

# Probing Word Translations in the Transformer and Trading Decoder for Encoder Layers

Hongfei Xu<sup>1,2</sup> Josef van Genabith<sup>1,2</sup> Qiuhui Liu<sup>3\*</sup> Deyi Xiong<sup>4</sup>

<sup>1</sup>Saarland University / Saarland, Germany

<sup>2</sup>German Research Center for Artificial Intelligence / Saarland, Germany

<sup>3</sup>China Mobile Online Services / Henan, China

<sup>4</sup>Tianjin University / Tianjin, China

hfxunlp@foxmail.com, Josef.Van\_Genabith@dfki.de,

liuqhano@foxmail.com, dyxiong@tju.edu.cn

## Abstract

Due to its effectiveness and performance, the Transformer translation model has attracted wide attention, most recently in terms of probing-based approaches. Previous work focuses on using or probing source linguistic features in the encoder. To date, the way word translation evolves in Transformer layers has not yet been investigated. Naively, one might assume that encoder layers capture source information while decoder layers translate. In this work, we show that this is not quite the case: translation already happens progressively in encoder layers and even in the input embeddings. More surprisingly, we find that some of the lower decoder layers do not actually do that much decoding. We show all of this in terms of a probing approach where we project representations of the layer analyzed to the final trained and frozen classifier level of the Transformer decoder to measure word translation accuracy. Our findings motivate and explain a Transformer configuration change: if translation already happens in the encoder layers, perhaps we can increase the number of encoder layers, while decreasing the number of decoder layers, boosting decoding speed, without loss in translation quality? Our experiments show that this is indeed the case: we can increase speed by up to a factor 2.3 with small gains in translation quality, while an 18-4 deep encoder configuration boosts translation quality by +1.42 BLEU (En-De) at a speed-up of 1.4.

## 1 Introduction

Neural Machine Translation (NMT) has achieved great success in the last few years. The popular Transformer (Vaswani et al., 2017) model, which outperforms previous RNN/CNN based translation models (Bahdanau et al., 2014; Gehring et al., 2017), is based on multi-layer self-attention networks and can be parallelized effectively.

Recently, a wide range of studies related to the Transformer have been conducted. For example, Bisazza and Tump (2018) perform a fine-grained analysis of how various source-side morphological features are captured at different levels of an NMT encoder. Surprisingly, they do not find any correlation between the accuracy of source morphology encoding and translation quality. Morphological features are only captured in context and only to the extent that they are directly transferable to target words. Voita et al. (2019a) study how information flows across Transformer layers and find that representations differ significantly depending on the objectives (machine translation, standard left-to-right language models and masked language modeling). Tang et al. (2019) find that encoder hidden states outperform word embeddings significantly in word sense disambiguation. However, to the best of our knowledge, to date there is no study about how the Transformer translation model transforms individual source tokens into corresponding target tokens (i.e., word translations), and specifically, which role each Transformer layer plays in word translation, and at which layer a word is translated.

To investigate the roles of Transformer layers in translation, in this paper, we adopt probing approaches (Adi et al., 2017; Hupkes et al., 2018; Conneau et al., 2018) and propose to measure the word translation accuracy of output representations of individual Transformer layers by probing how capable they are at translating words. Probing uses linear classifiers, referred to as “probes”, where a probe can only use the hidden units of a given intermediate layer as discriminating features. Moreover, these probes cannot affect the training phase of a model, and they are generally added after training (Alain and Bengio, 2017). In addition to analyzing the role of each encoder/decoder layer, we also analyze the contribution of the source context and the decoding history in translation by testing the effects of the masked self-attention sub-layer and

\* Corresponding author.

