

# Fixed Encoder Self-Attention Patterns in Transformer-Based Machine Translation

Alessandro Raganato, Yves Scherrer and Jörg Tiedemann

University of Helsinki

{name.surname}@helsinki.fi

## Abstract

Transformer-based models have brought a radical change to neural machine translation. A key feature of the Transformer architecture is the so-called multi-head attention mechanism, which allows the model to focus simultaneously on different parts of the input. However, recent works have shown that most attention heads learn simple, and often redundant, positional patterns. In this paper, we propose to replace all but one attention head of each encoder layer with simple fixed – *non-learnable* – attentive patterns that are solely based on position and do not require any external knowledge. Our experiments with different data sizes and multiple language pairs show that fixing the attention heads on the encoder side of the Transformer at training time does not impact the translation quality and even increases BLEU scores by up to 3 points in low-resource scenarios.

## 1 Introduction

Models based on the Transformer architecture (Vaswani et al., 2017) have led to tremendous performance increases in a wide range of downstream tasks (Devlin et al., 2019; Radford et al., 2019), including Machine Translation (MT) (Vaswani et al., 2017; Ott et al., 2018). One main component of the architecture is the multi-headed attention mechanism that allows the model to capture long-range contextual information. Despite these successes, the impact of the suggested parametrization choices, in particular the self-attention mechanism with its large number of attention heads distributed over several layers, has been the subject of many studies, following roughly four lines of research.

The first line of research focuses on the interpretation of the network, in particular on the analysis of attention mechanisms and the interpretability of the weights and connections (Raganato and Tiedemann, 2018; Tang et al., 2018; Mareček and Rosa,

2019; Voita et al., 2019a; Brunner et al., 2020). A closely related research area attempts to guide the attention mechanism, e.g. by incorporating alignment objectives (Garg et al., 2019), or improving the representation through external information such as syntactic supervision (Pham et al., 2019; Currey and Heafield, 2019; Deguchi et al., 2019). The third line of research argues that Transformer networks are over-parametrized and learn redundant information that can be pruned in various ways (Sanh et al., 2019). For example, Voita et al. (2019b) show that a few attention heads do the “heavy lifting” whereas others contribute very little or nothing at all. Similarly, Michel et al. (2019) raise the question whether 16 attention heads are really necessary to obtain competitive performance. Finally, several recent works address the computational challenge of modeling very long sequences and modify the Transformer architecture with attention operations that reduce time complexity (Shen et al., 2018; Wu et al., 2019; Child et al., 2019; Sukhbaatar et al., 2019; Dai et al., 2019; Indurthi et al., 2019; Kitaev et al., 2020).

This study falls into the third category and is motivated by the observation that most self-attention patterns learned by the Transformer architecture merely reflect positional encoding of contextual information (Raganato and Tiedemann, 2018; Kovaleva et al., 2019; Voita et al., 2019a; Correia et al., 2019). From this standpoint, we argue that most attentive connections in the encoder do not need to be learned at all, but can be replaced by simple predefined patterns. To this end, we design, analyze and experimentally compare intuitive and simple fixed attention patterns. The proposed patterns are solely based on positional information and do not require any learnable parameters nor external knowledge. The fixed patterns reflect the importance of locality and pose the question whether encoder self-attention needs to be learned

at all to achieve state-of-the-art results in machine translation.

This paper provides the following contributions:

- We propose fixed – *non-learnable* – attentive patterns that replace the learnable attention matrices in the encoder of the Transformer model.
- We evaluate the proposed fixed patterns on a series of experiments with different language pairs and varying amounts of training data. The results show that fixed self-attention patterns yield consistently competitive results, especially in low-resource scenarios.
- We provide an ablation study to analyze the relative impact of the different fixed attention heads and the effect of keeping one of the eight heads learnable. Moreover, we also study the effect of the number of encoder and decoder layers on translation quality.
- We assess the translation performance of the fixed attention models through various contrastive test suites, focusing on two linguistic phenomena: subject-verb agreement and word sense disambiguation.

Our results show that the encoder self-attention in Transformer-based machine translation can be simplified substantially, reducing the parameter footprint without loss of translation quality, and even improving quality in low-resource scenarios.

Along with our contributions, we highlight our key findings that give insights for the further development of more lightweight neural networks while retaining state-of-the-art performance for MT:

- The encoder can be substantially simplified with trivial attentive patterns at training time: only preserving adjacent and previous tokens is necessary.
- Encoder attention heads solely based on locality principles may hamper the extraction of global semantic features beneficial for the word sense disambiguation capability of the MT model. Keeping one learnable head in the encoder compensates for degradations, but this trade-off needs to be carefully assessed.
- Position-wise attentive patterns play a key role in low-resource scenarios, both for related (German  $\leftrightarrow$  English) and unrelated (Vietnamese  $\leftrightarrow$  English) languages.

