

TranSFormer: Slow-Fast Transformer for Machine Translation

Bei Li^{1*}, Yi Jing¹, Xu Tan², Zhen Xing³, Tong Xiao^{1,4†} and Jingbo Zhu^{1,4}

¹School of Computer Science and Engineering, Northeastern University, Shenyang, China

²Microsoft Research Asia, ³Fudan University

⁴NiuTrans Research, Shenyang, China

{libei_neu, jingyi_neu}@outlook.com, xuta@microsoft.com

zxing20@fudan.edu.cn, {xiaotong, zhujingbo}@mail.neu.edu.cn

Abstract

Learning multiscale Transformer models has been evidenced as a viable approach to augmenting machine translation systems. Prior research has primarily focused on treating subwords as basic units in developing such systems. However, the incorporation of fine-grained character-level features into multiscale Transformer has not yet been explored. In this work, we present a Slow-Fast two-stream learning model, referred to as TranSFormer, which utilizes a “slow” branch to deal with subword sequences and a “fast” branch to deal with longer character sequences. This model is efficient since the fast branch is very lightweight by reducing the model width, and yet provides useful fine-grained features for the slow branch. Our TranSFormer shows consistent BLEU improvements (larger than 1 BLEU point) on several machine translation benchmarks.

1 Introduction

Transformer (Vaswani et al., 2017) has demonstrated strong performance across a range of natural language processing (NLP) tasks. Recently, learning multiscale Transformer models has been evidenced as a promising approach to improving standard Transformer. Previous research on this line can be broadly categorized into two streams: one learns local fine-grained features by using a fixed-length window (Yang et al., 2019; Hao et al., 2019; Guo et al., 2020), linguistic-inspired local patterns (Li et al., 2022b), and a hybrid approach that combines convolution and self-attention models (Gulati et al., 2020) or run in parallel (Zhao et al., 2019); the other learns sequence representations by considering multiple subword segmentation/merging schemas (Wu et al., 2020).

Despite the attractiveness of these approaches, previous work is based on an assumption that sub-




Figure 1: Overview of the proposed TranSFormer, an encoder-decoder paradigm where the encoder involves two separated branches. One is the Slow-branch with subword-level input features and the other is the Fast-Branch with character-level features. H_f is far smaller than H_s for computation efficiency.

words are the basic units in sequence modeling, and therefore ignores smaller, more fine-grained character-level features. In fact, the benefits of using characters have long been appreciated, and character-based models have been discussed in several sub-fields of NLP, such as language modeling (Xue et al., 2022) and machine translation (Lee et al., 2017; Li et al., 2021; Gao et al., 2020). But there are still important problems one needs to address in multi-scale Transformer. The first of these is the computational challenge of dealing with long sequences. For example, when we represent an English text as a character sequence, the length of this sequence is in general $5 \times$ longer than that of the subword sequence. We therefore need to consider this length difference in model design. The second problem is that, from a multiscale learning perspective, learning text representations with features at different levels is not just making use of the syntactic hierarchy of language. To better model the problem, we need some mechanism to describe the interactions among these different linguistic units.

In this study, we aim to exploit the potential of character-level representations in multiscale sequence models while maintaining computational

*The work was done when the first author was an intern at Microsoft Research Asia.

† Corresponding author.

efficiency. Drawing inspiration from the SlowFast convolutional models in video classification (Feichtenhofer et al., 2019), we propose the SlowFast Transformer (TranSFormer) model, which utilizes a fast, thin branch to learn fine-grained character-level features and a slow, wide branch to capture correlations among subword features. A cross-granularity attention layer is placed between the self-attention and feedforward sublayers to make exchanges of cross-granularity information. This enables the slow branch to be aware of fine-grained features while providing optimized high-level representations of the input sequence to the fast branch.

We also make use of character-to-word boundary information to model the interactions among neighboring characters in a word. Additionally, we develop a boundary-wise positional encoding method to better encode the positional information within words for the fast branch. Through a series of extensive experiments on the WMT’14 English-German, WMT’17 Chinese-English and WMT’16 English-Romanian tasks, we demonstrate that TranSFormer yields consistent performance gains while having a negligible increase in the number of parameters and computational cost. As a bonus, our TranSFormer is robust to errors caused by suboptimal tokenization or subword segmentation.

2 Related Work

Multiscale Transformer Learning multiscale Transformer is a promising way to acquire for further improvements in the machine translation task. A feasible way is to model global and local patterns to enhance Transformer models (Shaw et al., 2018; Yang et al., 2018, 2019; Zhao et al., 2019). These work mainly modeled the localness within a fixed window size upon subword input features. Apart from these, Wu et al. (2018) partitioned the input sequence according to phrase-level prior knowledge, and build attention mechanism upon phrases. Similarly, Hao et al. (2019) proposed a multi-granularity self-attention mechanism, designed to allocate different attention heads to phrases of varying hierarchical structures. Perhaps the most related work to ours is UMST (Li et al., 2022b). They re-defined the sub-word, word and phrase scales specific to sequence generation, and modeling the correlations among scales. However, more fine-grained character-level scale is not explored in the previous work due to the serve challenge for encoding long

character sequences.

