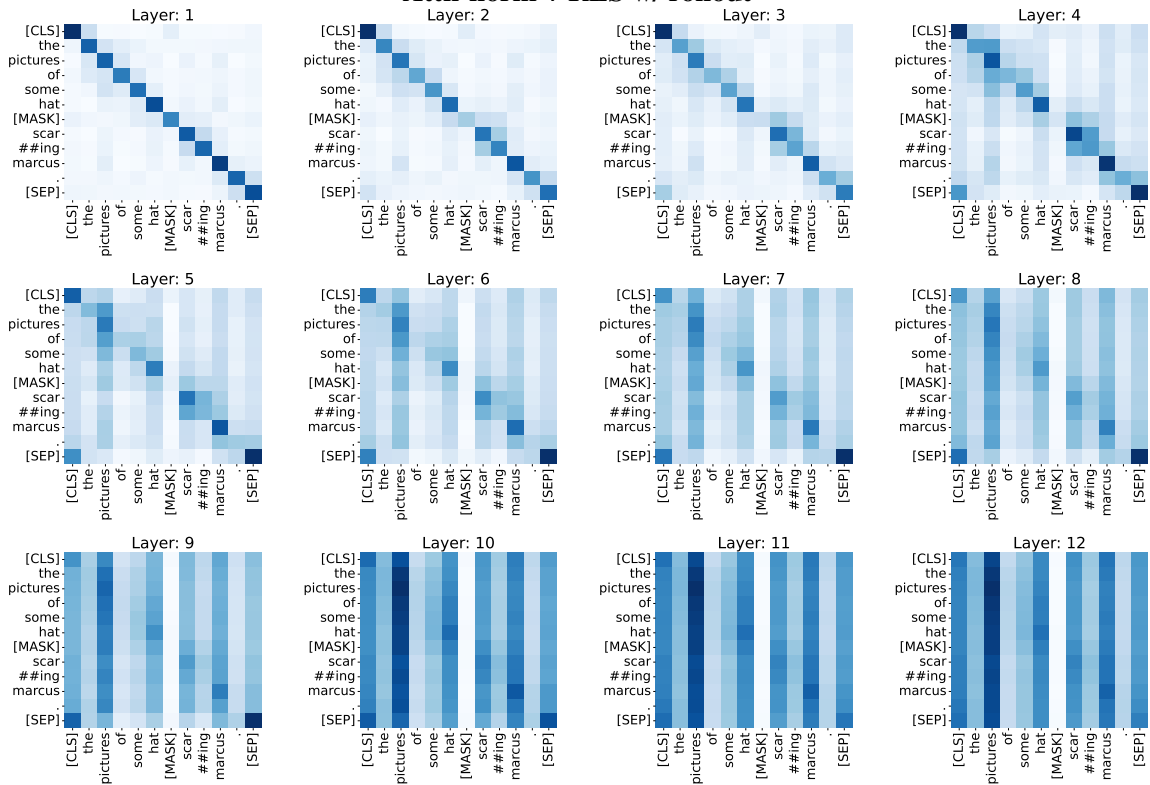


Figure A.10: Kobayashi et al. (2021)'s scores (Attn-norm + RES) aggregated by rollout method across layers.




Figure A.11: Kobayashi et al. (2021)'s scores (Attn-norm + RES + LN) across layers.

