Transcormer: Transformer for Sentence Scoring with Sliding Language Modeling

Anonymous Author(s)
Affiliation
Address
email

Abstract

Sentence scoring aims at measuring the likelihood score of a sentence and is widely used in many natural language processing scenarios, like reranking, which is to select the best sentence from multiple candidates. Previous works on sentence scoring mainly adopted either causal language modeling (CLM) like GPT or masked language modeling (MLM) like BERT, which have some limitations: 1) CLM only utilizes unidirectional information for the probability estimation of a sentence without considering bidirectional context, which affects the scoring quality; 2) MLM can only estimate the probability of partial tokens at a time and thus requires multiple forward passes to estimate the probability of the whole sentence, which incurs large computation and time cost. In this paper, we propose *Transcormer* – a Transformer model with a novel *sliding language modeling* (SLM) for sentence scoring. Specifically, our SLM adopts a triple-stream self-attention mechanism to estimate the probability of all tokens in a sentence with bidirectional context and only requires a single forward pass. SLM can avoid the limitations of CLM (only unidirectional context) and MLM (multiple forward passes) and inherit their advantages, and thus achieve high effectiveness and efficiency in scoring. Experimental results on multiple tasks demonstrate that our method achieves better performance than other language modelings.

1 Introduction

Sentence scoring is to measure the quality of a sentence via language modeling. It has been widely used in many natural language processing (NLP) scenarios, especially like reranking task. Reranking aims to select the most suitable candidate from multiple candidates produced by the text generation models (e.g., machine translation or speech recognition models) based on the sentence score. Precise sentence scoring is able to select a better candidate. Therefore, how to estimate sentence score is critical for the application scenarios.

With the development of deep learning, language modeling (LM) [1, 2, 3, 4] has been considered as the most widely used technique for sentence rescoring, since it can produce probability estimation of the whole sentence by computing the probability of each token and summing up their values as the sentence score. Specifically, causal language modeling (CLM) and masked language modeling (MLM) are the most representative LMs. CLM is conditioned on unidirectional context for next token prediction. And MLM is to mask some tokens in the input sequence to utilize bidirectional context for prediction. Some works [5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16] pre-trained these LMs on large-scale corpus and thus enhance the capacity and generalization of LMs. For example, GPT [8, 9, 10]

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

adopted CLM to pre-train a large Transformer [17] model for natural language generation (NLG) tasks. BERT [5] and its descendants [6, 7, 14] leveraged MLM for pre-training to handle natural language understanding (NLU) tasks. Inspired by these successes of CLM and MLM in NLG and NLU tasks, some works [18, 19, 20, 21, 22] have tried to use these pre-trained LMs for reranking on machine translation or speech recognition tasks and achieved some promising results.

However, we notice that both CLM and MLM still remain some deficiencies in calculating sentence scores. For instance, some works [22] utilized GPT-style model for ASR reranking within single inference, yet GPT model can only extract unidirectional information due to the limitation of CLM, without considering the whole sentence semantics, and thus affect the sentence score. To utilize bidirectional context, some works [19, 20, 21] applied BERT model for rescoring. However, the nature of MLM needs model to mask some tokens in the sentence for prediction, which means it requires BERT model to forward multiple times that each forward only masks one token for prediction. As a result, it is time-consuming to directly adopt MLM for scoring sentence. Overall, we find that for calculating sentence score, CLM only needs once inference but only use unidirectional information and MLM is costly in computing sentence score although it can use bidirectional context. Therefore, we raise a natural question: is it possible to design a language modeling to use bidirectional context for sentence scoring with only once inference pass.