## 2 Related work

Attention mechanisms in Neural Machine Translation (NMT) were first introduced in combination with Recurrent Neural Networks (RNNs) (Cho et al., 2014; Bahdanau et al., 2015; Luong et al., 2015), between the encoder and decoder. The Transformer architecture extended the mechanism by introducing the so-called self-attention to replace the RNNs in the encoder and decoder, and by using multiple attention heads (Vaswani et al., 2017). This architecture rapidly became the *de facto* state-of-the-art architecture for NMT, and more recently for language modeling (Radford et al., 2018) and other downstream tasks (Strubell et al., 2018; Devlin et al., 2019; Bosselut et al., 2019). The Transformer allows the attention for a token to be spread over the entire input sequence, multiple times, intuitively capturing different properties. This characteristic has led to a line of research focusing on the interpretation of Transformer-based networks and their attention mechanisms (Raganato and Tiedemann, 2018; Tang et al., 2018; Mareček and Rosa, 2019; Voita et al., 2019a; Vig and Belinkov, 2019; Clark et al., 2019; Kovaleva et al., 2019; Tenney et al., 2019; Lin et al., 2019; Jawahar et al., 2019; van Schijndel et al., 2019; Hao et al., 2019b; Rogers et al., 2020). As regards MT, recent work (Voita et al., 2019b) suggests that only a few attention *heads* are specialized towards a specific role, e.g., focusing on a syntactic dependency relation or on rare words, and significantly contribute to the translation performance, while all the others are dispensable.

At the same time, recent research has attempted to bring the mathematical formulation of self-attention more in line with the linguistic expectation that attention would be most useful within a narrow local scope, e.g. for the translation of phrases (Hao et al., 2019a). For instance, Shaw et al. (2018) replace the sinusoidal position encoding in the Transformer with relative position encoding to improve the focus on local positional patterns. Several studies modify the attention formula to bias the attention weights towards local areas (Yang et al., 2018; Xu et al., 2019; Fonolosa et al., 2019). Wu et al. (2019) and Yang et al. (2019) use convolutional modules to replace parts of self-attention, making the overall networks computationally more efficient. Cui et al. (2019) mask out certain tokens when computing attention, which favors local attention patterns and prevents redun-

dancy in the different attention heads. All these contributions have shown the importance of localness, and the possibility to use lightweight convolutional networks to reduce the number of parameters while yielding competitive results (Wu et al., 2019). In this respect, our work is orthogonal to previous research: we focus only on the original Transformer architecture and investigate the replacement of learnable encoder self-attention by fixed, non-learnable attentive patterns.

Recent analysis papers have identified a certain number of functions to which different self-attention heads tend to specialize. Voita et al. (2019b) identifies three types of heads: *positional heads* point to an adjacent token, *syntactic heads* point to tokens in a specific syntactic relation, and *rare word heads* point to the least frequent tokens in a sentence. Correia et al. (2019) identify two additional types of heads: *BPE-merging heads* spread weight over adjacent tokens that are part of the same BPE cluster or hyphenated words, and *interrogation heads* point to question marks at the end of the sentence. In line with these findings, we design our fixed attention patterns and train NMT models without the need of learning them.

In concurrent work, You et al. (2020) propose to replace learnable attention weights in Transformer-based NMT with hard-coded Gaussian distributions. This paper is complementary and differs in several respects: while You et al. (2020) consider three fixed patterns across the encoder-decoder architecture, we focus only on the encoder self-attention but present seven fixed patterns that cover additional known properties of self-attention. We study the relative impact of each of them and analyze their performance with respect to different numbers of encoder-decoder layers, and as semantic feature extractor for lexical ambiguity phenomena. Furthermore, in contrast to You et al. (2020), we show that our fixed patterns have a clear beneficial effect in low-resource scenarios.

### 3 Methodology

In this section, we briefly describe the Transformer architecture and its self-attention mechanism, and introduce the fixed attention patterns.

#### 3.1 Self-attention in Transformers

The Transformer architecture follows the so-called encoder-decoder paradigm where the source sentence is encoded in a number of stacked encoder

blocks, and the target sentence is generated through a number of stacked decoder blocks. Each encoder block consists of a multi-head self-attention layer and a feed-forward layer.

For a sequence of token representations  $\mathbf{H} \in \mathbb{R}^{n \times d}$  (with sequence length  $n$  and dimensionality  $d$ ), the self-attention model first projects them into queries  $\mathbf{Q} \in \mathbb{R}^{n \times d}$ , keys  $\mathbf{K} \in \mathbb{R}^{n \times d}$  and values  $\mathbf{V} \in \mathbb{R}^{n \times d}$ , using three different linear projections. Then, the attention energy  $\xi_i$  for position  $i$  in the sequence is computed by taking the scaled dot product between the query vector  $\mathbf{Q}_i$  and the key matrix  $\mathbf{K}$ :

$$\xi_i = \text{softmax} \left( \frac{\mathbf{Q}_i \mathbf{K}^\top}{\sqrt{d}} \right) \in \mathbb{R}^n \quad (1)$$

The attention energy is then used to compute a weighted average of the values  $\mathbf{V}$ :

$$\text{Att}(\xi_i, \mathbf{V}) = \xi_i \mathbf{V} \in \mathbb{R}^d \quad (2)$$

For multi-head attention with  $h$  heads, the query, key and value are linearly projected  $h$  times to allow the model to jointly attend to information from different representations. The attention vectors of the  $h$  heads are then concatenated. Finally, the resulting multi-head attention is fed to a feed-forward network that consists of two linear layers with a ReLU activation in between. This multi-head attention is often called *encoder self-attention*, as it builds a representation of the input sentence that is attentive to itself.