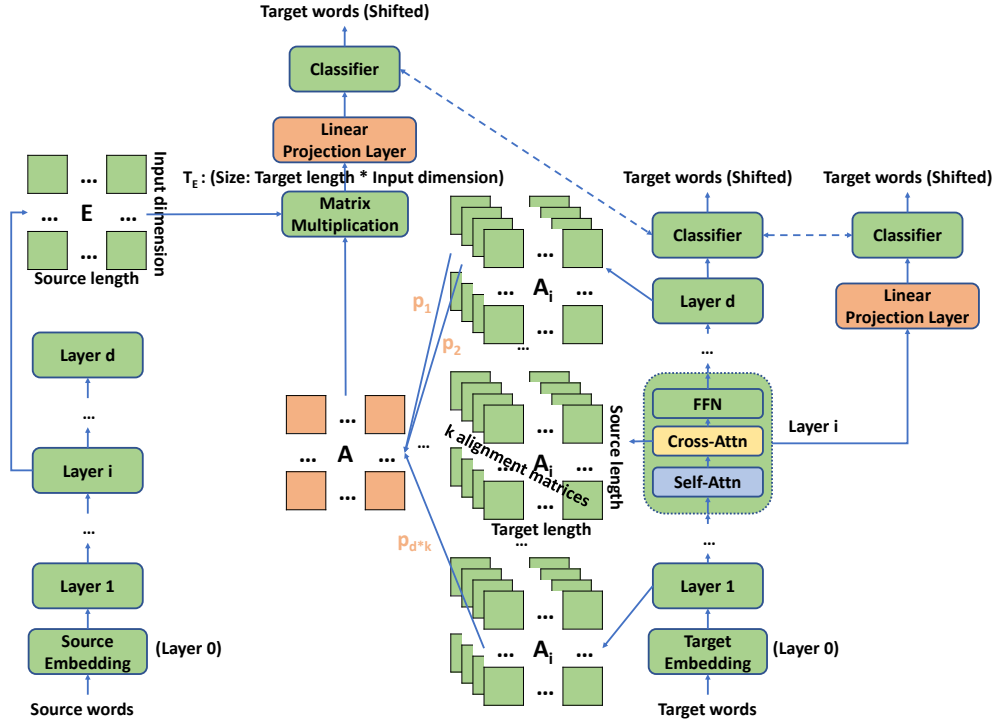


Figure 1: Analyzing word translations of Transformer layers. Green indicates layers of the trained Transformer model frozen for analysis. Orange indicates parameters of the linear projection layer and weights of alignment matrices  $A_i$  trained on the training set. Dashed arrows indicate shared modules. When analyzing the separate effects of source contexts or decoding history in a decoder layer, one of the cross-attention (in yellow) or self-attention sub-layers (in blue) of the analyzed decoder layer are bypassed by a residual connection (Section 2.2). Layers are independently analyzed. Target words (Shifted): the reference translation is one-position right-shifted compared to decoder input, i.e., predicting the next word with the current word as input.

the cross-attention sub-layer in decoder layers.

We present empirical results for how word translation is performed in each encoder/decoder layer, and how the alignment modeling (cross-attention sub-layers) and language modeling (masked self-attention sub-layers) contribute to the performance in each decoder layer. Our analysis demonstrates how word translation evolves across encoder/decoder layers and provides insights into the impact of the source “encoding” and the decoding history on the translation of target tokens. It reveals the existence of target translations in encoder states (and even source word embeddings) and the translation performed by encoder layers.

Based on our findings, we show that the proper use of more encoder layers with fewer decoder layers can significantly boost decoding speed without harming quality. Recently, Kasai et al. (2021) independently and similar to our encoder-decoder layer trading approach, compare the performance and speed of a 12-layer encoder 1-layer decoder with Non-Autoregressive Translation (NAT) approaches, and show that a one-layer autoregressive decoder

can yield state-of-the-art accuracy with comparable latency to strong non-autoregressive models. Our analysis explains why using a deep encoder with a shallow decoder is feasible, and we show that some encoder-decoder depth configurations deliver both increased speed and increased translation quality.

## 2 Probing Layer-wise Word Translation

To analyze word translation accuracy of the Transformer, we first freeze a trained Transformer model so its behavior is consistent in how it performs in translation during our analysis. We then extract output representations of the particular layer analyzed, apply a linear projection layer to extract features related to translation and feed the projected representations to the frozen decoder classifier of the trained Transformer. Our approach is minimally invasive in that only the linear projection layer and the weights of the alignment matrix  $A$  responsible for combining frozen cross-attention alignment matrices from the decoder are trained and updated on the training set, with the original Transformer being frozen. Thus the projection layer will only

transform between vector spaces without generating new features for the word translation, and the alignment matrix  $A$  will only combine frozen cross-attention alignment matrices. A high-level illustration of our analysis approach for encoder/decoder layers is shown in Figure 1.

## 2.1 Analysis of Encoder Layers

Analyzing word translation accuracy of encoder layers requires us to align source tokens with corresponding target tokens. We use the frozen alignment matrices computed by cross-attention sub-layers in decoder layers to align source tokens with target tokens (Figure 1). As there are multiple matrices produced by each sub-layer (due to the multi-head attention mechanism) and multiple decoder layers, we have to ensemble them into one matrix of high alignment accuracy using weights. Assume there are  $d$  decoder layers with  $k$  attention heads in each multi-head attention sub-layer, which results in  $d * k$  alignment matrices  $A_1, \dots, A_{d*k}$ . We use a  $d * k$  dimension weight vector  $w$  to combine all attention matrices. The weight vector is normalized by softmax to a probability distribution  $p$ :

$$p_i = \frac{e^{w_i}}{\sum_{j=1}^{d*k} e^{w_j}} \quad (1)$$

where  $i$  indicates the  $i$ th element in  $w$ .

Then we use  $p$  as the weights of the corresponding attention matrices and merge them into one alignment matrix  $A$ .

$$A = \sum_{i=1}^{d*k} A_i * p_i \quad (2)$$

$w$  is trained with the linear projection layer through backpropagation on the frozen Transformer.