To address the above issues, in this paper, we introduce *Transcormer*, a Transformer model designed for sentence scoring. More specifically, our Transcormer leverages a novel language modeling, named as *sliding language modeling* (SLM), that allows us to produce the probability of all tokens within single inference pass, and simultaneously utilizes bidirectional context. To fulfill this target, we innovatively design a triple-stream self-attention mechanism, which consists of two content streams (a forward stream and a backward stream) and one query stream. By employing some specific-designed mask strategies on the attention matrix, our method allows each token in the query stream to leverage all token information except itself for estimating probability (the tokens before and after it) to avoid any information leakage. To the best of our knowledge, SLM is the first language modeling tailored for sentence scoring. We pre-train our SLM on large-scale corpus, and then evaluate it on multiple datasets for reranking. Experimental results demonstrate that our Transcormer improves the baseline by up to + 0.8/0.6 BLEU score on low-resource/rich-resource machine translation tasks and gives 20% relative improvements on ASR datasets.

The contributions of our paper can be summarized as:

- We analyze the pros and cons of CLM and MLM when using them for scoring sentences, and thus propose Transcormer with a new sliding language modeling, which uses bidirectional context for probability estimation within a single pass.
- We introduce a novel triple-stream self-attention mechanism in SLM, which has two content streams to collect forward/backward semantics, and a query stream to estimate the probability of each token in a sentence.
- Experimental results on multiple datasets demonstrate the effectiveness and efficiency of our SLM for sentence scoring.

2 Background

2.1 Sentence Scoring

Sentence scoring has a long-term development history in NLP applications, especially in reranking tasks (*e.g.*, reranking for machine translation [23, 24, 25, 26] or speech recognition [27]). Generally, given n-best candidates generated by text generation models, reranking aims at scoring each candidate to select the best answer. Early works [23, 24, 25, 26] mainly used statistical LM to calculate the sentence scores and some works [28, 29] tried to combine statistical LMs with RNN-based LM for sentence scoring. Recently, directly using end-to-end neural network based LM has became the de facto approach for scoring and has been widely used in many NLP tasks [2, 1, 30]. Specifically,

causal language modeling (CLM) and masked language modeling (MLM) are the most representative language modeling methods, which are famous as GPT [8, 9, 10] and BERT [5] respectively. The details of CLM and MLM are described in below.

CLM CLM is to model sequence based on the previous tokens in a chain-style rule. Assuming the target sequence is $\mathbf{y} = \{y_1, \cdots, y_{|y|}\}$, where y_i is the i-th token, so the probability estimation function of CLM is $\sum_{i=1}^{|y|} \log P(y_i|y_{<i})$. Hence, CLM can obtain the probability of each token of the sentence in a single pass, but can only capture unidirectional information.

MLM Unlike CLM, MLM use bidirectional context for prediction. Specifically, MLM replaces a random subset of tokens as special symbol $[\mathrm{MASK}]$ for prediction. Assuming the masked subset as $\mathbf{y}_{\mathcal{S}}$ and the corrupted sequence as $\mathbf{y}_{\setminus\mathcal{S}}$, the objective of MLM is to optimize $\sum_{i=1}^{|\mathcal{S}|} \log P(y_{\mathcal{S}_i}|\mathbf{y}_{\setminus\mathcal{S}})$. However, MLM only supports partial prediction as each time we can only mask a subset for prediction. To calculate the sentence score, a kind of solution [19, 20, 21] is to forward multiple times and only mask one token each time. Therefore, the MLM for sentence scoring is formulated as: $\sum_{i=1}^{|\mathbf{y}|} \log P(y_i|\mathbf{y}_{\setminus y_i})$. We can find that the cost of MLM for scoring needs $|\mathbf{y}|$ inference passes, which is too time-consuming. To alleviate this problem, some works attempted to use stochastic estimation [19] or distillation [18] to approximately estimate the probability of each toke produced by MLM, but the approximated score cannot precisely estimate the token probability. Therefore, how to calculate sentence score efficiently is the main challenge in MLM.

Overall, CLM only needs a single forward pass to estimate the probability of all tokens but cannot extract bidirectional context, while MLM leverages bidirectional information but needs multiple inference passes. Consequently, we raise a natural question: is it possible to design a pre-trained language model to support all token prediction in a single pass and simultaneously leverage bidirectional information? This is exactly the motivation of our method.

