
Learning Multiscale Transformer Models for Sequence Generation

Bei Li^{*1} Tong Zheng^{*1} Yi Jing^{*1} Chengbo Jiao² Tong Xiao^{1,2} Jingbo Zhu^{1,2}

Abstract

Multiscale feature hierarchies have been witnessed the success in the computer vision area. This further motivates researchers to design multiscale Transformer for natural language processing, mostly based on the self-attention mechanism. For example, restricting the receptive field across heads or extracting local fine-grained features via convolutions. However, most of existing works directly modeled local features but ignored the word-boundary information. This results in redundant and ambiguous attention distributions, which lacks of interpretability. In this work, we define those scales in different linguistic units, including sub-words, words and phrases. We built a multiscale Transformer model by establishing relationships among scales based on word-boundary information and phrase-level prior knowledge. The proposed **Universal MultiScale Transformer**, namely UMST, was evaluated on two sequence generation tasks. Notably, it yielded consistent performance gains over the strong baseline on several test sets without sacrificing the efficiency.

1. Introduction

Transformers (Vaswani et al., 2017) have achieved remarkable success on a wide range of tasks in natural language processing (NLP) (Devlin et al., 2019; Dai et al., 2019; Wang et al., 2019; Dehghani et al., 2019). Given an input sequence, it uses self-attention mechanisms to establish global interactions over all positions without recurrences or convolutions.

Despite great potential on most of NLP tasks, the Transformer backbones still have a major shortcoming that it does not model local fine-grained features explicitly. For

^{*}Equal contribution ¹School of Computer Science and Engineering, Northeastern University, Shenyang, China ²NiuTrans Research, Shenyang, China. Correspondence to: Tong Xiao <xiaotong@mail.neu.edu.cn>.

example, given a word in a sentence, the surrounding words have more potential correlations with it than distant words. Also, replacing the self-attention mechanism by dynamic convolutions even shows competitive favorable results with Transformer (Wu et al., 2019). This observation motivates researchers to combine local and global features for a multiscale representation of features.

Along this thread of research, researchers attempt to incorporate local-patterns into the Transformer model. These works could be broadly categorized into two aspects. One is to enhance the self-attention mechanism via varying the receptive field across heads, such as convolutional self-attention (Yang et al., 2019), multiscale self-attention (Guo et al., 2020), multi-granularity self-attention (Hao et al., 2019). The other is placing the convolution networks in sequence (Gulati et al., 2020) or parallelly (Zhao et al., 2019) in self-attention. The core idea is to endow the Transformer to learn multiscale features rather than the single global interaction.

Multiscale feature hierarchy is one of the mainstream paradigms to handle the high-resolution tasks in CV (Lin et al., 2017). More recently, Transformer has been successfully applied to the computer vision area, such as image classification (ViT) (Dosovitskiy et al., 2021), object-detection (DETR) (Carion et al., 2020). In particular, multiscale vision Transformers has received great attention due to its superior performance and computation efficiency (Wang et al., 2021; Fan et al., 2021a; Li et al., 2021b). The input of ViT is a sequence of patches partitioned from an image, which is similar to the sequential input in NLP systems. But these systems follow a multi-stage paradigm where the channel is increased as the resolution is decreased.

In this work, we revisit the design of multiscale Transformer based for sequence generation. Although the aforementioned Transformer variants have made great efforts to boost the models' performance, most of existing works ignore the word-boundaries when modeling the local features. This is because the input sequence is segmented into a series of sub-words (Sennrich et al., 2016; Kudo & Richardson, 2018) to prevent these systems from the out-of-vocabulary problem. However, most of previous techniques, e.g. convolutional self-attention (Yang et al., 2019), multiscale self-attention (Guo et al., 2020) directly enhance the feature extraction of neighboring elements but simply discard those

language units that naturally exist. This leads to redundant and ambiguous attention map as the relationship between sub-words and words is lack of interpretability.

Our main contributions could be summarized as follows:

- We re-define the concept of scale for NLP, including scales of sub-word, word and phrase. Our intention is to leverage the word boundaries and phrase-level prior knowledge to compensate for the sub-word features.
- Given a sequence, each token is regarded as an individual. The correlations among the three scales are described in three-fold: inter-individual, intra-group and inter-group. We show that learning multiscale features is helpful.
- We propose an **Universal MultiScale Transformer**, namely UMST - a universal approach to extracting features from different scales. Also, UMST is flexible and provides an opportunity to incorporate other prior knowledge.

We comprehensively evaluate the proposed UMST model on the commonly used sequence generation tasks, including machine translation and abstractive summarization. Experimental results show that UMST achieve favorable results than the Transformer and recently multiscale Transformer variants without sacrificing the efficiency. Notably, it yields 0.88 and 0.44 BLEU point gains compared with Transformer under the base and big configurations. It significantly outperforms the Transformer over 1 Rouge point gain in average. UMST is orthogonal to the low-level modeling techniques, e.g. it gains an additional improvement of 0.4 BLEU points when working with relative positional representations (Shaw et al., 2018).

2. Preliminary

Transformer (Vaswani et al., 2017) follows the encoder-decoder paradigm (Sutskever et al., 2014). There are six identical blocks with individual initialization on both the encoder and decoder sides. A standard encoder block consists of two operations. The first is a multi-head self-attention operation for modeling the correlations over all positions. The other is a two-layer feed-forward network for modeling relations within an element in a point-wise fashion. Inter-twining these operations with the residual connection (He et al., 2016) and layer normalization (Lei Ba et al., 2016) allows transformers to be optimized stably.

Here, we take the Pre-Norm Transformer (Wang et al., 2019) as an instance. Let x and y be the input and output. A Pre-Norm Transformer block can be described as: $y = \bar{x} + \text{FFN}(\text{LN}(\bar{x}))$ and $\bar{x} = x + \text{SAN}(\text{LN}(x))$. SAN and FFN denote the multi-head attention sub-layer and the

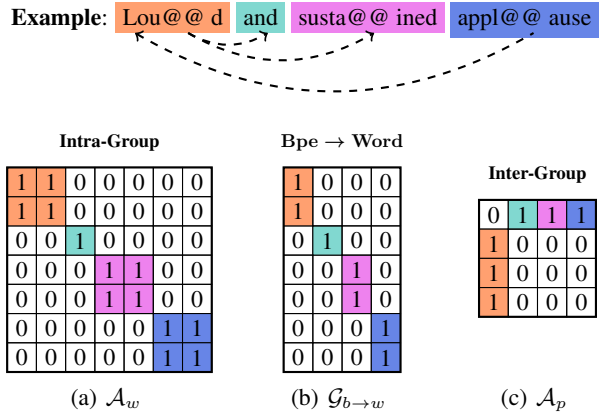


Figure 1. A detailed example of three matrices, which portrays three corresponding relationships among scales. Blocks in same color consist of a word who has been segmented into sub-words.

feed-forward sub-layer, respectively. LN denotes the layer normalization function, where $\text{LN}(x) = \frac{x - \mu}{\delta} \circ \gamma + \beta$. μ and δ are the mean and standard deviation of a feature, and \circ denotes the element-wise dot.

Given the input sequence $x \in \mathbb{R}^{L \times d}$, where L denotes the sequence length and d denotes the hidden size. SAN first obtains the transformation matrix Q, K, V via three sets of projections, where $Q = xW_q, K = xW_k$ and $V = xW_v$. W_q, W_k and W_v are trainable parameters. Then, SAN models the global correlations among all positions via $\text{SAN} = \text{ATTN} \cdot V$, where ATTN can be obtained by:

$$\text{ATTN} = \text{Softmax}\left(\frac{Q \cdot K^T}{\sqrt{d_k}}\right) \quad (1)$$

For simplification, we take the number of heads to 1 as an instance. Then FFN is formalized as: $\text{FFN} = \max(\bar{x}W_1 + b_1, 0)W_2 + b_2$, where W_1 and W_2 are transformation matrices and b_1 and b_2 are bias.