Character-level NMT Fully character-level neural machine translation originates from recurrent machine translation system in Lee et al. (2017). They built a fully character-level encoder-decoder model, and utilize convolution layers to integrate information among nearby characters. Cherry et al. (2018) show the potential of character-level models which can outperform subword-level models under fully optimization. This contributes to their greater flexibility in processing and segmenting the input and output sequences, though modeling such long sequences is time-consuming. More recently, several studies analyze the benefits of character-level systems in multilingual translation scenarios (Gao et al., 2020), low-resource translation and translating to typologically diverse languages (Li et al., 2021). But these methods all simply view characters as basic units in language hierarchy, and it is still rare to see the effective use of multi-scale learning on character-based language features.

Multi-Branch Transformer The utilization of multi-branch architectures has been extensively studied in Transformer models. Early efforts in this area include the Weighted Transformer (Ahmed et al., 2017), which replaced the vanilla self-attention by multiple self-attention branches. Subsequently, the Multi-attentive Transformer (Fan et al., 2020) and Multi-Unit Transformer (Yan et al., 2020) have advanced this design schema by incorporating branch-dropout and switching noise inputs, respectively. Additionally, Wu et al. (2020) investigated the potential advantages of utilizing dual cross-attention mechanisms to simultaneously attend to both Sentencepiece (Kudo and Richardson, 2018) and subword (Sennrich et al., 2016). In this work, we take a forward step to exploit the potential of character features. We argue that a lightweight branch is sufficient to encode useful fine-grained features, an aspect that has not been previously investigated.

3 Method

The proposed TranSFormer follows a encoder-decoder paradigm (see Figure 1) which involves two encoder branches operating at different input granularities. The original subword encoder, which has a large model capacity for fully learning correlations among input individuals, is defined as the slow branch. The other branch, designed to handle

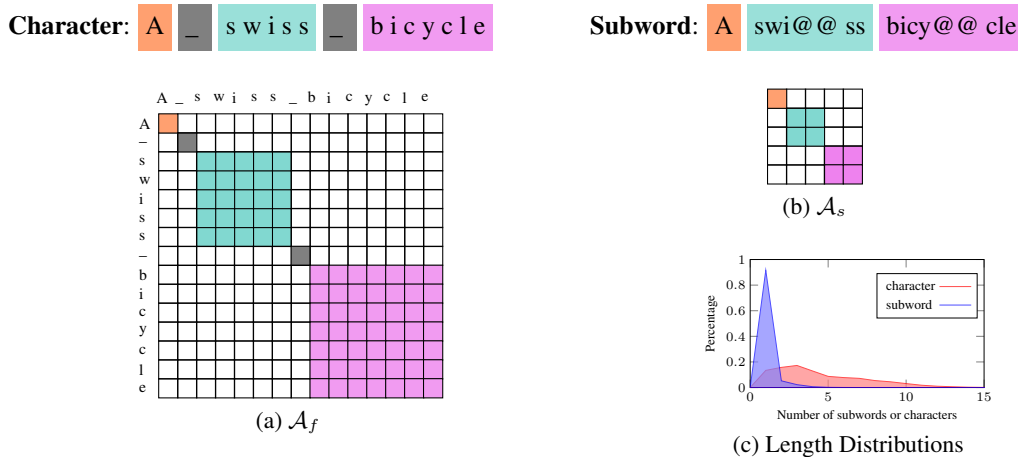


Figure 2: An example of the character-to-word boundaries (a), subword-to-word boundaries (b), along with the distribution of character/subword counts within a single word for the source language in the WMT En-De task (c). Note that characters/subwords in same color belong to the same word, and “_” denotes the word boundary among characters. Similar trends would be observed in Romanian and Mandarin.

character-level representations using a thin encoder to efficiently capture correlations among characters, is referred to as the fast branch. Our goal is to use the fast branch to learn a fine-grained but less precise representation to complement the slow branch. In the following sections, we will elaborate the core design of Slow branch, Fast branch and the cross-granularity attention, respectively.

3.1 The Slow Branch for Subwords

We use the standard Transformer as the slow branch due to its strong ability to model global interactions among input sequences. The input of the slow branch is the mixture of subwords and words since some high-frequency words have not been further divided into subwords. Following the suggestions in Li et al. (2022b), we adopt a graph convolutional network to model the inner correlations among words through the adjacency matrix \mathcal{A}_s . To this end, the Slow branch then encodes the enhanced representation via the self-attention mechanism, $\text{SAN} = \text{Softmax}(\frac{Q \cdot K^T}{\sqrt{d_k}}) \cdot V$, where Q, K, V are obtained through three independent projection matrix, such as W_q, W_k, W_v . A point-wise feed-forward network is followed, $\text{FFN} = \max(xW_1 + b_1, 0)W_2 + b_2$, where W_1 and W_2 are transformation matrices and b_1 and b_2 are bias matrices. To bridge the gap between two granularities, we sandwich a new sublayer between the self-attention and the feed-forward network, to accomplish the feature interaction between the slow and fast branches. A straightforward idea is to employ a cross-attention similar with encoder-decoder

attention in the decoder side. We will discuss more details in the Section 3.3.

3.2 The Fast Branch for Characters

To enhance the efficiency of modeling long character-level inputs, we propose the use of a fast branch with a tiny hidden size. The hidden size is a critical factor in the computation of the self-attention network (Vaswani et al., 2017), and by reducing it, we can achieve faster computation. To the best of our knowledge, this is the first attempt to design multiscale Transformer models that considers character-level features, as the long input sequence has previously hindered such exploration.