The decoder follows the same architecture as the encoder with multi-head attention mechanisms and feed-forward networks, with two main differences: i) an additional multi-head attention mechanism, called *encoder-decoder attention*, connects the last encoder layer to the decoder layers, and ii) future positions are prevented from being attended to, by masking, in order to preserve the auto-regressive property of a left-to-right decoder.

The *base* version of the Transformer, the standard setting for MT, uses 6 layers for both encoder and decoder and 8 attention heads in each layer. In this work, we focus on the *encoder self-attention* and replace the learned attention energy  $\xi$  by fixed, predefined distributions for all but one head.

#### 3.2 Fixed self-attention patterns

The inspection of encoder self-attention in standard MT models yields the somewhat surprising result that positional patterns, such as “previous token”,

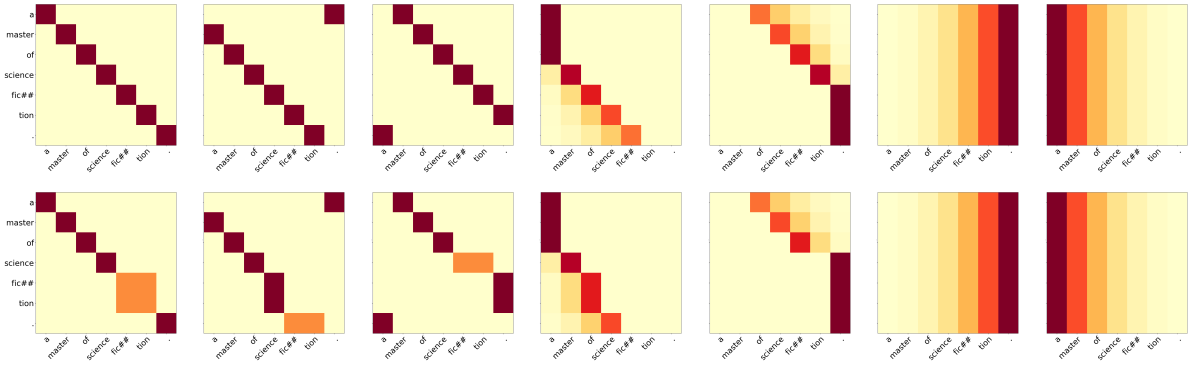


Figure 1: Token-based (upper row) and word-based (lower row) fixed attention patterns for the example sentence “a master of science fic## tion.”. The word-based patterns treat the subwords “fic##” and “tion” as a single token.

“next token”, or “last token of the sentence”, are key features across all layers and remain even after pruning most of the attention heads (Voita et al., 2019a,b; Correia et al., 2019). Instead of costly learning these trivial positional patterns using millions of sentences, we choose seven predefined patterns, each of which takes the place of an attention head (see Figure 1, upper row).

Given the  $i$ -th word within a sentence of length  $n$ , we define the following patterns:

1. the current token: a fixed attention weight of 1.0 at position  $i$ ,
2. the previous token: a fixed attention weight of 1.0 at position  $i - 1$ ,
3. the next token: a fixed attention weight of 1.0 at position  $i + 1$ ,
4. the larger left-hand context: a function  $f$  over the positions 0 to  $i - 2$ ,
5. the larger right-hand context: a function  $f$  over the positions  $i + 2$  to  $n$ ,
6. the end of the sentence: a function  $f$  over the positions 0 to  $n$ ,
7. the start of the sentence: a function  $f$  over the positions  $n$  to 0.

For illustration, the attention energies for patterns 2 and 4 are defined formally as follows:

$$\xi_{i,j}^{(2)} = \begin{cases} 1 & \text{if } j = i - 1 \\ 0 & \text{otherwise} \end{cases}$$

$$\xi_{i,j}^{(4)} = \begin{cases} f^{(4)}(j) & \text{if } j \leq i - 2 \\ 0 & \text{otherwise} \end{cases}$$

where

$$f^{(4)}(j) = \frac{(j + 1)^3}{\sum_{j=0}^{i-2} (j + 1)^3}$$

The same function is used for all patterns, changing

only the respective start and end points.<sup>1</sup> These predefined attention heads are repeated over all layers of the encoder. The eighth attention head always remains learnable.<sup>2</sup>

It is customary in NMT to split words into subword units, and there is evidence that self-attention treats split words differently than non-split ones (Correia et al., 2019). Therefore, we propose a second variant of the predefined patterns that assigns the same attention values to all parts of the same word (see lower row of Figure 1).

## 4 Experiments

We perform a series of experiments to evaluate the fixed attentive encoder patterns, starting with a standard German  $\leftrightarrow$  English translation setup (Section 4.1) and then extending the scope to low-resource and high-resource scenarios (Section 4.2). We use the OpenNMT-py (Klein et al., 2017) library for training, the *base* version of Transformer as hyper-parameters (Vaswani et al., 2017), and compare against the reference using sacreBLEU (Papineni et al., 2002; Post, 2018).<sup>3</sup>

### 4.1 Results: Standard scenario

To assess the general viability of the proposed approach and to quantify the effects of different numbers of encoder and decoder layers, we train models on a mid-sized dataset of 2.9M training sentences from the German  $\leftrightarrow$  English WMT19 news translation task (Barrault et al., 2019), using newstest2013 and newstest2014 as development and test data, respectively. We learn truecasers and Byte-Pair En-

<sup>1</sup>Pattern 5 and 7 are flipped versions of pattern 4 and 6, respectively.