After we obtain the alignment matrix  $A$ , instead of selecting the target token with the highest alignment weight as the translation of a source token, we perform matrix multiplication between the encoded source representations  $E$  (size: source sentence length \* input dimension) and the alignment matrix  $A$  (size: source sentence length \* target sentence length) to transform/re-order source representations to the target side  $T_E$ :

$$T_E = A^T \times E \quad (3)$$

where  $A^T$  and  $\times$  indicate the transpose of  $A$  and matrix multiplication.

Thus  $T_E$  has the same length as the gold translation sequence, and the ground-truth target sequence can be used directly as the translation represented by  $T_E$ .

Though source representations are transformed to the target side, we suggest this does not involve any target side information as the pre-trained Transformer is frozen and the transformation does not introduce any representation from the decoder side. We do not retrieve target tokens with the highest alignment score as word translations of corresponding source tokens because translation may involve zero/one/multiple source token(s) to zero/one/multiple target token(s) alignments, and we suggest that using a soft alignment (attention weights) may lead to more reliable gradients than a hard alignment.

## 2.2 Analysis of Decoder Layers

The analysis of the prediction accuracy of the decoder is simpler than the encoder, as we can directly use the shifted target sequence (teacher forcing) without the requirement to bridge different sequence lengths between the source sentence and the target while analyzing the encoder. We use the output representations of the analyzed layer, and evaluate its prediction accuracy after projection.

However, as studied by Li et al. (2019a), the decoder involves two kinds of “translation”. One (performed by the self-attention sub-layer) translates the history token sequence to the next token, another (performed by the cross-attention sub-layer) translates by attending source tokens. We additionally analyze the effects of these two kinds of translation on predicting accuracy by dropping the corresponding sub-layer (either cross- or masked self-attention) of the analyzed decoder layer (i.e., we only compute the other sub-layer and the feed-forward layer where only the residual connection is kept as the computation of the skipped sub-layer).

## 3 Analysis Experiments

### 3.1 Settings

We first trained a Transformer base model for our analysis on the popular WMT 14 English to German news translation task to compare with Vaswani et al. (2017). We employed a  $512 * 512$  parameter matrix as the linear projection layer. The source embedding matrix, the target embedding matrix and the weight matrix of the classifier were tied. Parameters were initialized under the Lipschitz con-

Layer	Encoder		Decoder					
	Acc	$\Delta$	Acc	$\Delta$	-Self attention		-Cross attention	
	Acc	$\Delta$	Acc	$\Delta$	Acc	$\Delta$	Acc	$\Delta$
0	40.73		13.72					
1	41.85	1.12	20.52	6.80	17.46	-3.06	16.47	-4.05
2	43.75	1.90	26.06	5.54	21.03	-5.03	22.91	-3.15
3	45.49	1.74	34.13	8.07	26.68	-7.45	27.79	-6.34
4	47.14	1.65	55.00	20.87	39.43	-15.57	35.32	-19.68
5	48.35	1.21	66.14	11.14	62.60	-3.54	55.84	-10.30
6	49.22	0.87	70.80	4.66	70.13	-0.67	69.03	-1.77

Table 1: Word translation accuracy of Transformer layers on the WMT 14 En-De task.

Layer	Encoder		Decoder					
	Acc	$\Delta$	Acc	$\Delta$	-Self attention		-Cross attention	
	Acc	$\Delta$	Acc	$\Delta$	Acc	$\Delta$	Acc	$\Delta$
0	41.87		16.26					
1	43.61	1.74	25.73	9.47	23.31	-2.42	18.89	-6.84
2	45.26	1.65	32.55	6.82	27.10	-5.45	26.82	-5.73
3	46.68	1.42	40.80	8.25	34.05	-6.75	32.84	-7.96
4	47.88	1.20	55.60	14.80	47.29	-8.31	40.48	-15.12
5	48.73	0.85	64.39	8.79	62.41	-1.98	55.69	-8.70
6	49.39	0.66	67.10	2.71	66.93	-0.17	66.31	-0.79

Table 2: Word translation accuracy of Transformer layers on the WMT 15 Cs-En task.

straint (Xu et al., 2020) to ensure the convergence of deep encoders. We implemented our approaches based on the Neutron implementation (Xu and Liu, 2019) of the Transformer translation model.

We applied joint Byte-Pair Encoding (BPE) (Sennrich et al., 2016b) with 32k merge operations. We only kept sentences with a maximum of 256 sub-word tokens for training. The concatenation of newstest 2012 and newstest 2013 was used for validation and newstest 2014 as the test set.

The number of warm-up steps was set to 8k.<sup>1</sup> The model was trained for 100k training steps with around 25k target tokens in each batch. We followed all the other settings of Vaswani et al. (2017).

We averaged the last 5 checkpoints saved with an interval of 1,500 training steps. For decoding, we used a beam size of 4, and evaluated tokenized case-sensitive BLEU.<sup>2</sup> The averaged model achieved a

BLEU score of 27.96 on the test set.