While the fast branch may not be as powerful as the slow branch, it is still effective in learning fine-grained features. Our initial experiments have yielded two notable findings: 1) a slow branch with hidden size of 32 is sufficient for transferring fine-grained knowledge to the slow branch, and 2) cross-granularity fusion is crucial for the slow branch, while removing the reversed fusion in the fast branch has only a moderate effect on performance. We would ablate this settings in the Section 4.2. To further improve the modeling ability, we introduce several techniques as follows:

Char Boundary Information The use of word-boundary information has been shown to effectively reduce the redundant correlations among subwords, as demonstrated in (Li et al., 2022b). This leads to the consideration of character-level modeling, which poses a more challenging problem

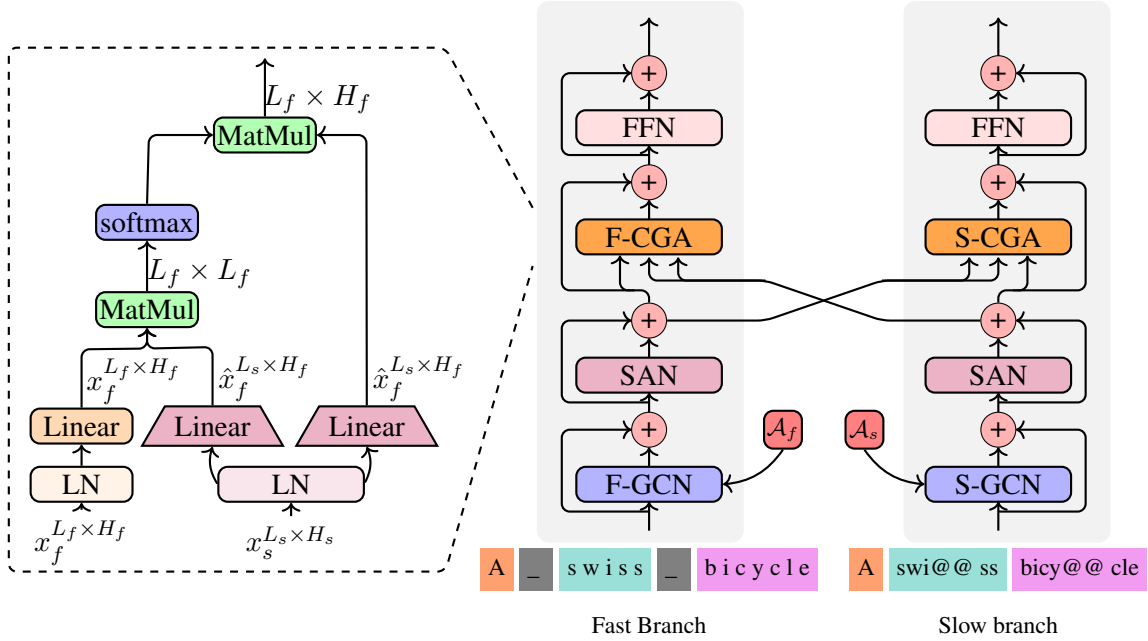


Figure 3: The encoder architecture of the proposed TransFormer, a two-branch network with fast branch (left) and slow branch (right). \mathcal{A}_f and \mathcal{A}_s are adjacency matrices of fast and slow branches. Here, we omit the layer normalization for simplification. Actually, we follow the pre-normalization strategy as its stability.

due to the greater number of characters typically present within a word in comparison to subwords. The statistical analysis in Figure 2c further evidences it that a significant proportion of words contain more than 5 characters, while a much smaller number are divided into subwords. Thus, model may be unable to discern the distinction between the same character that belongs to the same word and that of distinct words.

To address this issue, we propose the use of a character-level graph convolution network (GCN) to learn local, fine-grained features while also allowing each character to be aware of its proximity to other characters. GCN(Kipf and Welling, 2017) is a suitable choice for this task as it aggregates feature information from the neighbors of each node to encapsulate the hidden representation of that node. The computation can be described as:

$$\text{GCN}_{\text{Fast}} = \sigma(\tilde{D}_f^{-\frac{1}{2}} \tilde{\mathcal{A}}_f \tilde{D}_f^{-\frac{1}{2}} \cdot x W_f^g), \quad (1)$$

$\tilde{\mathcal{A}}_f = \mathcal{A}_f + \mathcal{I}_L$ denotes the adjacency matrix of the undirected graph with self-connections. Here \mathcal{I}_L denotes the identity matrix. \tilde{D}_f is the degree matrix of the adjacency matrix $\tilde{\mathcal{A}}_f$. W_f^g is a linear transformation which is a trainable parameter. The character-level encoder architecture is illustrated in Figure 3.

Boundary-wised Positional Encoding To further enhance the relative positional representation among characters, we design a boundary-wised positional encoding (PE) method. Our intuition is to provide each character with the ability to recognize which characters belong to the same word. Thus we restrict the relative window within each word, as illustrated in Figure 4. Vanilla relative positional encoding (Shaw et al., 2018) models the correlations in a fixed window size $2k + 1$. Here we set $k = 3$, positions exceed k would be masked. Differently, the proposed boundary-wised PE is utilized to enhance the inner relative positional information among characters within each word. In our preliminary experiments, boundary-wised PE is helpful for stable training.