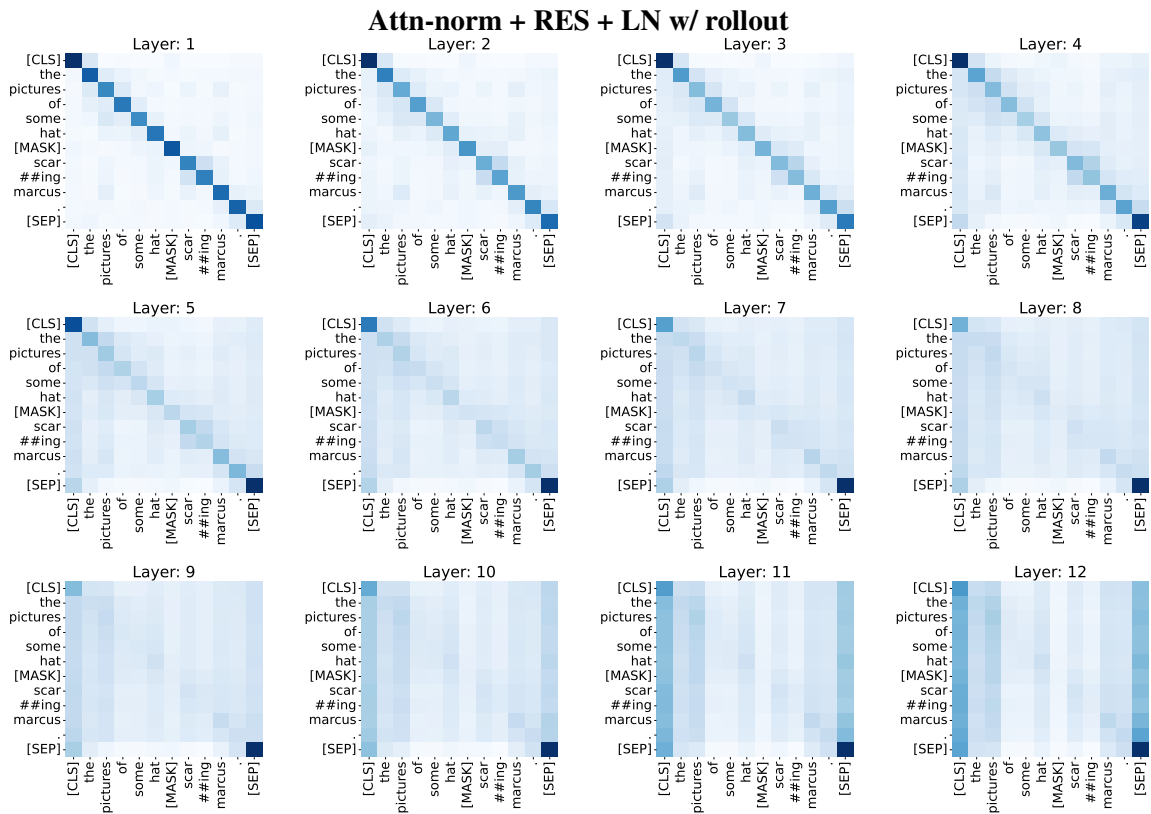


Figure A.12: Kobayashi et al. (2021)'s scores (Attn-norm + RES + LN) aggregated by rollout method across layers.