The projection matrix and the weight vector  $w$  of 48 elements for alignment were trained on the training set with the frozen Transformer. We monitored the accuracy on the development set, and report results on the test set.

### 3.2 Analysis

The analysis results of the trained Transformer are shown in Table 1. Layer 0 stands for the embedding layer. “Acc” indicates the prediction accuracy. “-Self attention” and “-Cross attention” in the decoder layer analysis mean bypassing the computation of the masked self-attention sub-layer and the cross-attention sub-layer respectively of the analyzed decoder layer using a residual connection. In our layer analysis of the encoder and decoder, “ $\Delta$ ” indicates improvements in word translation accuracy of the analyzed layer over the previous layer. While analyzing the self-attention and cross-attention sub-layers, “ $\Delta$ ” is the accuracy loss when we remove the computation of the corresponding sub-layer.

The results of the encoder layers in Table 1 show

<sup>1</sup><https://github.com/tensorflow/tensor2tensor/blob/v1.15.4/tensor2tensor/models/transformer.py#L1818>.

<sup>2</sup><https://github.com/moses-smt/mosesdecoder/blob/master/scripts/generic/multi-bleu.perl>.

that: 1) encoder layers already perform word translation, and the translation even starts at the embedding layer with unexpectedly high accuracy. 2) With the stacking of encoder layers, the word translation accuracy improves, and improvements brought about by different layers are relatively similar, indicating that all encoder layers are useful.

Surprisingly, analyzing decoder layers, Table 1 shows that: 1) shallow decoder layers (0, 1, 2 and 3) perform significantly worse compared to the corresponding encoder layers (all the way up until the 4th decoder layer, where a word translation accuracy which surpasses the embedding layer of the encoder is achieved); 2) The improvements brought about by different decoder layers are quite different. Specifically, the relative performance increases between the low-performance decoder layers (0, 1, 2 and 3) are low as well, while layers 4 and 5 bring more improvements than the others.

While analyzing the effects of the source context (“-Cross attention” prevents informing translation by the source “encoding”) and the decoding history (the self-attention sub-layer is responsible for the target language re-ordering, and “-Self attention” prevents using the decoding history in the analyzed decoder layer), Table 1 shows that in shallow decoder layers (layer 1-3), the decoding history is as important as the source “encoding”, while in deep decoder layers, the source “encoding” plays a more vital role than the decoding history. Overall, our results provide new insights on the importance of translation already performed by the encoder.

Since the English-German translation shares many sub-words naturally ( $\sim 13.89\%$  source sub-words including punctuations exist in the subword set of the corresponding target translation in the training set), we additionally provide results on the WMT 15 Cs-En task in Table 2. Table 2 confirms our observations reported in Table 1.

Zhang and Bowman (2018); Hewitt and Liang (2019); Voita and Titov (2020) articulate concerns about analyses with probing accuracies, as differences in accuracies fail to reflect differences in representations in several “sanity checks”. Specifically, Zhang and Bowman (2018) compare probing scores for trained models and randomly initialized ones, and observe reasonable differences in the scores only when reducing the amount of classifier training data. However, we argue that in our work, we use the frozen classifier of the pre-trained Transformer decoder as our probing classifier, and the

Layer	BLEU 1	$\Delta$	BLEU	$\Delta$
0	33.1		7.92	
1	35.7	2.6	8.99	1.07
2	41.0	5.3	11.05	2.06
3	43.3	2.3	11.89	0.84
4	46.8	3.5	13.13	1.24
5	48.1	1.3	13.34	0.21
6	48.6	0.5	13.45	0.11
FULL	62.0	13.4	33.26	19.81

Table 3: Translation performance of encoder layers on the WMT 14 En-De task.

introduced linear projection, as well as the alignment matrix  $A$ , are much smaller and weaker than the frozen classifier and the rest of the frozen Transformer components. Thus we suggest that our approach is minimally invasive and that our analysis is less likely to be seriously affected by this issue even though we use a large training set. To empirically verify this, we apply our analysis approach on a randomly initialized encoder and evaluate word translation accuracies obtained by the source embedding layer and last encoder layer, while the alignment between the source and the target is still from the pre-trained model. Both the source embedding layer and the last encoder layer resulted in the same accuracy of 23.66. Compared to the corresponding values (40.73 and 49.22) in Table 1, the gap between the randomly initialized layers and the pre-trained layers in accuracy is significant, and the gap between accuracy improvements from the representation extracted from the source embedding layer and propagated through all intermediate layers to the last encoder layer of pre-trained layers (8.49) and randomly initialized layers (0.00) is also significant. Thus, we suggest our analysis is robust.

### 3.3 Translation from Encoder Layers without Using Decoder Layers

Since our approach extracts features for translation from encoder states while analyzing them, is it possible to perform word translation with only these features from encoder layers without using the decoder except the frozen classifier?