3.3 Cross-Granularity Fusion

As depicted in Figure 3, the computation of the two branches in our TransFormer architecture is separated, with each branch operating independently of the other's representation. To facilitate communication between the branches, we propose the utilization of a cross-granularity information fusion method within each encoder block. This method can be implemented through various options. Given the lengths of the slow and fast branches as L_s and L_f , and the hidden sizes as H_s and H_f , respectively, the goal is to seamlessly

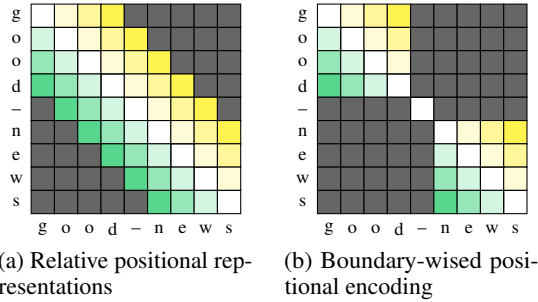


Figure 4: Comparison of the relative positional encoding (a) and our proposed boundary-wised positional encoding method (b). Note that characters in a dark color means the mask for exceeding the max window.

integrate cross-scale information within each encoder block between the two branches. In the field of machine translation, it is straightforward to employ cross-attention mechanisms, such as encoder-decoder attention (Vaswani et al., 2017) or context-aware cross-attention (Voita et al., 2018), to capture correlations between the representations.

Our default strategy is to employ a cross-granularity attention mechanism (namely CGA) sandwiched between the self-attention and feed-forward network. The architecture is plotted in Figure 3. x_f and x_s denote the representation of the fast and slow branches, respectively. The challenge remains here is the mismatched feature shape between x_s and x_f . Here, we take the Fast branch as an instance, we first normalize x_s via $\hat{x}_f = \text{LN}(x_s)$. $\text{LN}(\cdot)$ denotes the layer normalization for stable optimization. Then \hat{x}_f is fed into CGA of the fast branch, the formulation is as follows:

$$\text{ATTN}_f = \text{Softmax}\left(\frac{x_f W_f^q \cdot (\hat{x}_f W_f^k)^\top}{\sqrt{d_f^k}}\right),$$

$$\text{CGA} = \text{ATTN}_f \cdot \hat{x}_f W_f^v, \quad (2)$$

where the query is derived from the residual output of SAN in the Fast branch via $x_s \cdot W_f^q$. The key and value are derived from the Slow branch via $\hat{x}_f W_f^k$ and $\hat{x}_f W_f^v$, respectively. It is worthy to note that, the shape of W_f^k and $W_f^v \in \mathbb{R}^{H_s \times H_f}$, to reduce the hidden size. Detailed transformation could be found in the left part of Figure 3.

It is important to note that our proposed method of cross-granularity fusion is bidirectional, as opposed to the lateral connections used in the Slow-Fast (Feichtenhofer et al., 2019). Other alternative methods would be discussed in Section 4.3.

3.4 Interactions Between Encoder and Decoder

In vanilla Transformer, the key and value of the encoder-decoder attention on the decoder side derives from the encoder output, however, there are two branches in our TransFormer (See Figure 1). It is worthy to investigate how to effectively leverage the multi-granularity representations. Our default strategy is to regard the fast branch as an auxiliary to provide fine-grained features for the slow branch, thus only the output of the slow branch is exposed to the decoder. Besides this, there are also several feasible options. For example, we can fuse the outputs of two branches as the final encoder output, or building a double-branch encoder-decoder attention to attend two branches independently. We compares this options in our experiments.

4 Experiments

4.1 Experimental Setups

Datasets The present study examines the performance of our proposed TransFormer on several machine translation datasets: the WMT’14 English-German (En-De), WMT’16 English-Romanian (En-Ro) and WMT’17 Chinese-English (Zh-En) datasets. The En-De dataset comprises approximately 4.5 million tokenized sentence pairs, which were preprocessed following the same procedure as in Ott et al. (2018) to yield a high-quality bilingual training dataset. For validation, we use the *newstest2016* set, while the *newstest2014* set served as the test data. The En-Ro dataset consists of 610K bilingual sentence pairs, and we adopt the same preprocessing scripts as in Lee et al. (2018); Kasai et al. (2020), using a joint source and target BPE factorization with a vocabulary size of 40K. The *newsdev2016* set is used for validation, while the *newstest2016* set served as the test set. For the Zh-En task, we collect all the available parallel data for the WMT17 Chinese-English translation task, consisting 15.8M sentence pairs from the UN Parallel Corpus, 9M sentence pairs from the CWMT Corpus and about 332K sentence pairs from the News Commentary corpus. After carefully data filtering setups in Hassan et al. (2018), there are left 18M bilingual pairs. *newsdev2017* and *newstest2017* are served as the validation and test sets, respectively.