GlobEnc without rollout

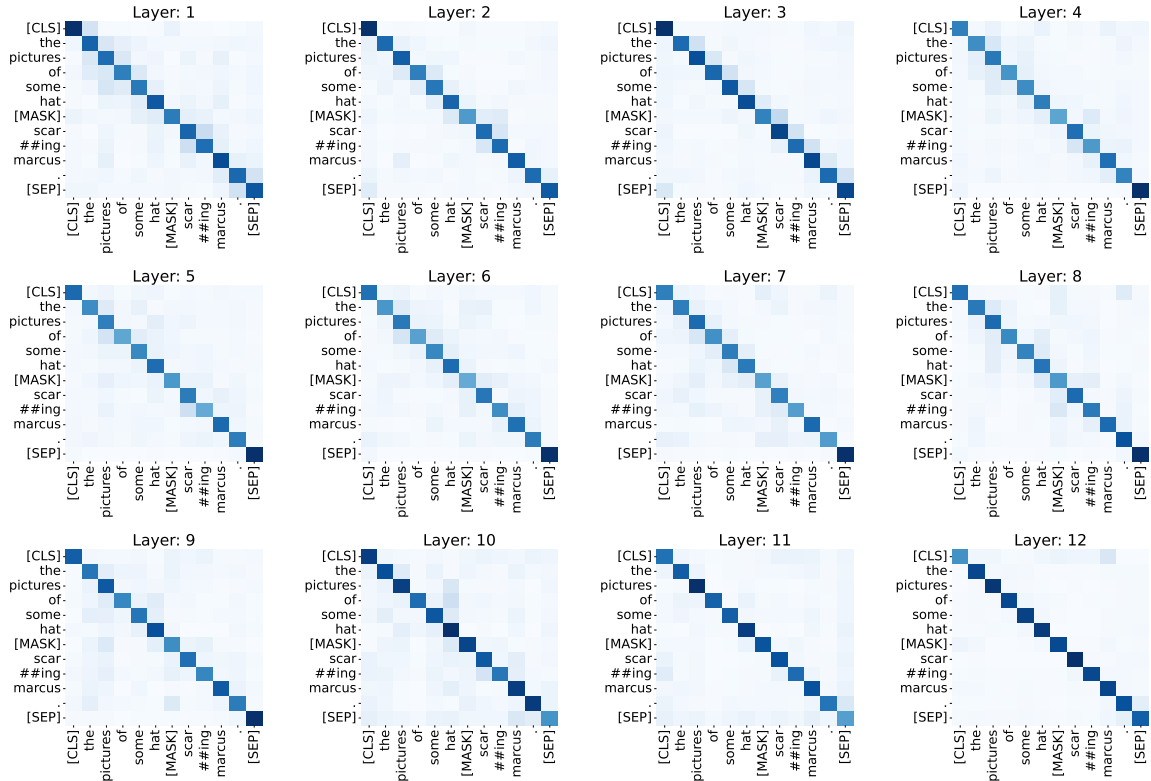


Figure A.13: Modarressi et al. (2022)'s scores (GlobEnc \rightarrow) without aggregation (rollout) across layers.

GlobEnc (w/ rollout)

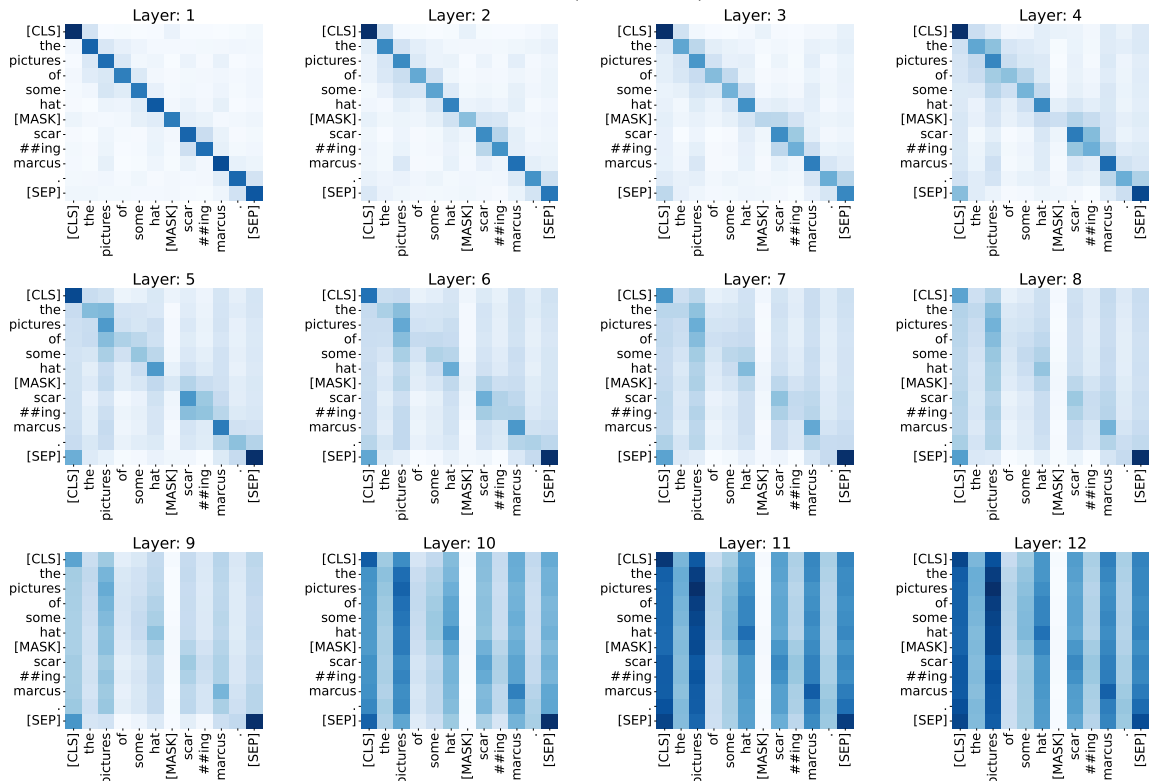


Figure A.14: Modarressi et al. (2022)'s scores (GlobEnc) across layers (the rollout method is inherently incorporated in GlobEnc).