To test this question, we feed output representations from an encoder layer to the corresponding linear projection layer, and feed the output of the linear projection layer directly to the frozen decoder classifier, and retrieve tokens with the highest probabilities as “translations”. Even though

such “translations” from encoder layers have the same length and the same word order as source sentences, individual source tokens are translated to the target language to some extent. We evaluated BPEized<sup>3</sup> case-insensitive BLEU and BLEU 1 (1-gram BLEU, indicates the word translation quality), and results are shown in Table 3. “FULL” is the performance of the whole Transformer model (decoding with a beam size of 4). “ $\Delta$ ” means the improvements obtained by the introduced layer (or the decoder for “FULL”) over the previous layer.

Table 3 shows that while there is a significant gap in BLEU scores between encoder layers and the full Transformer, the gap in BLEU 1 is relatively smaller than in BLEU. It is reasonable that encoder layers achieve a comparably high BLEU 1 score but a low BLEU score overall, as they perform word translation in the same order as the source sentence without any word re-ordering of the target language. We suggest that the BLEU 1 score achieved by only the source embedding layer (i.e., translating with only embeddings) is surprising and worth noting.

### 3.4 Discussion

Our probing approach involves crucial information from the decoder (encoder-decoder attention from all decoder layers). However, we argue that probe training requires supervision. For the decoder, we can directly use gold references. On the encoder side, parallel data does not provide word translations for source tokens, and we have to generate this data by aligning target tokens to source tokens. One choice is extracting alignments by taking an argmax of alignment matrices or using toolkits like fastalign (Dyer et al., 2013). In this case, probe training does not involve attention matrices, but this has drawbacks: multiple/no target tokens may align to one source token. We use soft aggregation to preserve more information (other attention possibilities besides the highest are kept) and to alleviate error propagation. We argue that the use of attention matrices is only to bring supervision (word translations) from the target side to the source side, which is inevitable. Decoder representations cannot flow back to the frozen encoder.

Our paper also empirically reveals the impact of attention matrices: 1) In Section 3.3, where after the training of source probes, we decode target tokens with only encoder layers, the trained probe

<sup>3</sup>Since there is no re-ordering of the target language performed, which makes the merging of translated sub-word units in the source sentence order pointless.

(without involving cross-attention networks) and the pre-trained classifier. 2) In the last paragraph of Section 3.2, we train probes with alignment matrices from the pre-trained model but a frozen random encoder, showing the effects of cross-attention matrices on the probe.

## 4 Trading Decoder for Encoder Layers

### 4.1 Motivation

From our analysis of the 6-layer Transformer base model (Table 1), we find that in contrast to the improvements of the word translation accuracy with increasing depth on the encoder side, some decoder layers contribute significantly fewer improvements than others (i.e., layers 4 and 5 bring more word translation accuracy improvements than those from layers 1, 2, 3 and 6 in Table 1). This suggests that there might be more “lazy” layers in the decoder than in the encoder, which means that it might be easier to compress the decoder than the encoder, and further we conjecture that simply removing some decoder layers while adding the same number of encoder layers may even improve the translation quality of the transformer. Motivations targeting efficiency include:

- Each decoder layer has one more cross-attention sub-layer than an encoder layer, and increasing encoder layers while decreasing the same number of decoder layers will reduce the number of parameters and computational cost;
- During inference, the decoder has to autoregressively compute the forward pass for every decoding step (the decoding of each target token), which prevents efficient parallelization, while encoder layers are non-autoregressively propagated and highly parallelized, and the acceleration caused by using fewer decoder layers with more encoder layers will be more significant in decoding, which is of practical value.

### 4.2 Results and Analysis

We examine the effects of reducing the number of decoder layers while adding corresponding numbers of encoder layers, and results are shown in Table 4. “Speed up” stands for the decoding acceleration compared to the 6-layer Transformer.

Table 4 shows that while the acceleration of trading decoder layers for encoder layers in training is small, in decoding it is significant. Specifically, the

Model	Depth		BLEU	Para. (M)	Train	Time	
	Encoder	Decoder				Decode (/s)	Speed up
Zhang et al. (2018a)	6	6	28.13	74.97	40h09m	29	1.52
Transformer	6	6	27.96	62.37	33h33m	44	1.00
	7	5	28.07	61.32	32h17m	38	1.16
	8	4	28.61	60.27	31h26m	31	1.42
	9	3	28.53	59.22	30h29m	25	1.76
	10	2	28.47	58.17	30h11m	19	2.32
	11	1	27.02	57.12	29h27m	13	3.38
	18	4	29.38	91.77	52h56m	32	1.38

Table 4: Effects of encoder/decoder depth on the WMT 14 En-De task. The decoding time is for the test set of 3,003 sentences with a beam size of 4.