Setups For the machine translation task, we mainly evaluate the proposed TransFormer on base and big configurations. The hidden size of slow

	Model	Enc.	Dec.	Base		Big	
				Param	BLEU	Param	BLEU
Multiscale	Transformer (Vaswani et al., 2017)	Sub	Sub	65M	27.30	213M	28.40
	Transformer	Char	Sub	63M	26.56	208M	28.05
	RPR (Shaw et al., 2018)	Sub	Sub	65M	27.60	213M	29.20
	CSAN (Yang et al., 2019)	Sub	Sub	88M	28.18	-	28.74
	Localness (Yang et al., 2018)	Sub	Sub	89M	28.11	267M	29.18
	MG-SA (Hao et al., 2019)	Sub	Sub	89M	28.28	271M	29.01
Double-Branch	UMST (Li et al., 2022b)	Sub	Sub	70M	28.51	242M	29.75
	Muse (Zhao et al., 2019)	Sub	Sub	-	-	233M	29.90
	Multi-Attentive (Fan et al., 2020)	Sub	Sub	-	-	325M	29.80
Character-level	Multi-Unit (Yan et al., 2020)	Sub	Sub	130M	29.30	-	-
	ConvTransformer (Gao et al., 2020) †	Char	Char	65M	23.47	-	-
	Fast Only (Hidden=32, L=6)	Char	Sub	42M	17.90 _(16.9)	-	-
	Fast Only (Hidden=512, L=6)	Char	Sub	64M	27.11 _(26.1)	211M	28.65 _(27.6)
Slow-Fast	Slow Only	Sub	Sub	63M	27.40 _(26.4)	211M	28.80 _(27.8)
	TransFormer (Hidden=32, L=6)	Char/Sub	Sub	66M	28.56 _(27.6)	231M	29.85 _(28.9)
	TransFormer + ODE (Li et al., 2022a)	Char/Sub	Sub	66M	29.30 _(28.3)	-	-

Table 1: Comparison with previous studies on the WMT En-De task. Models with † denote the re-implementing results based on our codebase within the same hyperparameters. BLEU at the right corner denotes the SacreBLEU.

(a) Previous work based on Big models		
System	Params	BLEU
Transformer-Big(Hassan et al., 2018)	-	24.20
CSAN (Yang et al., 2019)	-	25.01
Localness (Yang et al., 2018)	307M	25.03
UMST (Li et al., 2022b)	307M	25.23
(b) Our Big models		
subword-level Transformer-Big	261M	24.41
character-level Transformer-Big	258M	23.80
TransFormer (Hidden=64)	283M	25.55

Table 2: Results on WMT Zh-En. We compare several prior work of learning local patterns.

branch is 512/1024 for base and big, respectively. And the filter size in FFN is 2048/4096. In our default setting, a width of 32 slow branch is enough to learn fine-grained features, and the corresponding filter size is set to 128. We both employ residual dropout, attention dropout and activation dropout. All values are 0.1, except the residual dropout 0.3 for big counterparts.

Training and Evaluations The codebase is developed upon *Fairseq* (Ott et al., 2019). All experiments are conducted on 8 Tesla V100 GPUs. We use Adam (Kingma and Ba, 2015) optimizer with (0.9, 0.997), and the default learning rate schedule with 0.002 max value, 16,000 warmup steps. For machine translation tasks, BLEU scores are computed by *mult-bleu.perl*, and we also provide the SacreBLEU¹ for En-De. The beam size is 4 for En-De and 8 for Zh-en, and the length penalty is 0.6 and 1.3, respectively.

¹BLEU+case.mixed+numrefs.1+smooth.exp+tok.13a+version.1.2.12

Model	Param	BLEU
DELIGHT (Mehta et al., 2020)	53M	34.70
Baseline in MBART (Liu et al., 2020)	-	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 (Li et al., 2022b)	60M	34.81
ODE Transformer (Li et al., 2022a)	69M	34.94
TransFormer (Hidden=32)	59M	35.40

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

Results of En-De The results of the WMT En-De task under both base and big configurations are summarized in Table 1. As evidenced by the results, our TransFormer model demonstrates significant improvements in BLEU when compared to the Slow only model, with gains of up to 1.16/1.05 BLEU scores under the base/big configurations. Conversely, the Fast only baseline, which has a hidden size of 32, only attains a BLEU score of 17.90, leading to a considerable performance gap due to its limited capacity. However, it still contributes up to a 1.14 BLEU-point benefit to the Slow branch, indicating that the fine-grained correlations modeled by the Fast branch are complementary. Additionally, we present the results of prior works, which have employed both character-level and subword-level systems, and categorize them in terms of various aspects. TransFormer can beat or on par with prior works with less parameters. This indicates the investigation of character-level multiscale models is meaningful. Note that TransFormer is computationally efficient, only requiring additional 15% training cost and negligible inference latency. And TransFormer can also benefit from

Model	Width	En-De Param.	Zh-En Param.
Transformer	-	27.20	63.1M
TranSFormer	16	28.10	66.6M
TranSFormer	32	28.56	66.9M
TranSFormer	64	28.39	67.6M
TranSFormer	128	28.27	69.7M
TranSFormer	256	28.17	76.1M

(a) **Char-width:** Comparisons of various char width on performance.

Model	Fusion Methods	BLEU
Slow only	-	27.20
Fast only	-	17.90
TranSFormer	CGA	28.56
TranSFormer	Linear + Attention	28.47
TranSFormer	DS + Concat	27.82
TranSFormer	DS + Sum	27.96

(b) **Fusion Method:** Fusing Slow and Fast branches with several types of methods.

Model	Enc. Output	BLEU
TranSFormer	Slow Branch	28.56
TranSFormer	Fast Branch	23.11
TranSFormer	Both Slow and Fast	28.25

(c) **Interactions:** Figuring out the impact of various encoder-decoder interaction manners on performance.

Model	Input Granularity		BLEU
	Slow-Branch	Fast-Branch	
TranSFormer	Subword	Character	28.56
TranSFormer	Subword	Subword	27.50
TranSFormer	Subword	Sentencepiece	28.27
TranSFormer	Sentencepiece	Character	28.60

(d) **Input:** Figuring out the impact of various input granularities for Slow and Fast branch.