2.1. Definition of Scales and Relations

We start with the definition of linguistic-partitioned scale in this work. In natural language processing, it is natural to see a text in multiple levels, e.g. sub-words, words, phrases, sentences, paragraphs. We adopt this idea to define our linguistic-partitioned scales. In this work, we take sub-words, words and phrases as the basis of these scales. The sub-words are the lowest-level scale while the phrases are the highest-level scale.

We then define the relations among different scales. Take the byte-pair encoding (BPE) (Sennrich et al., 2016) as an instance. A BPE-partitioned sequence composed of multilevel contents, e.g. sub-words, words, and phrases. We regard a sub-word as an individual, and words as a group which composed of the corresponding sub-words. Note that

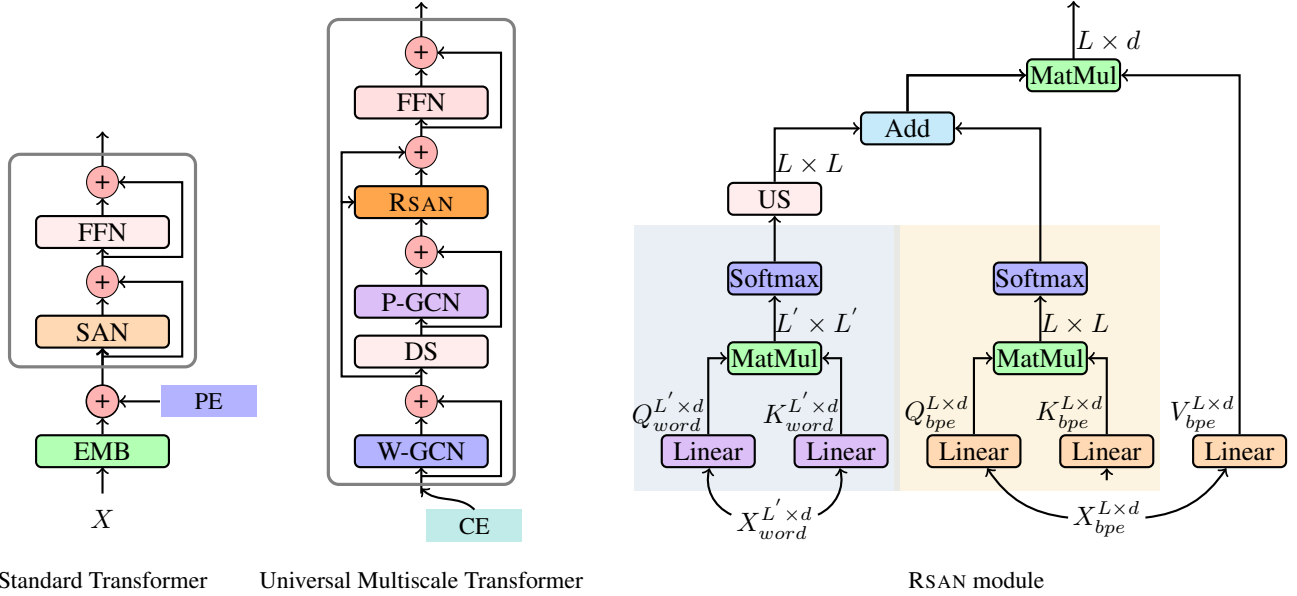


Figure 2. Model architectures of the standard Transformer and the proposed UMST. Note that the DS and US modules denote the down-sampling and up-sampling operations. PE and CE represent positional encoding and class embedding. Here, we omit the layer normalization for simplification. Actually, we follow the pre-normalization strategy as its stability.

some words share both identities mentioned above. Thus, we have the following relations:

- **Inter-individual relation.** It describes the relationship among input tokens, consisting of sub-words and words. This is the global relation in standard Transformer.
- **Intra-group relation.** It is a local fine-grained relation which endows sub-words to obtain the corresponding word boundary information.
- **Inter-group relation.** It is obtained by the dependency parsing which describes the relationship among groups.

Our intention is to use the intra-group and inter-group relations to compensate for the inter-individual interactions. To facilitate the information flow between them, we define three transformation matrices, \mathcal{A}_w , \mathcal{A}_p and $\mathcal{G}_{b \rightarrow w}$, where \mathcal{A}_w and \mathcal{A}_p denote the adjacency matrices of the individual and group, respectively. $\mathcal{G}_{b \rightarrow w}$ is a transformation matrix that indicates the mapping between the word (group) and its corresponding sub-words (individual). Figure 1 gives an example. Given a sequence $x \in \mathbb{R}^{L \times d}$, we define three matrices for describing the scales.

\mathcal{A}_w is the adjacency matrix in $\mathbb{R}^{L \times L}$ which satisfies: $\mathcal{A}_w^{i,j} = 1$ if and only if the sub-tokens i and j belong to the same word. \mathcal{A}_w is a block diagonal matrix where entries in the same color denote a group of sub-tokens of a word.

\mathcal{A}_p is an adjacency matrix in $\mathbb{R}^{L' \times L'}$. It satisfies: $\mathcal{A}_p^{i,j} = 1$ if only if the tokens i and j have a dependency edge. As is seen in Figure 1(c), dashed lines denote the dependency between two words, which can be easily obtained through open-source parsing tools, e.g. StanfordNLP (Qi et al., 2018), Berkeley Neural Parser (Kitaev & Klein, 2018). Notably, we build \mathcal{A}_p as an undirected graph. It is a symmetry matrix where corresponding elements are shaded with different colors. Similar with \mathcal{A}_w , it portrays the relation among words rather than sub-words.

$\mathcal{G}_{b \rightarrow w}$ is a matrix in $\mathbb{R}^{L \times L'}$. Its function is to solve the spatial size mismatch problem between word level and sub-word level. $\mathcal{G}_{b \rightarrow w}$ satisfies: $\mathcal{G}_{b \rightarrow w}^{i,j} = 1$ if only if the sub-token i is belonging to token j . $\mathcal{G}_{b \rightarrow w}$ is an orthogonal matrix of column vectors which guarantees the lossless rank of the matrix.

3. Universal Multiscale Transformer

As noted in § 2, SAN collects the interaction information over all positions, no matter if a token is a word or a sub-word. This results in ineffective interactions and ambiguous attention map. To remedy this issue, we redesign the Transformer architecture from the multiscale perspective. The overall architecture is depicted in Figure 2. As we can see, it slightly differs from the standard Transformer (Figure 2 (left)) since we employ two additional graph convolutional networks before computing the attention. The core of this

Table 1. Several operations to handle the local-pattern.

Operations	Efficiency	Expression	Stable
Average Pooling	✓		✓
GAT (Velickovic et al., 2018)		✓	
GCN (Kipf & Welling, 2017)	✓	✓	✓

design is to coverage the intra-group and inter-group interactions via W-GCN and P-GCN, respectively. Also, we replace the standard self-attention network by the proposed rectified self-attention mechanism.

3.1. Class Embedding

The standard Transformer receives a sequence of embedding as input. The embedding is a mixture of token embedding and position embedding. In order to distinguish which part of position implies sub-tokens, we introduce a binary class embedding, where BPE-level tokens have a value of 1, and word-level tokens have a value of 0. The embedding is randomly initialized by a normal distribution where the mean is 0 and the standard deviation is $1/\sqrt{d}$.