ALTI without rollout

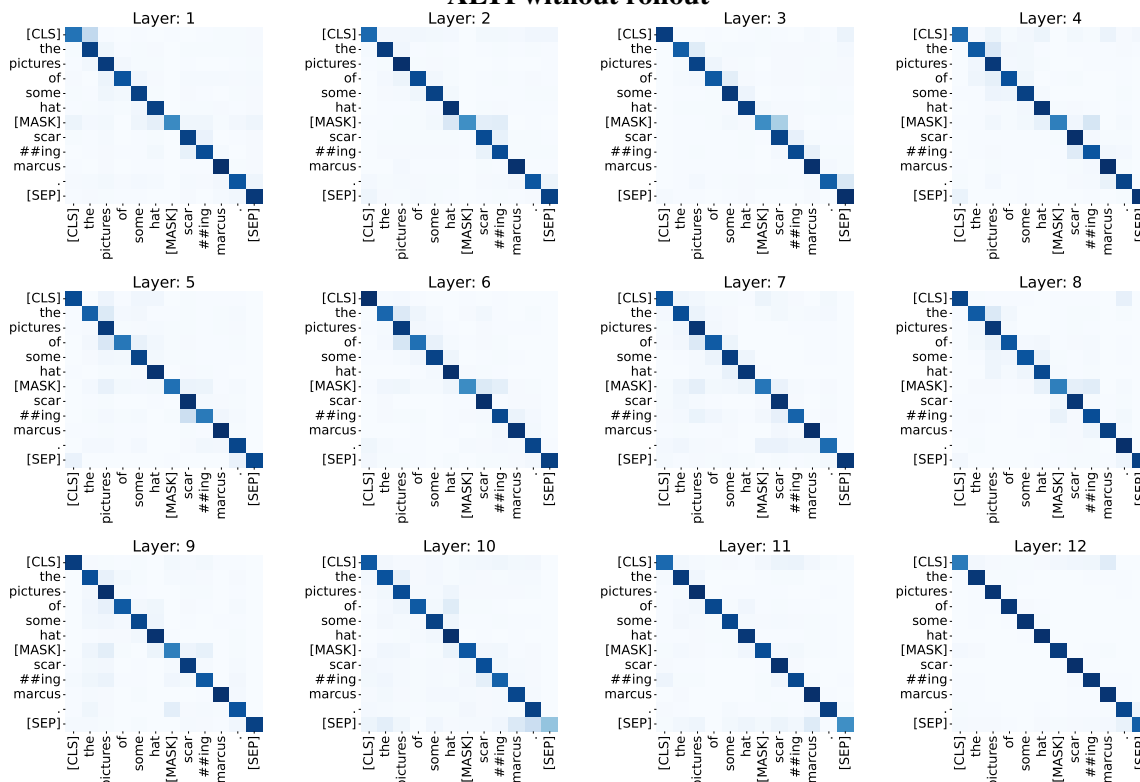


Figure A.15: Ferrando et al. (2022)'s scores (ALTI \rightarrow) without aggregation (rollout) across layers.

ALTI (w/ rollout)