Table 4: Ablations on TranSFormer design on the WMT En-De task. The evaluation metric is BLEU (%). We mainly ablate the experiments from the width of fast branch, various fusion methods, the interactions between encoder-decoder, and the input granularity.

advanced design, *e.g.*, another 0.74 improvement with ODE method (Li et al., 2022a).

4.2 Results

Results of Zh-En The WMT’17 Zh-En task poses a significant challenge due to the linguistic differences between Chinese and English. Additionally, the Chinese language owns less characters per word than English. Table 2 shows the results of our comparison of the TranSFormer with prior works. We observe TranSFormer yields a 1 BLEU point improvements than the subword-level systems. Our TranSFormer model demonstrates superior performance compared to previous work that models local patterns, while maintaining efficient computational requirements. We will exploit whether TranSFormer can gain more benefits when incorporating these techniques on the slow branch.

Results of En-Ro Furthermore, our empirical evaluations of the proposed TranSFormer architecture on the smaller WMT En-Ro dataset also demonstrate consistent improvements in BLEU scores as a result of the utilization of interactions among granularities. Notably, the TranSFormer model even outperforms the ODE Transformer (Li et al., 2022a), an advanced variant that leverages the advantages of high-order ordinary differential equations (ODE) solutions, by a substantial margin while incurring less computational cost.

4.3 Analysis

This section provides ablation studies of TranSFormer in terms of several core techniques.

Effect of width on Fast branch We first aim to explore TranSFormer under various widths of the Fast branch, including 16, 32, 64, 128 and 256. Results in Table 4a show that even a hidden size of 16 can provide helpful fine-grained features for the slow branch, and yielding almost 1 BLEU-point gains by bringing modest parameters. Empirically, a hidden of 32 and 64 deliver the best performance on base (En-De) and big (Zh-En) configurations, respectively. Further increasing the hidden layer dimension of the model results in no more gains, while requiring more computational cost.

Fusion methods between branches In addition to our proposed fusion method CGA, there are several alternative techniques that can be considered. The most straightforward one is to transform the hidden with a linear projection and then use a standard cross-attention. It delivers similar performance but consumes more parameters. Another option is to downsample the character-level representation from L_f to L_s , and then concatenate (namely DS + Concat) or sum (DS + Sum) the two representations. Although both of these methods have been found to outperform the Slow only baseline, they have not been found to be on par with CGA method. This may be due to the fact that downsampling may impede optimization due to the low compression ratio of text compared with images.

Various interaction methods In Table 4c, we present a summary of various promising options for interactions between the encoder and decoder. Empirical results indicate that utilizing the Slow

#	Model	BLEU
1	TranSFormer	28.56
2	w/ unidirectional CGA	28.41
3	w/o character-boundary	27.61
4	w/ replace boundary with random	26.79
5	w/o boundary-wised PE	28.13
6	w/ linear attention (Wang et al., 2020)	28.08

Table 5: Ablations on the fast branch in terms of several core design schemas.

#	Model	BLEU
1	TranSFormer (default: all blocks)	28.56
2	+ at the last encoder block (e.g., 6)	27.99
3	+ at bottom 3 blocks (e.g., 1/2/3)	28.10
4	+ at top 3 blocks (e.g., 4/5/6)	28.40
5	+ every 2 blocks (e.g., 1/3/5)	28.20

Table 6: Ablations on operating interactions between two branches in different levels.

branch as the output yields the highest performance. This can be attributed to the Fast branch’s ability to provide fine-grained features and low-level semantic knowledge as auxiliary information to the Slow branch. Additionally, while utilizing the Fast branch as the encoder output results in inferior performance compared to the baseline, it still yields a significant improvement over the Slow only baseline (17.90). This highlights the effectiveness of the TranSFormer model in leveraging interactions between different granularities. Furthermore, we also evaluated a two-stream approach in the decoder, in which one stream attends to the Slow branch and the other attends to the Fast branch, with a gated mechanism being used to fuse the features. However, this method was not sufficient to further improve performance. We attribute this to the negative interactions brought by the Fast branch, increasing the optimization difficulty.

Effect of various input granularities To ascertain whether the observed performance gains can be attributed to the complementary information provided by fine-grained character-level representations, we replaced the input of the fast branch with subword-level sequences, identical to that of the slow branch. The results presented in Table 4d demonstrate a degradation of up to 1 BLEU point. This can be attributed to the lack of distinct or complementary features provided by the fast branch and the limited capacity of the model in fully optimizing subword-level features. This observation further supports the hypothesis that the Slow-Fast design can learn complementary features for each granularity. Furthermore, we found that the TranS-

#	Model	50K	500K	1000K
1	TranSFormer	11.87	22.75	25.30
2	Character-only	10.50	20.50	22.50
3	Subword-only	7.00	22.00	23.50

Table 7: Comparison of different low-resource settings, including 50K, 500K, and 1000K training subsets sampled from the WMT En-De dataset.




Figure 5: BLEU against different BPE merging operations and model depths.

Former architecture with sentencepiece (Kudo and Richardson, 2018) as the fast branch input can also benefit from the two-branch design, due to the different segmentations. Additionally, our TranSFormer is a general design that can work well with a character-level fast branch and a sentencepiece-level slow branch, yielding a BLEU score of 28.60, even slightly better than the subword-level one.