<sup>2</sup>In Section 4.4, we present a contrastive system in which the eighth head is fixed as well.

<sup>3</sup>Signature: BLEU+case.lc+#.1+s.exp+tok.13a+v.1.2.11.

		Encoder+Decoder layers						
Encoder heads		1+1	2+1	3+1	4+1	5+1	6+1	6+6
EN-DE	8L	20.61	21.68	22.63	23.02	23.18	23.36	25.02
	7F <sub>token</sub> +1L	20.61	21.58	22.38	23.15	23.10	23.07	24.63
	7F <sub>word</sub> +1L	19.72*	21.43	21.81*	22.83	22.74	22.88*	24.85
	1L	18.14*	19.88*	21.42*	21.71*	22.63*	22.29*	23.87*
DE-EN	8L	25.66	27.28	27.88	28.62	28.71	29.31	30.99
	7F <sub>token</sub> +1L	24.90*	27.01	26.84*	28.09*	28.43	28.61*	30.61
	7F <sub>word</sub> +1L	25.03*	26.72*	27.38*	27.78*	27.82*	28.40*	30.69
	1L	23.76*	25.75*	26.96*	27.34*	27.44*	27.56*	30.17*

Table 1: BLEU scores for the German  $\leftrightarrow$  English (DE  $\leftrightarrow$  EN) standard scenario, for different configurations of learnable (L) and fixed (F) attention heads. Scores marked in gray with \* are significantly lower than the respective 8L model scores, at  $p < 0.05$ . Statistical significance is computed using the *compare-mt* tool (Neubig et al., 2019) with paired bootstrap resampling with 1000 resamples (Koehn, 2004).

coding (BPE) segmentation (Sennrich et al., 2016) on the training corpus, using 35 000 merge operations.

We train four Transformer models:

- 8L: all 8 attention heads in each layer are learnable,
- 7F<sub>token</sub>+1L: 7 fixed token-based attention heads and 1 learnable head per encoder layer,
- 7F<sub>word</sub>+1L: 7 fixed word-based attention patterns and 1 learnable head per encoder layer,
- 1L: a single learnable attention head per encoder layer.

Each model is trained in 7 configurations: 6 encoder layers with 6 decoder layers, and 1 to 6 encoder layers coupled to 1 decoder layer. BLEU scores are shown in Table 1.

Results for the most powerful model (6+6) show that the two fixed-attention models are almost indistinguishable from the standard model, whereas the single-head model yields consistently slightly lower results. It could be argued that the 6-layer decoder is powerful enough to compensate for deficiencies due to fixed attention on the encoder side. The 6+1 configuration, which uses a single layer decoder, shows indeed a slight performance drop for German  $\rightarrow$  English, but no significant difference in the opposite direction. Overall translation quality drops significantly with three and less encoder layers, but the difference between fixed and learnable attention models is statistically insignificant in most cases. The fixed attention models always outperform the model with a single learnable head, which shows that the predefined patterns

are indeed helpful. The (simpler) token-based approach seems to outperform the word-based one, but with higher numbers of decoder layers the two variants are statistically equivalent.

## 4.2 Results: Low-resource and high-resource scenarios

We hypothesize that fixed attentive patterns are especially useful in low-resource scenarios since intuitive properties of self-attention are directly encoded within the model, which may be hard to learn from small training datasets. We empirically test this assumption on four translation tasks:

- German  $\rightarrow$  English (DE $\rightarrow$ EN), using the data from the IWSLT 2014 shared task (Cettolo et al., 2014). As prior work (Ranzato et al., 2016; Sennrich and Zhang, 2019), we report BLEU score on the concatenated dev sets: tst2010, tst2011, tst2012, dev2010, dev2012 (159 000 training sentences, 7 282 for development, and 6 750 for testing).
- Korean  $\rightarrow$  English (KO $\rightarrow$ EN), using the dataset described in Park et al. (2016) (90 000 training sentences, 1 000 for development, and 2 000 for testing).<sup>4</sup>
- Vietnamese  $\leftrightarrow$  English (VI $\leftrightarrow$ EN), using the data from the IWSLT 2015 shared task (Cettolo et al., 2015), using tst2012 and tst2013 for development and testing, respectively (133 000 training sentences, 1 553 for development and 1 268 per testing).

<sup>4</sup><https://github.com/jungyeul/korean-parallel-corpora>

Enc. heads	DE-EN	KO-EN	EN-VI	VI-EN
8L	30.86	6.67	29.85	26.15
7F <sub>token</sub> +1L	<b>32.95</b>	8.43	31.05	<b>29.16</b>
7F <sub>word</sub> +1L	32.56	<b>8.70</b>	<b>31.15</b>	28.90
1L	30.22	6.14	28.67	25.03
Prior work	† 33.60	† 10.37	‡ 27.71	‡ 26.15

Table 2: BLEU scores obtained for the low-resource scenarios with 6+6 layer configuration. Results marked with † are taken from [Sennrich and Zhang \(2019\)](#), those marked with ‡ from [Kudo \(2018\)](#).