3.2. Intra-Group Interaction

Previous studies ignore the word-boundary information and directly modeling the local feature in a fixed window, which results in degeneration of expression and low interpretability. To tackle these problems, we firstly model the intra-group correlations where sub-words belonging to the same word have a high similarity among them.

Given the word boundary \mathcal{A}_w , there are several ways to gain the fusion representation, such as average pooling, graph attention networks (GAT) (Velickovic et al., 2018), and graph convolutional networks (GCN) (Kipf & Welling, 2017). Table 1 summarizes the properties of three operations considering the expression power, training stability and computation efficiency. In our preliminary experiments, the average pooling method is light and stable while having a poor expression ability. On the contrary, graph attention networks show considerable power to capture the intra-group dependency. However, stacking several attention-based modules leads to poor stability during training, especially when our model becomes deep. Finally, we choose GCN as the default model by taking these factors into account.

A GCN layer encapsulates each node’s hidden representation by aggregating feature information from its neighbors. Our intention is to let each sub-words receive its corresponding word-level information. We assume a sequence $x \in \mathbb{R}^{L \times d}$. Instead of feeding the input into SAN, we use a GCN (Kipf & Welling, 2017) to enhance the feature extraction within each group. The formulation of GCN can be




Figure 3. An example of the attention map transformation.

described as:

$$\text{GCN}_{\text{word}} = \sigma(\tilde{D}_w^{-\frac{1}{2}} \tilde{\mathcal{A}}_w \tilde{D}_w^{-\frac{1}{2}} \cdot xW_w) \quad (2)$$

$\tilde{\mathcal{A}}_w = \mathcal{A}_w + \mathcal{I}_L$ the adjacency matrix of the undirected graph with self-connections. Here \mathcal{I}_L denotes the identity matrix. \tilde{D}_w is the degree matrix of the adjacency matrix $\tilde{\mathcal{A}}_w$. W_w is a linear transformation which is a trainable parameter. σ denotes an activation function. We choose ReLU for the design of activation function: $\text{ReLU} = \max(0, \cdot)$. To further strengthen the fine-grained feature, we multiply the output of GCN_{word} by $\tilde{D}_w^{-\frac{1}{2}} \tilde{\mathcal{A}}_w \tilde{D}_w^{-\frac{1}{2}}$ at the left side.

Essentially, the output GCN_{word} still has a length L which is not changed. In this way, each sub-words not only keeps the original meaning, but also derives the corresponding word boundaries.

3.3. Inter-Group Interaction

Beyond the intra-group interaction, another conceive here is to make full use of the phrase-level prior knowledge. The dependency among words portrays a long-range relationship rather a local constraint within a varying window size. It is a kind of phrase-level prior which could be easily obtained through open-source tools. Notably, the phrase prior is built upon words rather than sub-words. Thus we need to down-sample the representation $X_{\text{word}}^{L' \times d}$ from the length L into L' . Here, we choose the averaging pooling operation to aggregate into the word-level embedding. Subsequently, we model the inter-group interaction via the adjacency matrix \mathcal{A}_p , where the formation is as follows:

$$\text{GCN}_{\text{phrase}} = \sigma(\tilde{D}_p^{-\frac{1}{2}} \tilde{\mathcal{A}}_p \tilde{D}_p^{-\frac{1}{2}} \cdot xW_p) \quad (3)$$

where \tilde{D}_p is the degree matrix of the adjacency matrix $\tilde{\mathcal{A}}_p$. W_p is a learnable linear projection. One can share the projection matrices W_w and W_p across all transformer blocks, which reduces the total parameters in a certain extent. Another choice is to use a shared projection matrix when handling the word-level GCN and the phrase-level GCN, but it slightly hurts the performance. All these discussions could be found in § 6.

3.4. Rectified Self-attention

The proposed RSAN module is a rectified multi-head self-attention via a double-branch manner. Formally, there are two kinds of input. The first input $X_{bpe}^{L \times d}$ contains the word-level interaction, and the other input $X_{word}^{L' \times d}$ is further enhanced by the phrase-level interaction. As depicted in Figure 2, the right branch (yellow) is similar with the original self-attention which computes the correlations over all positions. The attention map $\text{ATTN}_1 \in \mathbb{R}^{L \times L}$ could be easily obtained through Eq. 1, where the query and key are generated from $X_{bpe}^{L \times d}$. The left branch (blue) aims to extract higher-level inter-group interaction over all words. Likewise, we can obtain the attention map $\text{ATTN}_2 \in \mathbb{R}^{L' \times L'}$.

After obtaining two distributions, a next step is converting the attention map from $L' \times L'$ into $L \times L$ using the following transformation:

$$\text{ATTN}'_2 = \mathcal{G}_{b \rightarrow w} \cdot \text{ATTN}_2 \cdot (\mathcal{G}_{b \rightarrow w})^T \quad (4)$$

where $\mathcal{G}_{b \rightarrow w}$ is the transformation mapping. The data follow is shown in Figure 3. The attention map ATTN_1 portrays an inter-individual correlation whose $Q_{bpe}^{L \times d}$ and $K_{bpe}^{L \times d}$ have already been augmented via the injected word-boundaries. Meanwhile, the attention map ATTN_2 captures a high-level interaction among each word rather than the mixture of sub-words and words. Our intention here is to use ATTN_2 to enhance ATTN_1 .

There are several approaches to attaining the final attention map, e.g. sum of the two distributions, average pooling or gated mechanism. However, there is no significant performance gap in our preliminary experiments. Here, we choose the average pooling to guarantee the normalization.

The final output of the proposed rectified self-attention is the multiplication of the fused attention map ATTN'_2 with value $V_{bpe}^{L \times d}$. In this way, RSAN incorporate both the word-boundaries and the phrase-level prior knowledge.

4. Experimental Setups

We evaluate the proposed multiscale Transformer on two sequence generation benchmarks, including machine translation and abstractive summarization.

4.1. Dataset

Machine Translation We report results on two machine translation datasets, including a large-scale WMT’14 English-German (En-De) dataset, and a WMT’16 English-Romanian (En-Ro) dataset. For the En-De dataset, the training data consisted of approximately 4.5M tokenized sentence pairs, as in (Vaswani et al., 2017). We preprocess the dataset following the same setup in Ott et al. (2018)’s work,

Table 2. Comparison with previous studies on the WMT En-De task. The results with † denote our re-implementing results, since some of works have not been evaluated on MT.

Model	Base		Big	
	Param	BLEU	Param	BLEU
Transformer (Vaswani et al., 2017)	65M	27.30	213M	28.40
Scaling NMT (Ott et al., 2018)	-	-	210M	29.30
DLCL (Wang et al., 2019)	62M	27.30	-	-
MUSE (Zhao et al., 2019)	-	-	-	29.90
MG-SA (Hao et al., 2019)	89M	28.28	272M	29.01
Transformer †	65M	27.63	216M	29.31
MUSE† (Zhao et al., 2019)	68M	27.97	233M	29.11
MSMSA† (Guo et al., 2020)	65M	27.57	233M	28.84
TNT† (Han et al., 2021)	83M	28.48	-	-
UMST	70M	28.51	242M	29.75
UMST + RPR	70M	28.90	242M	30.15

obtaining a clean and high-quality bilingual training data containing 3, 894, 376 sentences. We selected *newstest2013* as the validation data and *newstest2014* as the test data. To validate the proposed method on various granularities, we segmented the sentences into sub-word units (Sennrich et al., 2016) with 7K (Xu et al., 2021), 15K and 32K merging operations. For the En-Ro dataset, it consisted of 610K bilingual sentence pairs. Also, we adopted the same scripts as those used in Lee et al. (2018); Kasai et al. (2020)’s work. We adopt a joint source and target BPE factorization with the vocabulary size of 40K. We use newsdev-2016 and newstest-2016 as the validation and test sets, respectively.