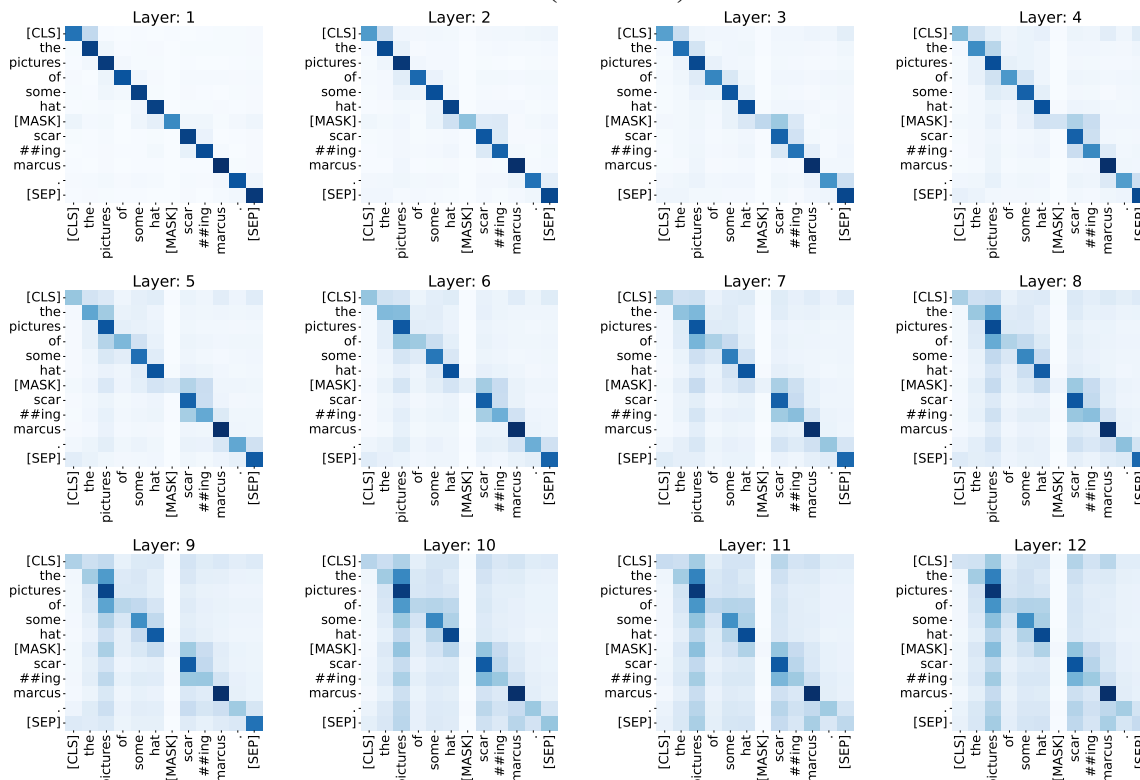


Figure A.16: Ferrando et al. (2022)'s scores (ALTI) across layers (the rollout method is inherently incorporated in ALTI).