Encoder heads	EN-DE	DE-EN
8L	26.75	<b>34.10</b>
7F <sub>token</sub> +1L	26.52	33.50
7F <sub>word</sub> +1L	<b>26.92</b>	33.17
1L	26.26	32.91

Table 3: BLEU scores obtained for the high-resource scenario with 6+6 layer configurations.

Low-resource scenarios can be sensible to the choice of hyperparameters ([Sennrich and Zhang, 2019](#)). Hence, we apply three of the most successful adaptations to all our configurations: reduced batch size (4k  $\rightarrow$  1k tokens), increased dropout (0.1  $\rightarrow$  0.3), and tied embeddings. Sentences are BPE-encoded with 30 000 merge operations, shared between source and target language, but independent for Korean  $\rightarrow$  English.

Results of the 6+6 layer configurations are shown in Table 2.<sup>5</sup> The models using fixed attention consistently outperform the models using learned attention, by up to 3 BLEU. No clear winner between token-based and word-based fixed attention can be distinguished though.

Our English  $\leftrightarrow$  Vietnamese models outperform prior work based on an RNN architecture by a large margin, but the German  $\rightarrow$  English and Korean  $\rightarrow$  English models remain below the heavily optimized models of [Sennrich and Zhang \(2019\)](#). However, we note that our goal is not to beat the state-of-the-art in a given MT setting but rather to show the performance of simple non-learnable attentive patterns across different language pairs and data sizes. Moreover, it is worth to mention that the Korean  $\rightarrow$  English dataset, being automatically created, includes some noise in the test data that

<sup>5</sup>The 6+1 models show globally lower scores, but similar relative rankings between models.

Disabled head	6+1 layers		6+6 layers	
	EN-DE	DE-EN	EN-DE	DE-EN
1 Current word	-0.15	0.11	0.12	-0.04
2 Previous word	<b>-5.72</b>	<b>-5.21</b>	<b>-3.05</b>	<b>-3.26</b>
3 Next word	<b>-1.80</b>	<b>-1.98</b>	<b>-2.08</b>	<b>-1.36</b>
4 Prev. context	<b>-4.73</b>	<b>-5.20</b>	<b>-1.42</b>	<b>-2.85</b>
5 Next context	-0.72	-0.34	-0.47	-0.66
6 Start context	-0.17	-0.12	0.14	0.13
7 End context	-0.02	0.12	-0.30	0.10
8 Learned head	<b>-2.22</b>	<b>-4.05</b>	-0.58	-0.78
	EN-VI	VI-EN	EN-VI	VI-EN
1 Current word	0.12	-0.14	0.16	-0.05
2 Previous word	<b>-2.32</b>	<b>-2.67</b>	<b>-2.71</b>	<b>-3.04</b>
3 Next word	<b>-1.12</b>	<b>-1.61</b>	<b>-1.35</b>	<b>-2.15</b>
4 Prev. context	<b>-4.11</b>	<b>-4.32</b>	<b>-2.82</b>	<b>-3.09</b>
5 Next context	-0.27	-0.50	-0.83	-0.77
6 Start context	-0.29	-0.08	-0.04	0
7 End context	0.28	-0.29	-0.23	-0.19
8 Learned head	-0.57	-0.88	-0.18	0.36

Table 4: BLEU score differences with respect to the full model with 8 enabled heads (values  $< -1$  in bold).

may impact the comparison.<sup>6</sup>

Finally, we also evaluate a high-resource scenario for German  $\leftrightarrow$  English with 11.5M training sentences.<sup>7</sup> Table 3 shows that the results of the fixed attention models do not degrade even when abundant training data allow all attention heads to be learned accurately.

### 4.3 Ablation study

We perform an ablation study to assess the contribution of each attention head separately. To this end, we mask out one attention pattern across all encoder layers at test time. Table 4 shows the differences compared to the full model, on the mid-sized German  $\leftrightarrow$  English and on the Vietnamese  $\leftrightarrow$  English models, both in the 6+1 and 6+6 layer configurations.

We find that heads 2, 3 and 4 (previous word, next word, previous context) are particularly important, whereas the impact of the remaining context heads is small. Head 1 (current word) is not useful in the token-based model, but shows slightly larger numbers in the word-based setting.

<sup>6</sup><https://github.com/jungyeul/korean-parallel-corpora/issues/1>

<sup>7</sup>This scenario uses the same benchmark from Section 4.1, increasing the training data with a filtered version of ParaCrawl (bicleaner filtered v3).

Enc. heads	#Param.	EN-DE	DE-EN	EN-VI	VI-EN
8L	91.7M	25.02	30.99	29.85	26.15
7F <sub>token</sub> +1L	88.9M	24.63	30.61	31.05	29.16
8F <sub>token</sub>	88.5M	24.64	30.56	31.45	28.97

Table 5: BLEU scores for the experiments with eight fixed attention heads and 6+6 layers. “#Param.” denotes the number of parameters for the EN-DE model.