Abstractive Summarization We also test the models’s ability to process long sequences on the CNN-DailyMail summarization task (Nallapati et al., 2016; Hermann et al., 2015). The preprocessed method is the same as in (Ott et al., 2019)’s work. We use a shared BPE with 10K operations, resulting in a vocabulary of 10, 420 entries.

4.2. Implementation Details

As discussed in § 3.3, here we choose the open-source parsing tool proposed by Stanford ¹ to extract the dependency tree. Due to the page limit, more experimental setups and model configurations could be found in Appendix A.

5. Results

5.1. Machine Translation

Results on WMT En-De Table 2 compares the proposed UMST with several strongly related systems. The results are evaluated on WMT’14 En-De newstest2014 test set.

¹<https://stanfordnlp.github.io/CoreNLP/>

Table 3. Results on the WMT En-Ro task.

Model	Param	BLEU
DELIGHT (Mehta et al., 2020)	53M	34.70
Baseline in MBART (Liu et al., 2020b)	-	34.30
Baseline in DISCO (Kasai et al., 2020)	-	34.16
Transformer † (Vaswani et al., 2017)	54M	34.21
TNT† (Han et al., 2021)	73M	34.00
UMST	60M	34.81
UMST + RPR	60M	35.31

Table 4. Results of the proposed UMST and previous studies on the CNN-DailyMail dataset.

Model	RG-1	RG-2	RG-L
DYNAMICCONV (Wu et al., 2019)	39.84	16.25	36.73
BOTTOM-UP (Gehrmann et al., 2018)	41.22	18.68	38.34
SURFACE (Liu et al., 2020a)	41.00	18.30	37.90
DMAN (Fan et al., 2021c)	40.98	18.29	37.88
Transformer†	40.55	17.81	37.47
UMST w/o inter-group interactions	41.62	18.65	38.28
UMST	41.82	18.91	38.54

We re-implement these methods in our codebase within the same setting for fair comparisons. First, we can see that the standard Transformer delivers a BLEU point of 27.63 and 29.31 under the base/big configurations. Our UMST outperforms it by 0.88 and 0.44 BLEU points, indicating that incorporating word-boundaries and phrase-level prior knowledge is indeed helpful. Also, UMST is orthogonal to previous local modeling methods (Yang et al., 2019; Shaw et al., 2018; Zhao et al., 2019). Here, we take RPR (Shaw et al., 2018) as an instance. When applying RPR into RSAN, our model can obtain an additional improvement of 0.4 BLEU points. Notably, UMST can beat previous multiscale variants under the same experimental setting. For example, it can beat TNT (Han et al., 2021) whose parameters equals to a 12-layer Transformer. This indicates that redefining scales from the linguistic unit perspective is valuable. More details and comparisons could be found in Appendix B.

Results on WMT En-Ro Similarly, UMST yileds a 34.81 BLEU point on the WMT En-Ro task (see Table 3), outperforming the previous results in terms of BLEU. Also, the RPR-enhanced UMST almost achieves the state-of-art with no use of pre-training models, it outperforms UMST by 0.5 BLEU points. We will further validate the combination with other local modeling methods in future work.

5.2. Abstractive Summarization

The results on abstractive summarization are listed in Table 4. Note that we choose the base configuration, and all sys-

Table 5. The ablation study on the WMT En-De testset.

Model	Depth	BLEU	Depth	BLEU
Transformer	6-6	27.63	12-6	28.67
UMST	6-6	28.51	12-6	29.49
w/o class-embedding	6-6	28.39	12-6	28.99
w/o intra-group interactions	6-6	27.87	12-6	failed
w/o inter-group interactions	6-6	28.06	12-6	29.37
replace GCN with pooling	6-6	27.96	12-6	28.89
replace GCN with GAT	6-6	28.11	12-6	failed

Table 6. Comparisons of shared parameters and several mappings on the WMT En-De task.

#	Model	Param	BLEU
0	Transformer	65M	27.63
1	UMST	70M	28.51
2	shared W_w and W_p (across blocks)	66M	28.33
3	shared W_w and W_p (within each block)	69M	28.39
4	shared Q, K in RSAN	67M	28.41
5	replace by nearest interpolation	70M	27.96
6	replace by linear interpolation	70M	27.61

tems consist of 6 blocks for both encoder and decoder. We can see that the proposed UMST outperforms the standard Transformer by a large margin (e.g. 1.27 Rouge-1 benefits) as the the decoupled multiscale attention receives benefits from both the intra-group (word boundaries) and inter-group (phrase prior) interactions. Note that the model can still attain nearly 1 rouge gains in terms of three metrics when removing the phrase-level prior knowledge. This observation again demonstrates the essential of word-boundaries.

6. Analysis

In this section, we present a detailed analysis to provide some insights on why UMST improves over Transformer.

Ablation studies We conduct a series of ablation studies to inspect UMST from different aspects. Table 5 summarizes the results of depth 6-6 and 12-6 on the WMT’14 En-De task. First, as discussed in § 3.1, the proposed class embedding is used to distinguish whether the current token is a word or a sub-word. Through the results, we find it plays a quite important role when the model goes deep. There is a 0.5 BLEU point drop when removing it.

In this work, both intra-group and inter-group interaction could be regarded as two kinds of local-pattern modeling. The former is to incorporate the word-boundaries, which acts within a continuous restricted window (inside a word). While the latter is a sparse and long-range dependency, containing high-level semantic information. Through the results




Figure 4. The comparison of BLEU against different encoder depths and BPE merging operations.

Table 7. Discussion on training and inference efficiency. The experiments are conducted based on Transformer-Base.

Model	Inference	Training	GPU
Transformer	142.33	1.00×	6.4G
UMST w/o word	147.60	1.18×	6.5G
UMST	136.67	1.25×	6.7G

we find that removing either of them causes BLEU degradations, especially for the intra-group interaction. It results in unstable training when optimizing a 12-layer system. This is because that inter-group interaction is established upon words, rather than sub-words. So UMST w/o intra-group interactions ignoring the word-boundary information brings noise signals to the training and even destroys the training stability. Moreover, inter-group interaction is more important for the shallow model. It causes a 0.43 BLEU drop when removing the inter-group interaction. However, only a 0.12 BLEU drop could be observed for deeper models. This is because deep models may gradually gain high-level semantic features through the stacked self-attention.

We then replace GCN with the average pooling and GAT. We see that replacing GCN with both alternatives suffer from a sharp decrease in performance. The situation becomes more serious for GAT when switching to a deeper scenario, that it even fails to converge. One reasonable explanation is that stacking multiple continuous attention networks raises the optimization difficulty and makes the training be fragile.

Effect of sharing parameters As discussed in § 3.3 and § 3.4, one can share the projection matrices W_w and W_p across all transformer blocks (#2, totally 2 matrices), share W_w and W_p within the block (#3, totally 6 matrices) or use the same series of Q and K transformation matrices inside of RSAN (#4). Table 6 shows that these strategies achieve almost comparable results with the default setting but consuming less parameters. And they all outperform the baseline by a large margin.




Figure 5. Illustration of constructing random matrices.