Layer	Encoder		Decoder					
	Acc	$\Delta$	Acc	$\Delta$	-Self attention		-Cross attention	
	Acc	$\Delta$	Acc	$\Delta$	Acc	$\Delta$	Acc	$\Delta$
0	40.48		14.04					
1	41.29	0.81	37.42	23.38	25.56	-11.86	20.40	-17.02
2	43.00	1.71	68.77	31.35	62.01	-6.76	40.67	-28.10
3	44.07	1.07						
4	45.86	1.79						
5	46.54	0.68						
6	47.46	0.92						
7	48.92	1.46						
8	49.58	0.66						
9	50.24	0.66						
10	50.35	0.11						

Table 5: Word accuracy analysis on Transformer with 10 encoder and 2 decoder layers on the WMT 14 En-De task.

Transformer with 10 encoder layers and 2 decoder layers is 2.32 times as fast as the 6-layer Transformer while achieving a slightly higher BLEU.

Can we use more than 12 encoder layers with a shallow decoder to benefit both translation quality and inference speed? Table 4 shows that the 18-4 model<sup>4</sup> brings about +1.42 BLEU improvements over the strong baseline, while being 1.38 times as fast in decoding. Comparing the 18-4 model to the 8-4 model, the time cost for using 10 more encoder layers only increases 1 second for translating the test set, suggesting that autoregressive decoding

speed is quite insensitive to the encoder depth.

Our results show that using more encoder layers with fewer but sufficient decoder layers can significantly boost the decoding speed with small gains in translation quality, and that a good choice in the distribution of encoder and decoder layers (18-4) can result in slightly faster decoding and a substantial increase in translation quality, which is simple but effective and valuable for back-translation (Sennrich et al., 2016a) and production applications.

We present the word accuracy analysis results of the 10 encoder layer - 2 decoder layer Transformer on the En-De task in Table 5. Comparing Table 5 with Table 1, we find that: 1) The differences in improvements (1.71 vs. 0.11) brought by individual layers of the 10-layer encoder are larger than those of the 6-layer encoder (1.90 vs. 0.87), indi-

<sup>4</sup>A full grid search over configurations is tedious and expensive. We take inspiration from Table 4 where going from 5 to 4 decoder layers brings about the biggest relative jump in translation quality. We explored a few configurations and find that using more than 18 encoder layers can still bring improvements, but the gains are relatively small.

Depth		En-De	En-Fr	Cs-En
Encoder	Decoder			
	6	27.96	40.13	28.69
10	2	28.47	40.49	28.87
18	4	<b>29.38<sup>†</sup></b>	<b>40.90<sup>†</sup></b>	<b>29.75<sup>†</sup></b>

Table 6: Verification of deep encoder and shallow decoder on WMT En-De, En-Fr and Cs-En tasks. <sup>†</sup> indicates significance at  $p < 0.01$ .

cating that there might now be some “lazy” layers in the 10-layer encoder; 2) Decreasing the depth of the decoder removes “lazy” decoder layers in the 6-layer decoder and makes decoder layers rely more on the source “encoding” (by comparing the effects of skipping the self-attention sub-layer and cross-attention sub-layer on performance).

### 4.3 Verification of Deep Encoder and Shallow Decoder on other Language Pairs

To investigate how a deep encoder with a shallow decoder will perform in other tasks, we conducted experiments on the WMT 14 English-French and WMT 15 Czech-English news translation tasks in addition to the WMT 14 English-German task. Results on newstest 2014 (En-De/Fr) and 2015 (Cs-En) respectively are shown in Table 6.

Table 6 shows that the 10-2 model consistently achieves higher BLEU scores than the 6-layer model, and the 18-4 model consistently leads to significant improvements in all 3 tasks.

## 5 Related Work

**Analysis of NMT Models.** Belinkov et al. (2020) analyze the representations learned by NMT models at various levels of granularity and evaluate their quality through relevant extrinsic properties. Li et al. (2019a) analyze the word alignment quality in NMT and the effect of alignment errors on translation errors. They demonstrate that NMT captures word alignment much better for those words mostly contributed from the source than those from the target. Voita et al. (2019b) evaluate the contribution of individual attention heads to the overall performance of the model and analyze the roles played by them in the encoder. Yang et al. (2019) propose a word reordering detection task to quantify how well the word order information is learned by Self-Attention Networks and RNN, and reveal that although recurrence structure makes the model more universally effective on learning word order,

learning objectives matter more in the downstream tasks such as machine translation. Tsai et al. (2019) regard attention as applying a kernel smoother over the inputs with the kernel scores being the similarities between inputs, and analyze individual components of the Transformer’s attention with the new formulation via the lens of the kernel. Tang et al. (2019) find that encoder hidden states outperform word embeddings significantly in word sense disambiguation. He et al. (2019) measure the word importance by attributing the NMT output to every input word and reveal that words of certain syntactic categories have higher importance while the categories vary across language pairs. Voita et al. (2019a) use canonical correlation analysis and mutual information estimators to study how information flows across Transformer layers. Early work by Bisazza and Tump (2018) performs a fine-grained analysis of how various source-side morphological features are captured at different levels of the NMT encoder. While they are unable to find any correlation between the accuracy of source morphology encoding and translation quality, they discover that morphological features are only captured in context and only to the extent that they are directly transferable to the target words, and suggest encoder layers are “lazy”. Our analysis offers an explanation for their results as the translation already starts at the source embedding layer, and possibly source embeddings already represent linguistic features of their translations.