The most interesting results concern the eighth, learned head. Its impact is significant, but in most cases lower than the three main heads listed above. Interestingly, disabling it causes much lower degradation in the 6+6 configurations, which suggests that a more powerful decoder can compensate for the absence of learned encoder representations.

#### 4.4 Eight fixed heads

The ablation study suggests that it is not crucial to keep one learnable head in the encoder layers, especially if the decoder is deep enough. Here, we assess the extreme scenario where the eighth attention head is fixed as well. The eighth fixed attentive pattern focuses on the last token, with a fixed weight of 1.0 at position  $n$ . Table 5 shows the results for the standard English  $\leftrightarrow$  German scenario and the low-resource English  $\leftrightarrow$  Vietnamese scenario.

Overall, the learnable attention head is completely dispensable across both language pairs. As shown in Section 4.3, the impact of having learnable attention heads on the encoder side is negligible. Moreover, we also note that as we replace attention heads with non-learnable ones, our configurations reduce the number of parameters without degrading translation quality.

## 5 Analysis

To further analyze the fixed attentive encoder patterns, we perform three targeted evaluations: i) on the sentence length, ii) on the subject-verb agreement task, and iii) on the Word Sense Disambiguation (WSD) task. The length analysis inspects the translation quality by sentence length. The subject-verb agreement task is commonly used to evaluate long-range dependencies (Linzen et al., 2016; Tran et al., 2018; Tang et al., 2018), while the WSD task addresses lexical ambiguity phenomena, i.e., words of the source language that have multiple translations in the target language representing different meanings (Marvin and Koehn, 2018; Liu




Figure 2: BLEU scores for different ranges of sentence lengths.

et al., 2018; Pu et al., 2018; Tang et al., 2019).

For both tasks, we use contrastive test suites (Sennrich, 2017; Popović and Castilho, 2019) that rely on the ability of NMT systems to score given translations. Broadly speaking, a sentence containing the linguistic phenomenon of interest is paired with the correct reference translation and with a modified translation with a specific type of error. A contrast is considered successfully detected if the reference translation obtains a higher score than the artificially modified translation. The evaluation metric corresponds to the accuracy over all decisions.

We conduct the analyses using the DE-EN models from Section 4.1, i.e., 8L, 7F<sub>token</sub>+1L, and 7F<sub>token</sub> (H8 disabled).

### 5.1 Sentence length analysis

To assess whether our fixed attentive patterns may hamper modeling of global dependencies supporting long sentences, we compute BLEU score by reference sentence length.<sup>8</sup> Despite the small performance gap between models, as we can see from Figure 2, long sentences benefit from having *learnable* attentive patterns. This is clearly shown by the 7F<sub>token</sub> (H8 disabled) model, which is consistently degraded in almost every length bin.

### 5.2 Subject-verb agreement

The predefined attention patterns focus on relatively small local contexts. It could therefore be argued that the fixed attention models would perform worse on long-distance agreement, and that

<sup>8</sup>We use the *compare-mt* toolkit by Neubig et al. (2019).




Figure 3: Subject-verb agreement accuracies of the EN-DE models. The x-axis shows distances between the subject and the verb.

disabling the learned head in particular would be catastrophic. We test this hypothesis by evaluating the models from Section 4.1 on the subject-verb agreement task of the English-German Lingeval test suite (Sennrich, 2017). Figure 3 plots the accuracies by distance between subject and verb. In the 6+6 layer configuration, no difference can be detected between the three examined scenarios. In the 6+1 layer configuration, the fixed-attention model does not seem to suffer from degraded results, whereas disabling the learned head leads to clearly lower results. This drop is due to the expected degradation of general translation quality (cf. Table 4, ablation study) and is not worse than the degradation observed by disabling one of the fixed local context heads.

### 5.3 Word sense disambiguation

It has been shown that the encoder of Transformer-based MT models includes semantic information beneficial for WSD (Tang et al., 2018, 2019). In this respect, a model with predefined fixed patterns may struggle to encode global semantic features. To this end, we evaluate our models on two German-English WSD test suites, ContraWSD (Rios Gonzales et al., 2017) and MuCoW (Raganato et al., 2019).<sup>9</sup>

Table 6 shows the performance of our models on the WSD benchmarks. Overall, the model with 6 decoder layers and fixed attentive patterns

<sup>9</sup>As MuCoW is automatically built using various parallel corpora, we discarded those ones included in our training. We only report the average result from the TED (Cettolo et al., 2013) and Tatoeba (Tiedemann, 2012) sources.

Encoder heads	ContraWSD		MuCoW	
	6+1	6+6	6+1	6+6
8L	<b>0.804</b>	0.831	<b>0.741</b>	0.761
7F <sub>token</sub> +1L	0.793	<b>0.834</b>	0.734	<b>0.772</b>
7F <sub>token</sub> (H8 disabled)	0.761	0.816	0.721	0.757

Table 6: Accuracy scores of the German-English models on the ContraWSD and MuCoW test suites.

(7F<sub>token</sub>+1L) achieves higher accuracy than the model with all learnable attention heads (8L), while the 1-layer decoder models show the opposite effect. It appears that having 6 decoder layers can effectively cope with WSD despite having only one learnable attention head. Interestingly enough, when we disable the learnable attention head (7F<sub>token</sub> H8 disabled), performance drops consistently in both test suites, showing that the learnable head plays a key role for WSD, specializing in semantic feature extraction.