Comparisons of up-sampling methods Except for the up-sampling operation depicted in Figure 3, which converts $ATTN_2$ to $ATTN'_2$, we also test two well-popular up-sampling methods in CV, including nearest interpolation and linear interpolation. Unfortunately, both methods behave much inferior to ours. Notably, the system with linear interpolation even underperforms the baseline. This observation indicates that the proposed UMST is a well-defined architecture with specifically designed components.

BLEU against encoder depths Figure 4 (left) plots the BLEU scores of UMST and Transformer against different encoder depths. Obviously, we observe UMST beats Transformer under all configurations, attaining almost a 0.76 BLEU gap in average. This observation further demonstrates the effectiveness of the proposed method.

BLEU against BPE operations Figure 4 (middle) compares the performances between UMST and Transformer using different vocabularies. As plotted in Figure 4 (right), more sentences are likely to be separated into sub-tokens when a vocabulary gets smaller. In a nutshell, the word-boundary information is more essential within a small vocabulary, where UMST can gain more benefits. Intuitively, UMST achieves consistent BLEU improvements within different-scale BPE merging operations. It gains the largest benefit when the BPE merging operation is 7K.

Discussion on efficiency Deployment of transformer models on devices is important for practical applications. Table 7 summarizes the comparison between UMST and Transformer in terms of the inference speed, training cost and

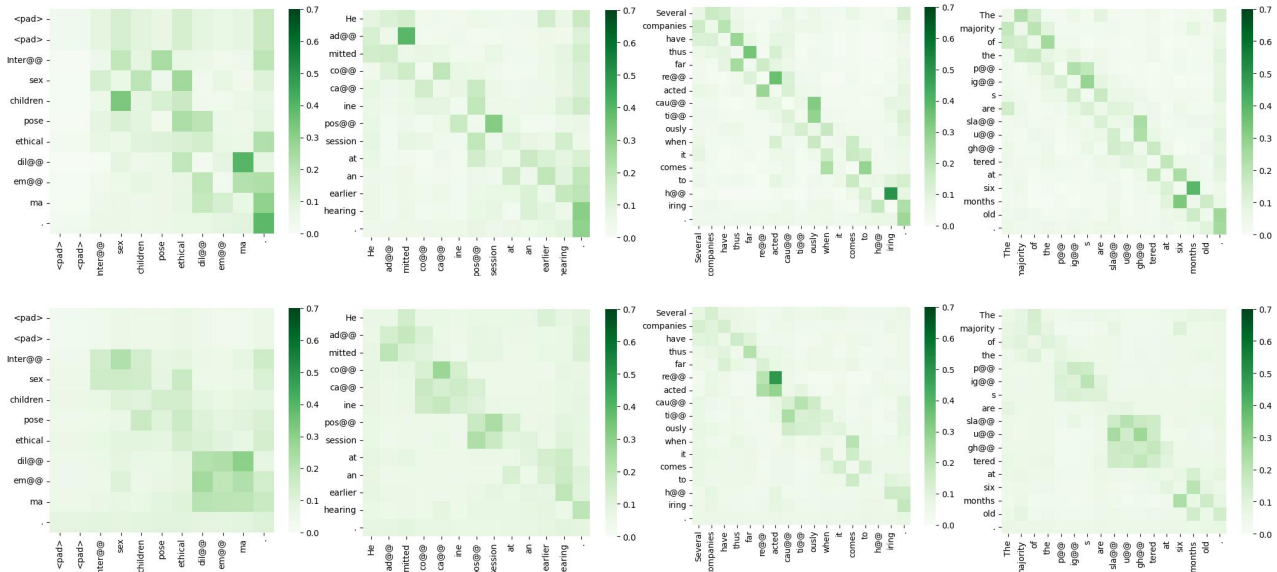


Figure 6. Quantitative examples of the attention distribution over several real cases. The above is the distribution generated by the standard Transformer, and the below is ours (UMST). Dark color means a higher value in the distribution.

Table 8. The results in terms of COMET and BLEURT on the WMT En-De task.

Model	BLEU	COMET	BLEURT
Transformer	27.63/28.67	48.24/51.43	71.03/72.01
UMST	28.51/29.49	49.04/52.55	71.31/72.63
UMST w/o inter-group	28.06/29.37	48.78/52.42	71.21/72.49
+ Random	26.98/27.11	46.10/43.04	70.54/70.48
+ Shuffle	28.08/28.96	48.36/51.80	71.01/72.38
UMST + Random P	28.20/-	48.90/-	71.28/-

GPU consumption. As we can see that UMST only brings $0.25\times$ extra training cost, with almost no inference latency, since we only apply our method in the encoder side. This is because the inference overhead mainly comes from the decoder side due to the auto-regressive decoding schema. Also, our UMST only consumes more than 0.3G GPU allocation, which is acceptable.

Evaluation with Other Metrics Previous work has pointed out that BELU may correlate poorly with human judgments. To eliminate the randomization, we also evaluate our proposed UMST on another two evaluation metrics: BLEURT (Sellam et al., 2020) and COMET (Rei et al., 2020). Table 8 displays the results. As the comparison of the first two lines in Table 8, we observe the same phenomenon with BLEU that UMST can also outperform the baseline over a large margin on these two metrics. This observation provides a strong evidence for the effectiveness.

Random Adjacency Matrices As analyzed above, UMST attains satisfactory performance gains with our well-designed components. However, it is still worthy to figure out whether the gains come from the proper usage of multiscale information or the enlarging model parameters. Here, we make a further in-depth analysis on it to answer this question. A natural idea is to use random adjacency matrices instead of the right one.

We elaborate two manners to construct the random adjacency matrices of \mathcal{A}_w and \mathcal{A}_p . For the \mathcal{A}_w , we provide two choices, namely Random \mathcal{A}_w and Shuffle \mathcal{A}_w . Random \mathcal{A}_w is a totally random matrix in which each entry is either 1 or 0. However, such a design of \mathcal{A}_w directly breaks the physical logic of BPE partitions, which are composed of several continuous elements. To remedy this problem, we design Shuffle \mathcal{A}_w , which is a shuffle of the original BPE partitions. The detailed operations are illustrated in Figure 5.

Table 8 summarizes the comparison of two random manners of \mathcal{A}_w ². As we expect, random initialization achieves a much weaker performance, which is even inferior to the baseline. This observations demonstrates the improvement mainly comes from the successful utilization of word-boundary information. On the other hand, shuffling the matrix maintains similar results with pure word-boundary information when the encoder depth is 6. This is because shuffling the matrix enables the model to learn implicit local fine-grained features. As we can see that part of cor-

²Note that we drop the inter-group interaction and replace the \mathcal{A}_w with the random versions of \mathcal{A}_w

rect word-boundary information is retained (Figure 5 a vs. c). However, it stills hinders the performance in terms of COMET (0.42) and BLEURT (0.20). For the robustness, we find the UMST can still deliver robust performance when injecting noise parser input (Random). Note that we keep the correctness of \mathcal{A}_w and replace the \mathcal{A}_p with the random version of \mathcal{A}_p . It obtains 0.14 BLEU gains compared with UMST w/o inter-group, though it still underperforms by the correct parser.

Quantitative Examples As we discussed in § 1, the standard Transformer tends to generate similar and dispersive attention distributions. This can be observed in Figure 6 (above), where attention maps of several sentences in the same batch share a similar distribution, which is not what we expected. For our model, the attention distributions tend to be more diverse and show distinct word boundaries, which follows their own linguistic characteristics. The results are shown in Figure 6 (below). For example, given a sentence “Intersex children pose ethical dilemma.” whose BPE-level partition is “Inter@@ sex children pose ethical dil@@ em@@ ma.”, the areas composed of “Inter@@”, “sex” and composed of “dil@@”, “em@@”, “ma” show darker color, which indicates strong relationship within the word. Similar phenomena can be observed in the other three cases. More results please refer to Appendix B.