**Analysis of BERT.** BERT (Devlin et al., 2019) uses the Transformer encoder, and analysis of BERT may provide valuable references for analyzing the Transformer. Jawahar et al. (2019) provide support that BERT networks capture structural information, and perform a series of experiments to unpack the elements of English language structure learned by BERT. Tenney et al. (2019) employ the edge probing task suite, and find that BERT represents the steps of the traditional NLP pipeline in an interpretable and localizable way, and that the regions responsible for each step appear in the expected sequence: POS tagging, parsing, NER, semantic roles, then coreference. Pires et al. (2019) present a large number of probing experiments, and show that Multilingual-BERT’s robust ability to generalize cross-lingually is underpinned by a multilingual representation.



**Accelerating Decoding.** Zhang et al. (2018a) propose average attention as an alternative to the self-attention network in the Transformer decoder to accelerate decoding. Wu et al. (2019) introduce lightweight convolution and dynamic convolutions. The number of operations required by their approach scales linearly in the input length, whereas self-attention is quadratic. Zhang et al. (2018b) apply cube pruning to neural machine translation to speed up translation. Zhang et al. (2018c) propose to adopt an n-gram suffix-based equivalence function into beam search decoding, which obtains similar translation quality with a smaller beam size, making NMT decoding more efficient. Non-Autoregressive Translation (NAT) (Gu et al., 2018; Libovický and Helcl, 2018; Wei et al., 2019; Shao et al., 2019; Li et al., 2019b; Wang et al., 2019; Guo et al., 2019) enables parallelized decoding, while there is still a significant quality drop compared to traditional autoregressive beam search, our findings on using more encoder layers might also be adapted to NAT. Recently, and independently of our work, Kasai et al. (2021) compare the performance and speed between a 12-layer encoder 1-layer decoder case with NAT approaches, and show that a one-layer autoregressive decoder yields state-of-the-art accuracy with comparable latency to strong non-autoregressive models. Our work explains why using a deep encoder with a shallow decoder is feasible, and we show that substantial increases in decoding speed are possible with small gains in translation quality, and that for some configurations (e.g., 18-4) significant translation quality increases with modest increases in decoding speed are possible.

## 6 Conclusion

We propose approaches for the analysis of word translation accuracy of Transformer layers to investigate how translation is performed. To measure word translation accuracy, our approach trains a linear projection layer that bridges representations from the frozen pre-trained analyzed layer and the frozen pre-trained classifier. While analyzing encoder layers, our approach additionally learns a weight vector to merge multiple attention matrices into one, and transforms the source “encoding” to the target shape by multiplying the merged alignment matrix. Both the linear projection layer and the weight vector are trained on the frozen Transformer. This is minimally invasive, and training the

new parameters does not account for the findings reported. For the analysis of decoder layers, we additionally analyze the effects of the source context and the decoding history in word prediction through bypassing the corresponding cross- and self-attention sub-layers. Our findings motivate and explain the benefits of trading decoder for encoder layers in our approach and that of Kasai et al. (2021).

Our analysis is the first to reveal the existence of target translations performed by encoder layers (including the source embedding layer). We show that increasing encoder depth while removing decoder layers can lead to significant BLEU improvements while boosting the decoding speed.

## Acknowledgements

We thank anonymous reviewers for their insightful comments. Hongfei Xu acknowledges the support of China Scholarship Council ([2018]3101, 201807040056). Josef van Genabith is supported by the German Federal Ministry of Education and Research (BMBF) under funding code 01IW20010 (CORA4NLP). Deyi Xiong is supported by the National Natural Science Foundation of China (Grant No. 61861130364), the Natural Science Foundation of Tianjin (Grant No. 19JCZDJC31400) and the Royal Society (London) (NAF\R1\180122).

## References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. [Fine-grained analysis of sentence embeddings using auxiliary prediction tasks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Guillaume Alain and Yoshua Bengio. 2017. [Understanding intermediate layers using linear classifier probes](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *arXiv preprint arXiv:1409.0473*.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2020. [On the linguistic representational power of neural machine translation models](#). *Computational Linguistics*, 46(1):1–52.