2.2 Multiple-Stream Self-Attention

The pioneer of multiple-stream self-attention is XLNet [11], which introduces two-streams self-attention to incorporate autoregressive pre-training for language understanding. It pre-trains Transformer [17] in an autoregressive manner via maintaining a content stream and a query stream. In details, for the t-th step, the content stream is able to capture the dependency from the tokens before the t-th step and itself (*i.e.*, $x_{\leq t}$), while the query stream is only allowed to view the tokens before the t-th step (*i.e.*, $x_{<t}$) to avoid information leakage. Besides, there also remain some other variants of two-stream self-attention [12, 31, 32]. For example, MPNet [12] used two-stream self-attention to build masked and permuted pre-training. ProphetNet [31] designed multiple query streams to predict N-gram future steps for sequence-to-sequence tasks, and ERNIE-GAN [32] proposed a multi-flow generation model, which includes two query streams for span and word prediction. We observe these works shared some similarities, that they used a single content stream and then used one or many query streams to predict more information. Different from these works, we introduce a triple-stream self-attention mechanism, which enables query stream to leverage two content streams for prediction, and thus enjoys the benefits of additional bidirectional context for estimating token probability.

3 Transcormer

To inherent the advantages of CLM and MLM for sentence scoring and avoid their limitations, we propose Transcormer – a Transformer model with a novel sliding language modeling for sentence scoring. First, we summarize that an ideal language modeling for sentence scoring should satisfy two requirements: 1) model should be able to use bidirectional context for effective probability estimation of each token; 2) model should produce the probability of all tokens in a sentence within a single inference pass for efficiency. Therefore, to fulfill these two requirements, we formulate a new language modeling, named as sliding language modeling (SLM), and describe it in Section 3.1. In SLM, we propose a Triple-Stream Self-Attention mechanism based on Transformer (please see

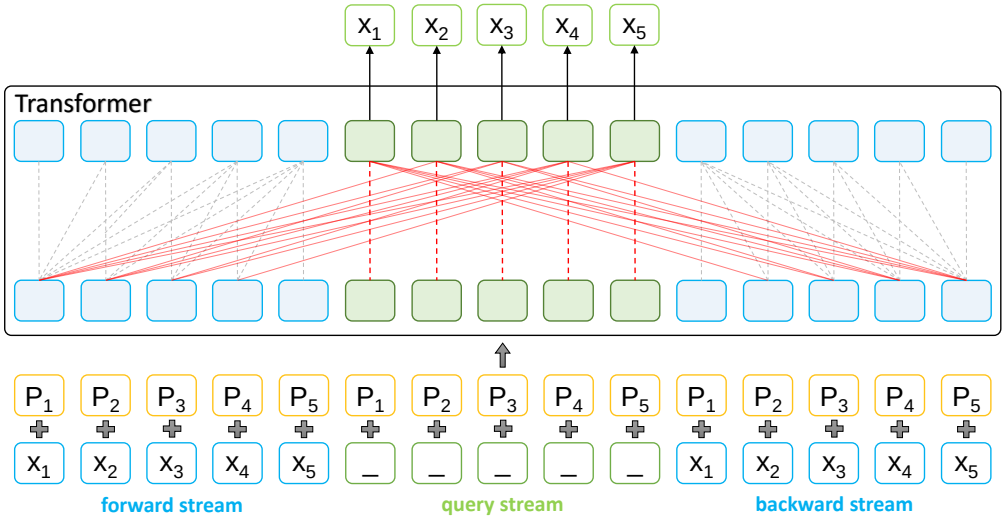

Figure 1: Transcormer with sliding language modeling. The left and the right (in blue) are forward and backward streams, respectively, and the middle (in green) is query stream. For query stream, the inputs are only the positional information. We use gray and red line to represent the allowed attended positions in the content and query streams.

Section 3.2 for details) to use bidirectional context for each token prediction and avoid information leakage. We also discuss the differences between SLM and other LMs in Section 3.3. Figure 1 presents the pipeline of our Transcormer for sentence scoring with SLM.