7. Related Work

Transformer serve as the primary backbone for natural language processing tasks, such as machine translation (Vaswani et al., 2017), language modeling (Baeviski & Auli, 2019; Dai et al., 2019), pretraining (Devlin et al., 2019; Brown et al., 2020), speech processing (Gulati et al., 2020). Learning deep Transformer models (Wang et al., 2019; Li et al., 2020; Brown et al., 2020; Li et al., 2021a) or improving the parameter efficiency (Bai et al., 2019; Li et al., 2022; Mehta et al., 2020) are effective to gain more performance gains. However, they only encode single-scale features instead of multiscale features that are widely used in computer vision.

Computer vision Multiscale modeling has been witnessed in the computer vision area since the emergence of pyramid feature network (Lin et al., 2017), which is the main backbone for object-detection, yielding the state-of-the-art performances compared with standard ResNet (He et al., 2016). More recently, Vision Transformer (ViT) shows strong ability to handle vision tasks compared with convolutional networks when the data is sufficient. It divides the input feature into several patches and recombines them in a sequential way without architecture modification. The success of ViT attract lots of researchers to design more effective backbones upon it. It is a natural idea to enable ViT

to capture patterns at different scales, such as incorporating both local and global patterns (Vaswani et al., 2021; Yang et al., 2021; Wu et al., 2021; Chen et al., 2021; Chu et al., 2021; Liu et al., 2021), multiscale vision Transformers (Fan et al., 2021a; Li et al., 2021b; Wang et al., 2021).

Natural language processing However, directly transferring the above techniques to natural language processing tasks is nontrivial. A potential explanation is that vision features are better at constructed multi-granularity representations as the consecutive pixels rather than discrete text inputs. As an alternative, researchers redesign the components in Transformer at different granularities. Concretely, Fan et al. (2021b) captured locally fine-grained information via restricting the receive filed of heads. Local pattern modeling has been a hot topic, recently. And several works (Hao et al., 2019; Zhao et al., 2019) have shown local information is indeed helpful as adjacent words are correlated in similar semantic spaces. Wu et al. (2018) developed a phrase-level self-attention mechanism to improve the computation efficiency. Wei et al. (2020) associated the decoder with the encoder with multi-granular dependencies in different space-scales. Another emerging thread of research is to incorporate various input features via combing the sentence-piece and BPE (Wu et al., 2020) or several BPE vocabularies in different sizes (Morishita et al., 2018).

8. Conclusions

In this work, we revisit the design of multiscale Transformer for sequence generation. First, we partition scales into sub-words, words and phrases from the linguistic unit perspective. Upon this, we establish interactions among different scales to incorporate the word boundaries and phrase-level information to compensate for the global interactions. Extensive results on two benchmarks, machine translation and abstractive summarization, show the superiority of our method. As another bonus, UMST provides an opportunity of incorporating the prior knowledge into Transformer, which may shed lights on the subsequent researchers. Our code is available at <https://github.com/libeiniu/UMST>.

Acknowledgments

This work was supported in part by the National Science Foundation of China (Nos. 61876035 and 61732005), the China HTRD Center Project (No. 2020AAA0107904) and Yunnan Provincial Major Science and Technology Special Plan Projects (Nos. 202002AD080001 and 202103AA080015). The authors would like to thank anonymous reviewers for their valuable comments. And thank Yufan Jiang for his helpful advice to improve the paper.

References

- Baevski, A. and Auli, M. Adaptive input representations for neural language modeling. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019. URL <https://openreview.net/forum?id=ByxZX20qFQ>.
- Bai, S., Kolter, J. Z., and Koltun, V. Deep equilibrium models. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, pp. 688–699, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/01386bd6d8e091c2ab4c7c7de644d37b-Abstract.html>.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., and Zagoruyko, S. End-to-end object detection with transformers. In Vedaldi, A., Bischof, H., Brox, T., and Frahm, J. (eds.), Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part I, volume 12346 of Lecture Notes in Computer Science, pp. 213–229. Springer, 2020.
- Chen, C., Panda, R., and Fan, Q. Regionvit: Regional-to-local attention for vision transformers. CoRR, abs/2106.02689, 2021. URL <https://arxiv.org/abs/2106.02689>.
- Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., and Shen, C. Twins: Revisiting spatial attention design in vision transformers. CoRR, abs/2104.13840, 2021. URL <https://arxiv.org/abs/2104.13840>.
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q., and Salakhutdinov, R. Transformer-XL: Attentive language models beyond a fixed-length context. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 2978–2988, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1285. URL <https://aclanthology.org/P19-1285>.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, L. Universal transformers. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019. URL <https://openreview.net/forum?id=HyzdRiR9Y7>.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. 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), pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., and Feichtenhofer, C. Multiscale vision transformers. CoRR, abs/2104.11227, 2021a. URL <https://arxiv.org/abs/2104.11227>.
- Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., and Feichtenhofer, C. Multiscale vision transformers. arXiv preprint arXiv:2104.11227, 2021b.
- Fan, Z., Gong, Y., Liu, D., Wei, Z., Wang, S., Jiao, J., Duan, N., Zhang, R., and Huang, X. Mask attention networks: Rethinking and strengthen transformer. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 1692–1701, Online, June 2021c. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.135. URL <https://aclanthology.org/2021.naacl-main.135>.
- Gehrmann, S., Deng, Y., and Rush, A. Bottom-up abstractive summarization. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language

- Processing, pp. 4098–4109, Brussels, Belgium, October–November 2018.
- Gulati, A., Qin, J., Chiu, C., Parmar, N., Zhang, Y., Yu, J., Han, W., Wang, S., Zhang, Z., Wu, Y., and Pang, R. Conformer: Convolution-augmented transformer for speech recognition. In Meng, H., Xu, B., and Zheng, T. F. (eds.), *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pp. 5036–5040. ISCA, 2020. URL <https://doi.org/10.21437/Interspeech.2020-3015>.
- Guo, Q., Qiu, X., Liu, P., Xue, X., and Zhang, Z. Multi-scale self-attention for text classification. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 7847–7854. AAAI Press, 2020. URL <https://aaai.org/ojs/index.php/AAAI/article/view/6290>.
- Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., and Wang, Y. Transformer in transformer. *arXiv preprint arXiv:2103.00112*, 2021.
- Hao, J., Wang, X., Shi, S., Zhang, J., and Tu, Z. 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)*, pp. 887–897, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1082. URL <https://aclanthology.org/D19-1082>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- Hermann, K. M., Kociský, T., Grefenstette, E., Espeholt, L., Kay, W., Suleyman, M., and Blunsom, P. Teaching machines to read and comprehend. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 1693–1701, 2015. URL <https://proceedings.neurips.cc/paper/2015/hash/afdec7005cc9f14302cd0474fd0f3c96-Abstract.html>.
- Kasai, J., Cross, J., Ghazvininejad, M., and Gu, J. Non-autoregressive machine translation with disentangled context transformer. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 5144–5155. PMLR, 2020. URL <http://proceedings.mlr.press/v119/kasai20a.html>.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=SJU4ayYgl>.
- Kitaev, N. and Klein, D. Constituency parsing with a self-attentive encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2676–2686, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1249. URL <https://www.aclweb.org/anthology/P18-1249>.
- Kudo, T. and Richardson, J. SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 66–71, Brussels, Belgium, November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-2012. URL <https://aclanthology.org/D18-2012>.
- Lee, J., Mansimov, E., and Cho, K. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1173–1182, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1149. URL <https://aclanthology.org/D18-1149>.
- Lei Ba, J., Kiros, J. R., and Hinton, G. E. Layer normalization. *ArXiv preprint*, abs/1607.06450, 2016. URL <https://arxiv.org/abs/1607.06450>.
- Li, B., Wang, Z., Liu, H., Jiang, Y., Du, Q., Xiao, T., Wang, H., and Zhu, J. Shallow-to-deep training for neural machine translation. In *Proceedings of the 2020*

- Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 995–1005, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.72. URL <https://aclanthology.org/2020.emnlp-main.72>.
- Li, B., Wang, Z., Liu, H., Du, Q., Xiao, T., Zhang, C., and Zhu, J. Learning light-weight translation models from deep transformer. In Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021, pp. 13217–13225. AAAI Press, 2021a. URL <https://ojs.aaai.org/index.php/AAI/article/view/17561>.
- Li, B., Du, Q., Zhou, T., Jing, Y., Zhou, S., Zeng, X., Xiao, T., Zhu, J., Liu, X., and Zhang, M. ODE transformer: An ordinary differential equation-inspired model for sequence generation. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 8335–8351, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.571. URL <https://aclanthology.org/2022.acl-long.571>.
- Li, Y., Wu, C., Fan, H., Mangalam, K., Xiong, B., Malik, J., and Feichtenhofer, C. Improved multiscale vision transformers for classification and detection. CoRR, abs/2112.01526, 2021b. URL <https://arxiv.org/abs/2112.01526>.
- Lin, C.-Y. ROUGE: A package for automatic evaluation of summaries. In Text Summarization Branches Out, pp. 74–81, Barcelona, Spain, 2004. Association for Computational Linguistics. URL <https://aclanthology.org/W04-1013>.
- Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., and Belongie, S. Feature pyramid networks for object detection. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2117–2125, 2017.
- Liu, X., Wang, L., Wong, D. F., Ding, L., Chao, L. S., and Tu, Z. Understanding and improving encoder layer fusion in sequence-to-sequence learning. ArXiv preprint, abs/2012.14768, 2020a. URL <https://arxiv.org/abs/2012.14768>.
- Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., Lewis, M., and Zettlemoyer, L. Multilingual denoising pre-training for neural machine translation. Transactions of the Association for Computational Linguistics, 8:726–742, 2020b. doi: 10.1162/tacl_a.00343. URL <https://aclanthology.org/2020.tacl-1.47>.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. CoRR, abs/2103.14030, 2021. URL <https://arxiv.org/abs/2103.14030>.
- Mehta, S., Ghazvininejad, M., Iyer, S., Zettlemoyer, L., and Hajishirzi, H. Delight: Very deep and light-weight transformer. ArXiv preprint, abs/2008.00623, 2020. URL <https://arxiv.org/abs/2008.00623>.
- Morishita, M., Suzuki, J., and Nagata, M. Improving neural machine translation by incorporating hierarchical subword features. In Proceedings of the 27th International Conference on Computational Linguistics, pp. 618–629, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://aclanthology.org/C18-1052>.
- Nallapati, R., Zhou, B., dos Santos, C. N., Gülçehre, Ç., and Xiang, B. Abstractive text summarization using sequence-to-sequence rnns and beyond. In Goldberg, Y. and Riezler, S. (eds.), Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning, CoNLL 2016, Berlin, Germany, August 11-12, 2016, pp. 280–290. ACL, 2016.
- Ott, M., Edunov, S., Grangier, D., and Auli, M. Scaling neural machine translation. In Proceedings of the Third Conference on Machine Translation: Research Papers, pp. 1–9, Brussels, Belgium, 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-6301. URL <https://aclanthology.org/W18-6301>.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. fairseq: A fast, extensible toolkit for sequence modeling. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations), pp. 48–53, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-4009. URL <https://aclanthology.org/N19-4009>.
- Qi, P., Dozat, T., Zhang, Y., and Manning, C. D. Universal dependency parsing from scratch. In Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies, pp. 160–170, Brussels, Belgium, October 2018. Association for Computational Linguistics. URL <https://nlp.stanford.edu/pubs/qi2018universal.pdf>.

- Rei, R., Stewart, C., Farinha, A. C., and Lavie, A. COMET: A neural framework for MT evaluation. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 2685–2702, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.213. URL <https://aclanthology.org/2020.emnlp-main.213>.
- Sellam, T., Das, D., and Parikh, A. BLEURT: Learning robust metrics for text generation. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 7881–7892, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.704. URL <https://aclanthology.org/2020.acl-main.704>.
- Sennrich, R., Haddow, B., and Birch, A. 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), pp. 1715–1725, Berlin, Germany, 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1162. URL <https://aclanthology.org/P16-1162>.
- Shaw, P., Uszkoreit, J., and Vaswani, A. 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), pp. 464–468, New Orleans, Louisiana, 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-2074. URL <https://aclanthology.org/N18-2074>.
- Sutskever, I., Vinyals, O., and Le, Q. V. Sequence to sequence learning with neural networks. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, pp. 3104–3112, 2014. URL <https://proceedings.neurips.cc/paper/2014/hash/a14ac55a4f27472c5d894ec1c3c743d2-Abstract.html>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, pp. 5998–6008, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Vaswani, A., Ramachandran, P., Srinivas, A., Parmar, N., Hechtman, B. A., and Shlens, J. Scaling local self-attention for parameter efficient visual backbones. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021, pp. 12894–12904. Computer Vision Foundation / IEEE, 2021. URL https://openaccess.thecvf.com/content/CVPR2021/html/Vaswani_Scaling_Local_Self-Attention..
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net, 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>.
- Wang, Q., Li, B., Xiao, T., Zhu, J., Li, C., Wong, D. F., and Chao, L. S. Learning deep transformer models for machine translation. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 1810–1822, Florence, Italy, 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1176. URL <https://aclanthology.org/P19-1176>.
- Wang, W., Xie, E., Li, X., Fan, D., Song, K., Liang, D., Lu, T., Luo, P., and Shao, L. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. CoRR, abs/2102.12122, 2021. URL <https://arxiv.org/abs/2102.12122>.
- Wei, X., Yu, H., Hu, Y., Zhang, Y., Weng, R., and Luo, W. Multiscale collaborative deep models for neural machine translation. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 414–426, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.40. URL <https://aclanthology.org/2020.acl-main.40>.
- Wu, F., Fan, A., Baevski, A., Dauphin, Y. N., and Auli, M. Pay less attention with lightweight and dynamic convolutions. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019. URL <https://openreview.net/forum?id=SkVhlh09tX>.
- Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., and Zhang, L. Cvt: Introducing convolutions to vision transformers. CoRR, abs/2103.15808, 2021. URL <http://arxiv.org/abs/2103.15808>.

- Wu, L., Xie, S., Xia, Y., Fan, Y., Lai, J., Qin, T., and Liu, T. Sequence generation with mixed representations. In Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research, pp. 10388–10398. PMLR, 2020. URL <http://proceedings.mlr.press/v119/wu20e.html>.
- Wu, W., Wang, H., Liu, T., and Ma, S. Phrase-level self-attention networks for universal sentence encoding. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 3729–3738, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1408. URL <https://aclanthology.org/D18-1408>.
- Xu, J., Zhou, H., Gan, C., Zheng, Z., and Li, L. Vocabulary learning via optimal transport for neural machine translation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pp. 7361–7373, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.571. URL <https://aclanthology.org/2021.acl-long.571>.
- Yang, B., Wang, L., Wong, D. F., Chao, L. S., and Tu, Z. 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), pp. 4040–4045, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1407. URL <https://aclanthology.org/N19-1407>.
- Yang, J., Li, C., Zhang, P., Dai, X., Xiao, B., Yuan, L., and Gao, J. Focal self-attention for local-global interactions in vision transformers. CoRR, abs/2107.00641, 2021. URL <https://arxiv.org/abs/2107.00641>.
- Zhao, G., Sun, X., Xu, J., Zhang, Z., and Luo, L. MUSE: parallel multi-scale attention for sequence to sequence learning. CoRR, abs/1911.09483, 2019. URL <http://arxiv.org/abs/1911.09483>.