Value Zeroing

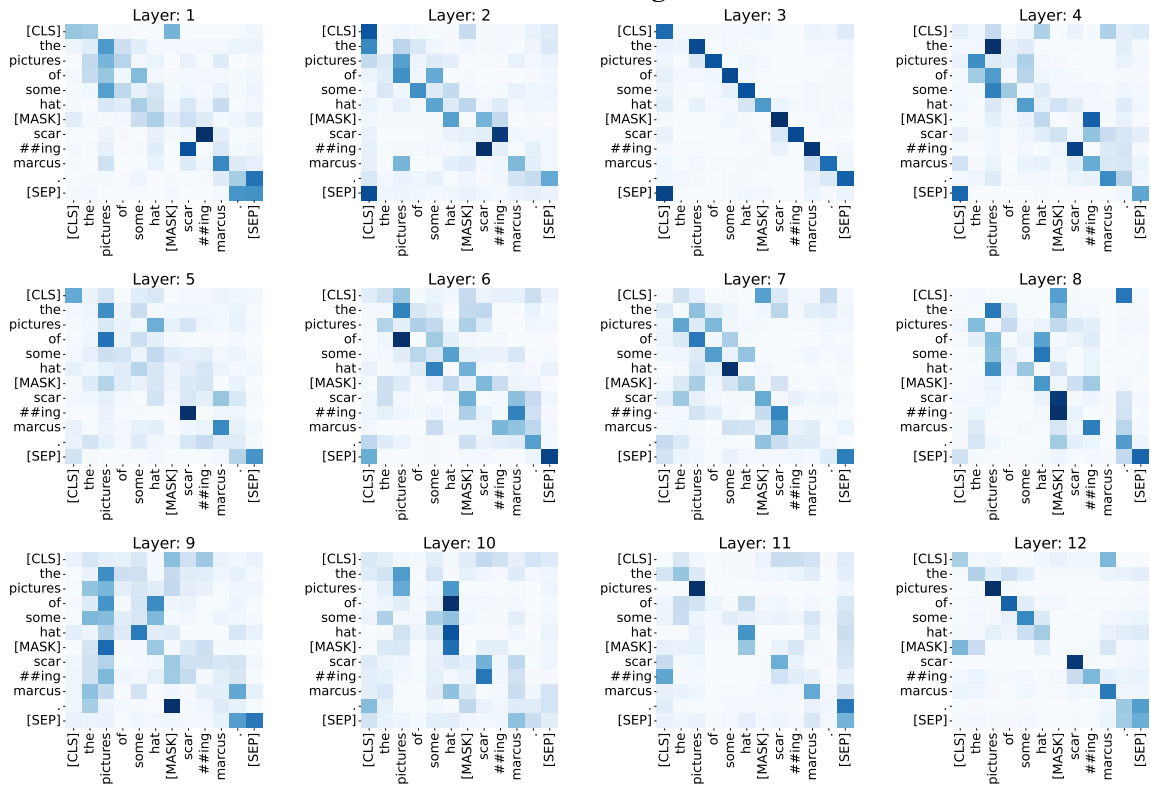


Figure A.17: Our scores (*Value Zeroing*) across layers.

Value Zeroing w/ rollout

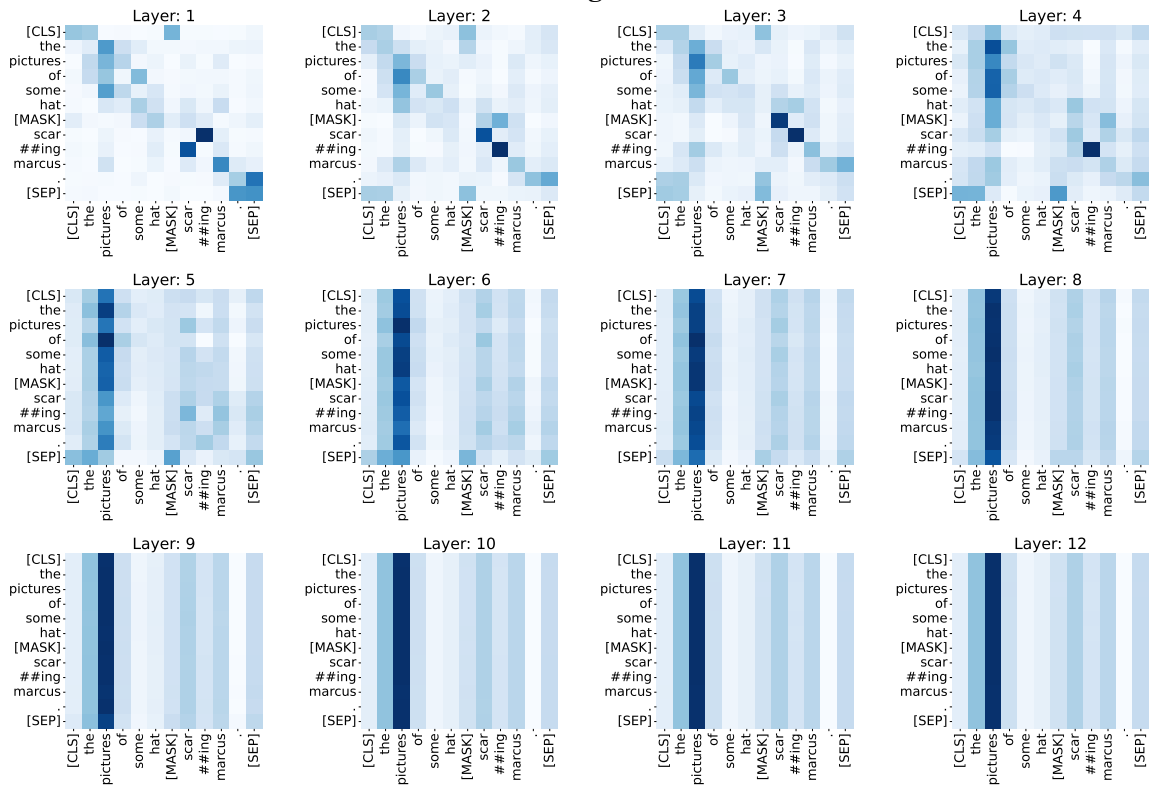


Figure A.18: Our global scores (*Value Zeroing*) aggregated by rollout method across layers.