- Arianna Bisazza and Clara Tump. 2018. [The lazy encoder: A fine-grained analysis of the role of morphology in neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2871–2876, Brussels, Belgium. Association for Computational Linguistics.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. [What you can cram into a single  \$\&\#\&^\*\$  vector: Probing sentence embeddings for linguistic properties](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. [A simple, fast, and effective reparameterization of IBM model 2](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1243–1252, International Convention Centre, Sydney, Australia. PMLR.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor O.K. Li, and Richard Socher. 2018. [Non-autoregressive neural machine translation](#). In *International Conference on Learning Representations*.
- Junliang Guo, Xu Tan, Di He, Tao Qin, Linli Xu, and Tie-Yan Liu. 2019. [Non-autoregressive neural machine translation with enhanced decoder input](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 3723–3730. AAAI Press.
- Shilin He, Zhaopeng Tu, Xing Wang, Longyue Wang, Michael Lyu, and Shuming Shi. 2019. [Towards understanding neural machine translation with word importance](#). 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 953–962, Hong Kong, China. Association for Computational Linguistics.
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). 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 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- Dieuwke Hupkes, Sara Veldhoen, and Willem Zuidema. 2018. Visualisation and ‘diagnostic classifiers’ reveal how recurrent and recursive neural networks process hierarchical structure. *J. Artif. Int. Res.*, 61(1):907–926.
- 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.
- Jungo Kasai, Nikolaos Pappas, Hao Peng, James Cross, and Noah Smith. 2021. [Deep encoder, shallow decoder: Reevaluating non-autoregressive machine translation](#). In *International Conference on Learning Representations*.
- Xintong Li, Guanlin Li, Lemao Liu, Max Meng, and Shuming Shi. 2019a. [On the word alignment from neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1293–1303, Florence, Italy. Association for Computational Linguistics.
- Zhuohan Li, Zi Lin, Di He, Fei Tian, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2019b. [Hint-based training for non-autoregressive 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 5708–5713, Hong Kong, China. Association for Computational Linguistics.
- Jindřich Libovický and Jindřich Helcl. 2018. [End-to-end non-autoregressive neural machine translation with connectionist temporal classification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3016–3021, Brussels, Belgium. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.

- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [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.
- Chenze Shao, Yang Feng, Jinchao Zhang, Fandong Meng, Xilin Chen, and Jie Zhou. 2019. [Retrieving sequential information for non-autoregressive neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3013–3024, Florence, Italy. 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.
- 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.
- Yao-Hung Hubert Tsai, Shaojie Bai, Makoto Yamada, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. [Transformer dissection: An unified understanding for transformer’s attention via the lens of kernel](#). 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 4344–4353, Hong Kong, China. 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 30*, pages 5998–6008. Curran Associates, Inc.
- 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 4396–4406, 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.
- Elena Voita and Ivan Titov. 2020. [Information-theoretic probing with minimum description length](#). *arXiv preprint arXiv:2003.12298*.
- Yiren Wang, Fei Tian, Di He, Tao Qin, ChengXiang Zhai, and Tie-Yan Liu. 2019. [Non-autoregressive machine translation with auxiliary regularization](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 5377–5384. AAAI Press.
- Bingzhen Wei, Mingxuan Wang, Hao Zhou, Junyang Lin, and Xu Sun. 2019. [Imitation learning for non-autoregressive neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1304–1312, Florence, Italy. Association for Computational Linguistics.
- Felix Wu, Angela Fan, Alexei Baevski, Yann Dauphin, and Michael Auli. 2019. [Pay less attention with lightweight and dynamic convolutions](#). In *International Conference on Learning Representations*.
- Hongfei Xu and Qiuhui Liu. 2019. [Neutron: An Implementation of the Transformer Translation Model and its Variants](#). *arXiv preprint arXiv:1903.07402*.
- Hongfei Xu, Qiuhui Liu, Josef van Genabith, Deyi Xiong, and Jingyi Zhang. 2020. [Lipschitz constrained parameter initialization for deep transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 397–402, Online. Association for Computational Linguistics.
- Baosong Yang, Longyue Wang, Derek F. Wong, Lidia S. Chao, and Zhaopeng Tu. 2019. [Assessing the ability of self-attention networks to learn word order](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3635–3644, Florence, Italy. Association for Computational Linguistics.
- Biao Zhang, Deyi Xiong, and Jinsong Su. 2018a. [Accelerating neural transformer via an average attention network](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1789–1798, Melbourne, Australia. Association for Computational Linguistics.

Kelly Zhang and Samuel Bowman. 2018. [Language modeling teaches you more than translation does: Lessons learned through auxiliary syntactic task analysis](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 359–361, Brussels, Belgium. Association for Computational Linguistics.

Wen Zhang, Liang Huang, Yang Feng, Lei Shen, and Qun Liu. 2018b. [Speeding up neural machine translation decoding by cube pruning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4284–4294, Brussels, Belgium. Association for Computational Linguistics.

Zhisong Zhang, Rui Wang, Masao Utiyama, Eiichiro Sumita, and Hai Zhao. 2018c. [Exploring recombination for efficient decoding of neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4785–4790, Brussels, Belgium. Association for Computational Linguistics.