## 6 Conclusion

In this work, we propose to simplify encoder self-attention of Transformer-based NMT models by replacing all but one attention heads with fixed positional attentive patterns that require neither training nor external knowledge.

We train NMT models on different data sizes and language directions with the proposed fixed patterns, showing that the encoder self-attention can be simplified drastically, reducing parameter footprint at training time without degradation in translation quality. In low-resource scenarios, translation quality is even improved. Our extensive analyses show that i) only adjacent and previous token attentive patterns contribute significantly to the translation performance, ii) the trainable encoder head can also be disabled without hampering translation quality if the number of decoder layers is sufficient, iii) encoder attention heads based on locality patterns are beneficial in low-resource scenarios, but may affect the semantic feature extraction necessary for addressing lexical ambiguity phenomena.

Apart from the consistent results given by our simple fixed encoder patterns, this work opens up potential further research for simpler and more efficient neural networks for MT, such as synthetic self-attention patterns (Tay et al., 2020).



## Acknowledgments



This work is part of the FoTran project, funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 771113).

The authors gratefully acknowledge the support of the CSC – IT Center for Science, Finland, for computational resources. Finally, We would also like to acknowledge NVIDIA and their GPU grant.

## References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. Findings of the 2019 conference on machine translation (WMT19). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. COMET: Commonsense transformers for automatic knowledge graph construction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics.
- Gino Brunner, Yang Liu, Damián Pascual, Oliver Richter, Massimiliano Ciaramita, and Roger Wattenhofer. 2020. On identifiability in transformers. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, Roldano Cattoni, and Marcello Federico. 2015. The IWSLT 2015 evaluation campaign. In *Proceedings of the International Workshop on Spoken Language Translation*.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. 2013. Report on the 10th IWSLT evaluation campaign. In *Proceedings of the International Workshop on Spoken Language Translation, Heidelberg, Germany*.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. 2014. Report on the 11th IWSLT evaluation campaign, IWSLT 2014. In *Proceedings of the International Workshop on Spoken Language Translation*, page 57, Hanoi, Vietnam.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. 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.
- Gonçalo M. Correia, Vlad Niculae, and André F. T. Martins. 2019. Adaptively sparse transformers. 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 2174–2184, Hong Kong, China. Association for Computational Linguistics.
- Hongyi Cui, Shohei Iida, Po-Hsuan Hung, Takehito Utsuro, and Masaaki Nagata. 2019. Mixed multi-head self-attention for neural machine translation. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 206–214, Hong Kong. Association for Computational Linguistics.
- Anna Currey and Kenneth Heafield. 2019. Incorporating source syntax into transformer-based neural machine translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 24–33, Florence, Italy. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Hiroyuki Deguchi, Akihiro Tamura, and Takashi Nishimiyama. 2019. Dependency-based self-attention for transformer NMT. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 239–246, Varna, Bulgaria. INCOMA Ltd.

- 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.
- José A. R. Fonollosa, Noe Casas, and Marta R. Costa-jussà. 2019. [Joint source-target self attention with locality constraints](#). *arXiv preprint arXiv:1905.06596*.
- Sarthak Garg, Stephan Peitz, Udhyakumar Nallasamy, and Matthias Paulik. 2019. [Jointly learning to align and translate with transformer models](#). 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 4453–4462, Hong Kong, China. Association for Computational Linguistics.
- Jie Hao, Xing Wang, Shuming Shi, Jinfeng Zhang, and Zhaopeng Tu. 2019a. [Multi-granularity self-attention for neural machine translation](#). 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 887–897, Hong Kong, China. Association for Computational Linguistics.
- Yaru Hao, Li Dong, Furu Wei, and Ke Xu. 2019b. [Visualizing and understanding the effectiveness 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 4143–4152, Hong Kong, China. Association for Computational Linguistics.
- Sathish Reddy Indurthi, Insoo Chung, and Sangha Kim. 2019. [Look harder: A neural machine translation model with hard attention](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3037–3043, Florence, Italy. Association for Computational Linguistics.
- Ganesh Jawahar, Benoît Sagot, and Djamel Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. [Reformer: The efficient transformer](#). *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada. Association for Computational Linguistics.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. 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 4364–4373, Hong Kong, China. Association for Computational Linguistics.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. 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.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Frederick Liu, Han Lu, and Graham Neubig. 2018. [Handling homographs in neural machine translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1336–1345, New Orleans, Louisiana. Association for Computational Linguistics.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- David Mareček and Rudolf Rosa. 2019. [From balustrades to pierre vinken: Looking for syntax in transformer self-attentions](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 263–275, Florence, Italy. Association for Computational Linguistics.