3.1 Sliding Language Modeling

Considering the pros and cons of CLM and MLM for scoring, we notice: 1) CLM can produce probability of all tokens within one forward pass, and thus obtain the unidirectional information of the whole sentence; 2) MLM for sentence scoring needs multiple inference passes, so that many context has been repeatedly calculated and cause a waste of computation. So, is it possible to reuse token information to build bidirectional context for token prediction?

Therefore, we propose sliding language modeling (SLM) to address the inherent flaws in previous LMs (*i.e.*, CLM and MLM) for sentence scoring. Specifically, we first maintain two individual streams to collect forward (left-to-right) context and backward (right-to-left) context. And for each token prediction, we decompose the sentence information as the past tokens (the tokens before it) and future tokens (the tokens after it) respectively. As a result, our SLM enforces each token to only capture the dependency from its past tokens and its future tokens concurrently, so that each token can utilize the whole sentence information (except itself) to estimate token probability. The objective function of SLM is formulated as:

$$\mathcal{L} = \sum_{i=1}^{|\mathbf{y}|} \log P(y_i|\mathbf{y}_{<i}, \mathbf{y}_{>i}; \theta), \tag{1}$$

where $\mathbf{y}_{<i}$ and $\mathbf{y}_{>i}$ respectively correspond to the tokens before the i-th token and after the i-th token, and θ represents the parameters of SLM. Thanks to such design, our SLM can utilize bidirectional context to produce the probability of each token within one forward pass, and thus satisfy the above requirements for sentence scoring. However, previous experiences [11, 33] pointed out that the states with bidirectional information will cause information leakage when propagating to the next layer [1]. Therefore, how to implement SLM to avoid information leakage and maintain different states together is still a troublesome problem.

[1]For example, assume the sequence has 3 tokens, and the hidden states of the i-th token at the first layer as h_i^1. So each position h_1^1, h_2^1 and h_3^1 should collect the information from positions $\{2, 3\}$, $\{1, 3\}$, $\{1, 2\}$. However, when h_1^1 and h_3^1 are delivered to h_2^2 at the second layer, it will cause a cyclic leakage as h_2^2 should not obtain information from position 2.

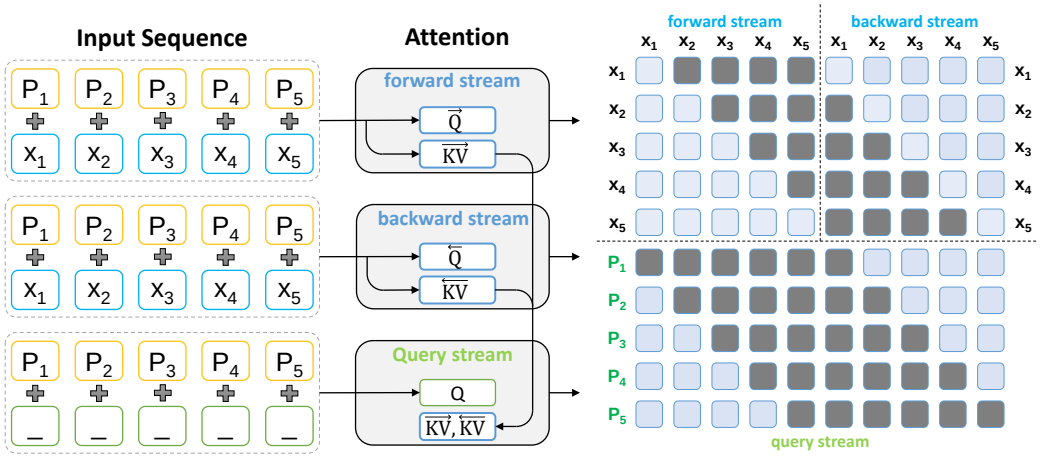

(a) (b)

Figure 2: (a) The structure of our triple-stream self-attention used in our sliding language modeling. The query stream reuse the hidden states from both forward and backward (content) stream as the key and value in attention. (b) The attention mask matrix used in our triple-stream self-attention. The above row is the attention matrix for the forward and backward stream and the below row is the attention matrix for the query stream. The cell in gray color means this position can not be attended.