A. Experimental Setups

A.1. Hyperparameters

We choose Pre-Norm Transformer as the backbone due to its training stability. All systems were trained via Adam optimizer (Kingma & Ba, 2015), where β_1 and β_2 were set to 0.9 and 0.997. The learning rate and warmup-step were $2e^{-3}/16000$ and $2e^{-3}/8000$ for the machine translation and abstractive summarization tasks, respectively. Our codebase is built based on Fairseq (Ott et al., 2019), and it would be publicly available soon.

For the machine translation task, we measured the performance in terms of BLEU using beam search strategy, where the beam width was 4 and the length penalty was 0.6. For the abstractive summarization task, we reported the Rouge-1, Rouge-2 and Rouge-L (Lin, 2004) for comparisons with previous works.

A.2. Configuration Details

We mainly built the model based on Transformer Deep and Big configurations. For the deep model, the hidden size is 512 and the filter size of FFN is 2048. We split the hidden space into 8 pieces for the multi-head attention mechanism. The values of dropout are set to 0.1, and so as the label smoothing. For the big model, the hidden size and the filter size are twice larger compared with the deep model. Note that the residual dropout is 0.3 for big models.

For fair comparisons, we re-implement several strongly related works under the same codebase. Also, we follow the setup in Ott et al. (2018)’s work to simulate a 128-gpu batching schema via the gradient accumulation strategy, where the max-token size is 9600 and every 8 steps to update the parameters. Note that Zhao et al. (2019) adopted cosine learning rate schema, we still use the standard decay schedule which is proportionally to the inverse square root of the current step. Thus the experiments listed may slightly different from theirs, but resulting in a fairer comparison.

B. Additional Experiments and Analysis

B.1. Comparisons with Previous Works

As we discussed in § 1, multiscale feature hierarchies have been successfully applied to convolutional and self-attentional networks, e.g. Transformer-in-Transformer (namely TNT) (Han et al., 2021), Multiscale Vision Transformer (namely MViT) (Fan et al., 2021a). Here, we follow their merits and redesign the Transformer architecture. The details are as follows:

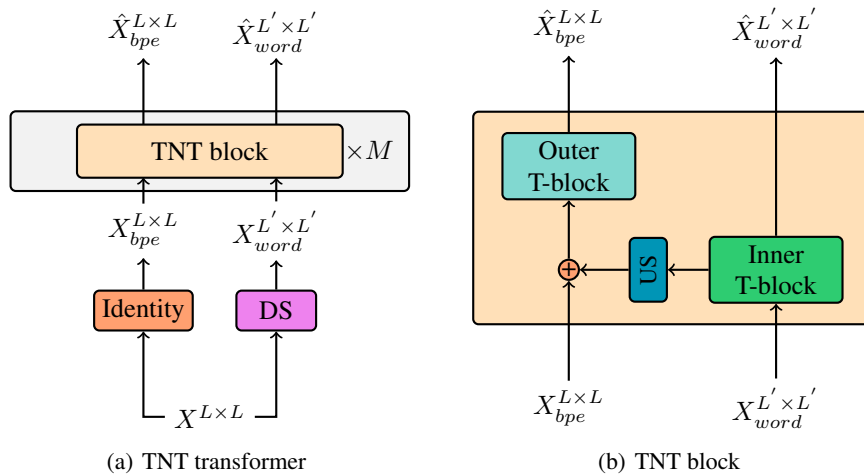


Figure 7. The architecture of the re-implementation for TNT.

TNT TNT is an improved variant of Vision Transformer (Dosovitskiy et al., 2021). It regards the local patches (e.g., 16×16) as “visual sentences” and they further split them into much smaller patches (e.g., 4×4) as “visual words”. There are two data flows in which one flow operates across the visual sentences and the other processes the visual words inside each sentence. Concretely, an inner transformer block is used to extract correlations among visual words and the output is then

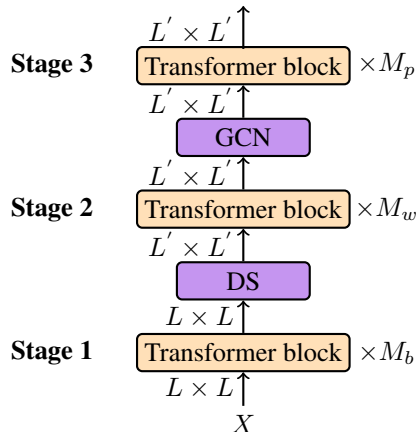


Figure 8. The architecture of Multi-stage Transformer.

added with the original visual sentences representations. At last, an outer transformer block computes the correlations among visual sentences. Also they employ residual connections among inner Transformer blocks to strengthen the information flow. Following this idea, we redesign the Transformer architecture and overall architecture is depicted in Figure 7. Our intention is to extract the inter-group and inter-individual interactions via transformer blocks and outer transformer blocks, respectively.

Multi-stage Transformer Also, for a more comprehensive comparison, we show the architecture follows a multi-stage fashion. As is shown in Figure 8, instead of keeping the length unchanged among different blocks, we partition multi-stages according to the scale. The lower-level is modeling correlations among the individual, which is a mixture sequence of words and sub-words. Subsequently, we used a down-sampling operation to convert individuals into group-level representations. After several stacked transformer blocks, a GCN is employed to embed the phrase-level dependency among words. For a fair comparison, M_b , M_w and M_p are set to 2.

Results Table 9 lists the results of UMST and the other two multiscale variants. As we discussed in § 5.1, TNT delivers a BLEU point of 28.48, which outperforms Transformer by a 0.85 BLEU point. To further validate whether the improvement comes from the additional parameters, we also share the parameters between the inner and outer blocks. As we can see, that it can still outperform the baseline by a BLEU point of 0.47. This observation further demonstrates the effectiveness of intra-group interactions. On the other hand, we find that Multi-stage Transformer behaves inferior to our default strategy. A potential explanation is the compression rate of a text is much less than that of an image.

Table 9. Comparisons of UMST with two multiscale variants.

Model	Param	BLEU
Transformer	65M	27.63
TNT (Han et al., 2021)	83M	28.48
TNT (shared inner and outer)	65M	28.10
Multi-stage Transformer	65M	27.96
UMST	70M	28.51

B.2. Quantitative Examples on Summarization

We have shown the quantitative examples on the machine translation task. Here, we choose a much longer sentence case on the abstractive summarization task. To make a clear presentation, we take parts from a long sentence and re-normalize the attention distribution. Through Figure 9, we find similar phenomenon with the machine translation task (see Figure 6), that sub-words belong to the same word have stronger relationship than distinct individuals. Also the use of inter-group interaction helps the model capture long-range dependencies. However, the attention map of the standard Transformer is mussy and much smoother. The finding here also convinces that multiscale Transformer is indeed helpful.




Figure 9. Quantitative examples of the attention distribution over several cases on abstractive summarization task. The above is the distribution generated by the standard Transformer, and the below is ours (UMST). Dark color means a higher value in the distribution.