- Rebecca Marvin and Philipp Koehn. 2018. [Exploring word sense disambiguation abilities of neural machine translation systems \(non-archival extended abstract\)](#). In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, pages 125–131, Boston, MA. Association for Machine Translation in the Americas.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) *Advances in Neural Information Processing Systems*.
- Graham Neubig, Zi-Yi Dou, Junjie Hu, Paul Michel, Danish Pruthi, and Xinyi Wang. 2019. [compare-mt: A tool for holistic comparison of language generation systems](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 35–41, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. [Scaling neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318. Association for Computational Linguistics.
- Jungyeul Park, Jeon-Pyo Hong, and Jeong-Won Cha. 2016. Korean Language Resources for Everyone. In *Proceedings of the 30th Pacific Asia Conference on Language, Information and Computation (PACLIC 30)*, pages 49–58, Seoul, Korea.
- Thuong Hai Pham, Dominik Macháček, and Ondřej Bojar. 2019. Promoting the knowledge of source syntax in transformer nmt is not needed. *Computación y Sistemas*, 23(3).
- Maja Popović and Sheila Castilho. 2019. Challenge test sets for MT evaluation. In *Proceedings of Machine Translation Summit XVII Volume 3: Tutorial Abstracts*, Dublin, Ireland. European Association for Machine Translation.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Xiao Pu, Nikolaos Pappas, James Henderson, and Andrei Popescu-Belis. 2018. [Integrating weakly supervised word sense disambiguation into neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 6:635–649.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#). *OpenAI Blog*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Alessandro Raganato, Yves Scherrer, and Jörg Tiedemann. 2019. [The MuCoW test suite at WMT 2019: Automatically harvested multilingual contrastive word sense disambiguation test sets for machine translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 470–480, Florence, Italy. Association for Computational Linguistics.
- 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.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Annette Rios Gonzales, Laura Mascarell, and Rico Sennrich. 2017. [Improving word sense disambiguation in neural machine translation with sense embeddings](#). In *Proceedings of the Second Conference on Machine Translation*, pages 11–19, Copenhagen, Denmark. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in bertology: What we know about how bert works. *arXiv preprint arXiv:2002.12327*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *5th Workshop on Energy Efficient Machine Learning and Cognitive Computing - NeurIPS*.
- Marten van Schijndel, Aaron Mueller, and Tal Linzen. 2019. [Quantity doesn’t buy quality syntax with neural language models](#). 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 5831–5837, Hong Kong, China. Association for Computational Linguistics.
- Rico Sennrich. 2017. [How grammatical is character-level neural machine translation? assessing MT quality with contrastive translation pairs](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 376–382, Valencia, Spain. Association for Computational Linguistics.

- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich and Biao Zhang. 2019. [Revisiting low-resource neural machine translation: A case study](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 211–221, Florence, Italy. 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.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, and Chengqi Zhang. 2018. Bi-directional block self-attention for fast and memory-efficient sequence modeling. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. [Linguistically-informed self-attention for semantic role labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5027–5038, Brussels, Belgium. Association for Computational Linguistics.
- Sainbayar Sukhbaatar, Edouard Grave, Piotr Bojanowski, and Armand Joulin. 2019. [Adaptive attention span in transformers](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 331–335, Florence, Italy. Association for Computational Linguistics.
- Gongbo Tang, Mathias Müller, Annette Rios, and Rico Sennrich. 2018. [Why self-attention? a targeted evaluation of neural machine translation architectures](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4263–4272, Brussels, Belgium. Association for Computational Linguistics.
- Gongbo Tang, Rico Sennrich, and Joakim Nivre. 2019. [Encoders help you disambiguate word senses in neural machine translation](#). 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 1429–1435, Hong Kong, China. Association for Computational Linguistics.
- Yi Tay, Dara Bahri, Donald Metzler, Da-Cheng Juan, Zhe Zhao, and Che Zheng. 2020. Synthesizer: Rethinking self-attention in transformer models. *arXiv preprint arXiv:2005.00743*.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in OPUS. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Ke Tran, Arianna Bisazza, and Christof Monz. 2018. [The importance of being recurrent for modeling hierarchical structure](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4731–4736, Brussels, Belgium. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Jesse Vig and Yonatan Belinkov. 2019. [Analyzing the structure of attention in a transformer language model](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 63–76, Florence, Italy. Association for Computational Linguistics.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019a. [The bottom-up evolution of representations in the transformer: A study with machine translation and language modeling objectives](#). 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 4395–4405, Hong Kong, China. Association for Computational Linguistics.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. 2019b. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Felix Wu, Angela Fan, Alexei Baevski, Yann N Dauphin, and Michael Auli. 2019. Pay less attention with lightweight and dynamic convolutions. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Mingzhou Xu, Derek F. Wong, Baosong Yang, Yue Zhang, and Lidia S. Chao. 2019. [Leveraging local and global patterns for self-attention networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3069–3075, Florence, Italy. Association for Computational Linguistics.

Baosong Yang, Zhaopeng Tu, Derek F. Wong, Fandong Meng, Lidia S. Chao, and Tong Zhang. 2018. [Modeling localness for self-attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4449–4458, Brussels, Belgium. Association for Computational Linguistics.

Baosong Yang, Longyue Wang, Derek F. Wong, Lidia S. Chao, and Zhaopeng Tu. 2019. [Convolutional self-attention networks](#). 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 4040–4045, Minneapolis, Minnesota. Association for Computational Linguistics.

Weiqiu You, Simeng Sun, and Mohit Iyyer. 2020. [Hard-coded gaussian attention for neural machine translation](#). In *Proceedings of ACL 2020*.