Ablations on fast branch designs It is hard to directly learn the tedious character sequence. The proposed character-boundary injection serves as a crucial component in addressing this challenge. Without this injection, the TranSFormer model suffers from a significant decrease in BLEU (#3). Furthermore, the situation is exacerbated when the boundary is replaced with a randomly initialized one (#4), emphasizing the importance of the proposed character-boundary injection. Also, both removing the boundary-wised positional encoding (#5) or replacing the vanilla attention by linear attention (Wang et al., 2020) (#6) lead to modest BLEU degradation. While, there is no significant impact when using unidirectional CGA (#2, from Fast to Slow).

Ablations on interactions in different levels Our default configuration permits the model to allocate interactions at each encoder layer. It is beneficial to determine how interaction frequency impacts the performance. Table 6 compares various interaction frequencies at different levels, including exclusively at the final encoder block, the bottom three blocks, the top three blocks and every two blocks. The experiments were conducted on




Figure 6: Results on MorphEval. Where P-Gender and P-Number stands for pron2nouns-gender and pron2nouns-number.

#	Model	FLOPs
1	Baseline	1.1G FLOPs
2	TransFormer	1.4G FLOPs
3	Baseline (Big)	3.9G FLOPs
4	TransFormer (Big)	5.0G FLOPs

Table 8: Comparison of FLOPs of different models.

WMT En-De. It is evident that the default configuration delivers optimal performance. Interactions conducted at the top three blocks demonstrate superior results compared to those at the bottom three blocks. Furthermore, performing fusion solely at the last encoder block proves insufficient for the model to learn multiscale interactions effectively.

BLEU v.s. Depth and BPE mergings Figure 5 plots the performance against model depths and BPE merging operations. The proposed TransFormer architecture demonstrates consistent performance improvements as a result of increased encoder depth. Furthermore, an empirical evaluation of the TransFormer against various byte-pair encoding (BPE) operations (Sennrich et al., 2016), on the slow branch of the model yields a statistically significant average gain of 1.1 BLEU scores over the Slow only baseline.

Low resource setting and morphological evaluation Li et al. (2021) has shown that character-level systems are better at handling morphological phenomena and show strong performance in low-resource scenarios than subword-level systems. Consequently, we evaluate how TransFormer behaves at these scenarios. For the low-resource setting, we randomly select subsets of 50K, 500K, and 1000K from the WMT En-De training corpus. TransFormer achieves respective BLEU scores of 11.87, 22.75, and 25.30, while the character-only and subword-only Transformers yield approximate scores of 10.50/7.00, 20.50/22.00, and 22.50/23.50. This empirical evidence demonstrates that TransFormer effectively amalgamates the benefits of

both character-level and subword-level features. Moreover, Figure 6 plots the performance on MorphEval(Burlot and Yvon, 2017) benchmark. TransFormer behaves better than subword solely in terms of Negation, Past, P-Gender and P-Number metrics.

Comparisons in Efficiency Table 8 compares the FLOPs between baseline and our TransFormer both in base and big configurations. Due to the light computation cost of the fast branch, TransFormer only brings additional 0.3G/1.1G FLOPs in base/big configurations, respectively. Note that the bulk of the additional computational cost is associated with the upsampling/downsampling operations within the cross-granularity attention mechanism. This process aligns the hidden size between the two representations.

5 Conclusions

In this work, we comprehensively leverage the potential of character-level features in multiscale sequence models while preserving high computational efficiency. To accomplish this, we propose a Slow-Fast Transformer architecture consisting of two branches in the encoder. The slow branch, akin to the vanilla Transformer, handles subword-level features, while the fast branch captures fine-grained correlations among characters. By leveraging the complementary features provided by the fast branch, our TransFormer demonstrates consistent improvements in BLEU scores on three widely-used machine translation benchmarks. Further in-depth analyses demonstrate the effectiveness of the TransFormer and its potential as a universal multiscale learning framework.

Acknowledgments

This work was supported in part by the National Science Foundation of China (No. 62276056), the National Key R&D Program of China, the China HTRD Center Project (No. 2020AAA0107904), the Natural Science Foundation of Liaoning Province of China (2022-KF-16-01), the Yunnan Provincial Major Science and Technology Special Plan Projects (No. 202103AA080015), the Fundamental Research Funds for the Central Universities (Nos. N2216016, N2216001, and N2216002), and the Program of Introducing Talents of Discipline to Universities, Plan 111 (No. B16009).

Limitations

The proposed TransFormer architecture employs a two-branch design, which separately encodes character-level and subword-level features. Our original design of the proposed cross-granularity attention is to acknowledge the correlation between subwords and characters that belong to the same word. For example, a cross-granularity Gaussian distribution to let subwords pay more attention to the corresponding characters. However, the variability of word boundary information across sentences presents a challenge in effectively batching them and achieving high computational efficiency. This is an area of ongoing research, and will be the focus of future work. On the other hand, our current evaluation of the TransFormer architecture is limited to machine translation tasks. It is worth exploring the potential of TransFormer in optimizing character sequences on natural language understanding tasks and other sequence generation tasks, such as abstractive summarization. These tasks are more challenging in terms of encoding longer sequences, but we believe that TransFormer can serve as a versatile backbone. We aim to verify its effectiveness on these tasks in the future.