3.2 Triple-Stream Self-Attention

Based on the Eqn 1 of our proposed SLM, model needs to maintain two states to collect forward and backward contexts for prediction, and we call these two states as the forward stream and backward stream respectively. To avoid information leakage, we additionally maintain an individual state for prediction and control it to only capture the dependency from the forward and backward streams. We name this state as the query stream. Therefore, we propose a novel Triple-Stream Self-Attention to update each stream, and the detailed design is described as following.

To fulfill our target, we choose Transformer [17] as our basic model, due to its flexibility in capturing global dependency. Assume the input sequence as $\{w_1, w_2, \cdots, w_n\}$ and its positions as $\{p_1, p_2, \cdots, p_n\}$, where w_i and p_i represents i-th token and its position in the sentence, and n is the token number. For the query stream, we only use the positions $\{p_1, p_2, \cdots, p_n\}$ as the input. For the forward and backward streams, we also maintain two individual states and both of them use the tokens plus its position (*i.e.*, $\{w_1 + p_1, w_2 + p_2, \cdots, w_n + p_n\}$) as the input. For the l-th layer calculation, we denote the forward and backward streams of the position i as $\vec{h}_i^l$ and $\overleftarrow{h}_i^l$, and they are updated as:

$$\vec{h}_i^l = \text{Attention}(\mathbf{Q} = \vec{h}_i^{l-1}, \mathbf{KV} = \vec{h}_{\leq i}^{l-1}; \theta), \tag{2}$$

$$\overleftarrow{h}_i^l = \text{Attention}(\mathbf{Q} = \overleftarrow{h}_i^{l-1}, \mathbf{KV} = \overleftarrow{h}_{\geq i}^{l-1}; \theta), \tag{3}$$

where $\text{Attention}(\cdot, \cdot)$ refers to the self-attention [17] in Transformer and $\mathbf{Q}$, $\mathbf{K}$, $\mathbf{V}$ denote the query, key and value in self-attention. Hence, $\vec{h}_i^l$ collects the information before the position i and itself, and $\overleftarrow{h}_i^l$ collects the information after the position i and itself. For the query stream, we denote q_i^l as its hidden states, and concatenate forward stream $\vec{h}_i^l$ and backward stream $\overleftarrow{h}_i^l$ of the current layer to be the key/value of query stream. So q_i^l is updated as:

$$q_i^l = \text{Attention}(\mathbf{Q} = q_i^{l-1}, \mathbf{KV} = \left[\vec{h}_{<i}^{l-1}, \overleftarrow{h}_{>i}^{l-1}\right]; \theta). \tag{4}$$

Here we find that q_i^l is required to only capture the dependency from the forward stream before the position i and the backward stream after the position i, rather than itself. Due to such design, the query stream is able to capture bidirectional context for estimating token probability and avoid information leakage, which is more effective than CLM in using context. More importantly, our triple-stream self-attention enables model to predict the probability of all tokens in a sentence within

LM	Model	Cost	Context	Scenario
CLM	GPT [9]	$\times 1$	forward	NLG
MLM	BERT [5]	$\times$n	bidirectional	NLU
Bi-LM	ELMO [34]	$\times 2$	forward + backward	NLU
SLM	Transcormer	$\times 3$	bidirectional	Scoring

Table 1: Comparisons between SLM and other LMs. We assume all LMs adopt the same architectures (*e.g.*, Transformer). The "Cost" column means the relative computations compared with CLM when calculating a sentence with n tokens. The "Context" column means the contextual information usage for prediction.

a single forward pass, which demonstrates more efficiency than MLM. Figure 2 presents the detailed design of our triple-stream self-attention. We can find that in the query stream, the masked matrix is like a sliding window to control each token to view its previous states in forward stream and its future states in backward stream. And that is why we name our model as sliding language modeling.