References

- Karim Ahmed, Nitish Shirish Keskar, and Richard Socher. 2017. [Weighted transformer network for machine translation](#). *CoRR*, abs/1711.02132.
- Franck Burlot and François Yvon. 2017. [Evaluating the morphological competence of machine translation systems](#). In *Proceedings of the Second Conference on Machine Translation*, pages 43–55, Copenhagen, Denmark. Association for Computational Linguistics.
- Colin Cherry, George Foster, Ankur Bapna, Orhan Firat, and Wolfgang Macherey. 2018. [Revisiting character-based neural machine translation with capacity and compression](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4295–4305, Brussels, Belgium. Association for Computational Linguistics.
- Yang Fan, Shufang Xie, Yingce Xia, Lijun Wu, Tao Qin, Xiang-Yang Li, and Tie-Yan Liu. 2020. [Multi-branch attentive transformer](#). *CoRR*, abs/2006.10270.
- Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2019. Slowfast networks for video recognition. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 6201–6210. IEEE.
- Yingqiang Gao, Nikola I. Nikolov, Yuhuang Hu, and Richard H.R. Hahnloser. 2020. [Character-level translation with self-attention](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1591–1604, Online. Association for Computational Linguistics.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented transformer for speech recognition](#). In *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pages 5036–5040. ISCA.
- Qipeng Guo, Xipeng Qiu, Pengfei Liu, Xiangyang Xue, and Zheng Zhang. 2020. [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*, pages 7847–7854. AAAI Press.
- Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. 2021. Transformer in transformer. *arXiv preprint arXiv:2103.00112*.
- Jie Hao, Xing Wang, Shuming Shi, Jinfeng Zhang, and Zhaopeng Tu. 2019. [Multi-granularity self-attention for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 887–897, Hong Kong, China. Association for Computational Linguistics.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, et al. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv preprint arXiv:1803.05567*.
- Jungo Kasai, James Cross, Marjan Ghazvininejad, and Jiatao Gu. 2020. [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*, pages 5144–5155. PMLR.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Thomas N. Kipf and Max Welling. 2017. [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.
- Taku Kudo and John Richardson. 2018. [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*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Jason Lee, Kyunghyun Cho, and Thomas Hofmann. 2017. [Fully character-level neural machine translation without explicit segmentation](#). *Transactions of the Association for Computational Linguistics*, 5:365–378.
- Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. [Deterministic non-autoregressive neural sequence modeling by iterative refinement](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182, Brussels, Belgium. Association for Computational Linguistics.
- Bei Li, Quan Du, Tao Zhou, Yi Jing, Shuhan Zhou, Xin Zeng, Tong Xiao, JingBo Zhu, Xuebo Liu, and Min Zhang. 2022a. [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)*, pages 8335–8351, Dublin, Ireland. Association for Computational Linguistics.
- Bei Li, Tong Zheng, Yi Jing, Chengbo Jiao, Tong Xiao, and Jingbo Zhu. 2022b. Learning multiscale transformer models for sequence generation. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, pages 13225–13241. PMLR.
- Jiahuan Li, Yutong Shen, Shujian Huang, Xinyu Dai, and Jiajun Chen. 2021. [When is char better than subword: A systematic study of segmentation algorithms 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 2: Short Papers)*, pages 543–549, Online. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Sachin Mehta, Marjan Ghazvininejad, Srinivasan Iyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2020. [Delight: Very deep and light-weight transformer](#). *ArXiv preprint*, abs/2008.00623.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [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)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. [Scaling neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468, New Orleans, Louisiana. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. [Context-aware neural machine translation learns anaphora resolution](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics.
- Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. [Linformer: Self-attention with linear complexity](#). *arXiv preprint arXiv:2006.04768*.
- Lijun Wu, Shufang Xie, Yingce Xia, Yang Fan, Jian-Huang Lai, Tao Qin, and Tie-Yan Liu. 2020. [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*, pages 10388–10398. PMLR.
- Wei Wu, Houfeng Wang, Tianyu Liu, and Shuming Ma. 2018. [Phrase-level self-attention networks for universal sentence encoding](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3729–3738, Brussels, Belgium. Association for Computational Linguistics.

- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. [ByT5: Towards a token-free future with pre-trained byte-to-byte models](#). *Transactions of the Association for Computational Linguistics*, 10:291–306.
- Jianhao Yan, Fandong Meng, and Jie Zhou. 2020. [Multi-unit transformers for neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1047–1059, Online. Association for Computational Linguistics.
- Baosong Yang, Zhaopeng Tu, Derek F. Wong, Fandong Meng, Lidia S. Chao, and Tong Zhang. 2018. [Modeling localness for self-attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4449–4458, Brussels, Belgium. Association for Computational Linguistics.
- Baosong Yang, Longyue Wang, Derek F. Wong, Lidia S. Chao, and Zhaopeng Tu. 2019. [Convolutional self-attention networks](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4040–4045, Minneapolis, Minnesota. Association for Computational Linguistics.
- Guangxiang Zhao, Xu Sun, Jingjing Xu, Zhiyuan Zhang, and Liangchen Luo. 2019. [MUSE: parallel multi-scale attention for sequence to sequence learning](#). *CoRR*, abs/1911.09483.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section Limitation
- A2. Did you discuss any potential risks of your work?
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract
- A4. Have you used AI writing assistants when working on this paper?

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
Not applicable. Left blank.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 4

C Did you run computational experiments?

Section 4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 5

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 4

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 4

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.