3.3 Discussion

To better understand our SLM, we analyze the advantages of our SLM over other LMs. The comparisons are listed in Table 1. We select three representative LMs for comparisons, which are CLM (BERT), MLM (GPT) and bidirectional LM (Bi-LM, used in ELMO [34] [2]) respectively. From Table 1, we have the following observations:

1. When compared with CLM, our SLM requires 3 $\times$ computations. However, our SLM can fully use the whole sentence information for prediction while CLM can only use unidirectional information. Even scaling CLM as 3 $\times$ parameters, it still can not use bidirectional context for prediction. This also demonstrates the effectiveness of our SLM in using context.
2. MLM is powerful at extracting bidirectional context but it needs $n\times$ inferences to calculate the whole sentence information limited to its masked prediction. Our SLM just needs a single inference and uses bidirectional information for prediction with only $3\times$ computations. Especially, our SLM shows higher efficiency compared with MLM when n is large.
3. Bi-LM can also extract forward and backward contextual information, but it just simply concatenates the forward and backward features for the final prediction, without any interactions. Instead, our SLM can iteratively fuse the bidirectional information thanks to our triple-stream self-attention mechanism.

Overall, the design of SLM is dedicated for sentence scoring, while CLM prefers NLG tasks and MLM/Bi-LM prefer NLU tasks.

4 Experiments

4.1 Experimental Setup

We adopt Transformer [17] as the backbone network. Following previous works [5], we adopt a base setting for our model as Transcormer$_{base}$ with 110M parameters, that consists of 12 transformer layers and each layer has 768 hidden size and 12 attention heads. During the pre-training, we use wikipedia plus bookcorpus (16GB) as the training corpus, to be consistent with previous works [5]. Our model is trained at the sentence-level (*i.e.*, one sentence per sample). We choose Adam [36] as the default optimizer with learning rate of $5e-4$, $\beta_1 = 0.9$, $\beta_2 = 0.98$ and $\epsilon = 1e-6$, and weight decay is set as 0.01. The learning rate warms up over the first 10,000 steps and then linearly decays. We set the batch size as 8192 tokens per batch, and the training step is 125,000 steps. We use 32 NVIDIA Tesla 32GB GPUs, with FP16 speedup. The total training needs 5.5 days for Transcormer$_{base}$. For more experimental settings, please refer to the Appendix.

[2] ELMO pre-trains a left-to-right and a right-to-left LSTM [35] and concatenates the outputs of each last unidirectional LSTM layer for prediction.

Model	IWSLT De	Es	It	Nl	Pl	Ro	Ru	Tr	WMT DeEn
Oracle	41.80	48.69	41.89	44.38	27.90	46.01	29.60	27.25	39.17
Baseline	34.77	41.20	34.95	37.73	22.67	38.73	24.21	21.65	32.54
CLM (GPT)	34.96	41.39	35.14	38.08	22.91	39.03	24.62	22.14	32.88
MLM (BERT)	35.14	41.54	**35.54**	38.14	23.00	39.21	24.65	22.36	33.07
Bi-LM (ELMO)	35.10	41.52	35.21	38.03	23.09	39.07	24.53	21.91	32.90
SLM (Transcormer)	**35.24**	**41.86**	35.52	**38.45**	**23.29**	**39.34**	**24.69**	**22.41**	**33.10**

Table 2: Reranking results on IWSLT and WMT tasks, and all LMs have the same model architecture as Transcormer. All results are reported in BLEU score. All LMs are pre-trained in the wikipedia + bookcorpus (16GB) with the same optimization. The last row is the oracle score from the generated candidates.

4.2 Experiments on Neural Machine Translation

We choose IWSLT14 dataset [37], which includes multiple translation tasks from low-resource languages to English, and WMT14 English-German dataset [3] for evaluation. We adopt Transformer [17] as the machine translation model with 6-6 transformer layers to generate multiple candidates for reranking, with a beam size of 10. The hidden size and attention head are set as 512/1024 and 8/16 for IWSLT and WMT tasks respectively. During the reranking, we combine the original score produced by the machine translation model and the LM score with a hyper-parameter λ, by following previous experiences [3, 18]. The hyper-parameter λ is tuned on dev set with a range of $[0.0, 2.0]$, and then select the best λ to evaluate the test set. The results are reported in Table 2, in terms of BLEU [38]. Each task will be evaluated by five times based on different pre-trained checkpoints, and report the mean value, with a variance of 0.05. From Table 2, we have the following observations:

- Our Transcormer$_{base}$ can obtain better performance than CLM and Bi-LM [4] in both low-resource IWSLT and large-scale WMT tasks, which indicates the importance of bidirectional context for sentence scoring, and further validate the ability of our SLM in utilizing bidirectional information.
- When compared with MLM, our Transcormer$_{base}$ also achieves comparable results. Considering that the computation of MLM for scoring is linear to the input length and needs n inference passes, our SLM show higher efficiency with only $3\times$ computations in a single pass to maintain three streams.

Overall, all of the results reveal that our model is more effective in using contextual information for probability estimation and more efficient with only a single forward pass, especially for long sentences.

Model	dev-clean	dev-other	test-clear	test-other
Baseline	2.80	6.90	3.06	7.05
CLM (GPT)	2.47	6.13	2.73	6.33
MLM (BERT)	2.30	5.65	2.59	5.90
Bi-LM (ELMO)	2.41	5.92	2.63	6.12
SLM (Transcormer)	**2.23**	**5.54**	**2.49**	**5.72**
Oracle	1.45	4.23	1.59	4.19

Table 3: Reranking results on LibrSpeech dataset. All results are reported in WER.

4.3 Experiments on Automatic Speech Recognition

We choose LibrSpeech [39] to evaluate the performance of our model for reranking on ASR task. We train a Conformer model [40] on LibrSpeech, which has 12 encoder layers and 6 decoder layers with

[3] Here we only evaluate German→English direction as our model is trained on English domain.

[4] Here, we replace LSTM as Transformer in Bi-LM to keep the consistence in architecture.

512 hidden size and 8 attention heads, and the beam size is set as 10. In addition, we use SpecAug [41] as a data augmentation technology to further improve the accuracy of ASR system. We use word error rate (WER) to evaluate the performance of ASR tasks. We follow the same tuning technique used in NMT tasks for hyper-parameter λ, but with a larger range as $[0.0, 5.0]$. The results are reported in Table 3. From Table 3, we find that our model can give nearly 20% relative improvements over the baseline and also outperform other LMs, including CLM, MLM and Bi-LM. The results on ASR task further demonstrates the generalization and effectiveness of our SLM in sentence scoring.

5 Analyses

In this section, we conduct some method analyses on our proposed SLM and CLM/MLM. Besides, we also provide more analyses in Appendix.

Model	IWSLT		WMT	LibriSpeech	
	De	Es	DeEn	dev-clean	dev-other
Baseline	34.77	41.20	32.54	2.80	6.90
CLM (GPT)	34.96	41.39	32.88	2.47	6.13
SLM (Transcormer$_{small}$)	**35.05**	**41.58**	**32.94**	**2.48**	**5.95**

Table 4: Comparisons between SLM (Transcormer) and CLM (GPT) under the same computation.

5.1 Comparable Computation between CLM and SLM

As aforementioned, our SLM needs $3\times$ computations when compared with CLM. To make a fair comparison, we also pre-train a Transcormer$_{small}$ with 34M parameters in total, which consists 6 transformer layers and each layer has 512 hidden size and 8 attention heads. Hence, our Transcormer$_{small}$ has approximately $\frac{1}{3}$ parameters of Transcormer$_{base}$, and has the similar computations as GPT$_{base}$. We conduct experiments on three NMT tasks and an ASR task (LibriSpeech dataset) for comparisons and the results are listed in Table 4. We can find that even under the same computation, our model still outperforms CLM and this result further validates the necessity of using bidirectional context for sentence scoring. Besides, considering that our Transcormer$_{small}$ has fewer parameters, our model is also friendly to the device deployments (*e.g.*, CPU).

Model	Cost	PPL	dev-clean	dev-other	test-clear	test-other
Baseline	-	-	2.80	6.90	3.06	7.05
MLM ($k=1$)	$\times n$	4.26	2.30	5.65	2.59	5.90
MLM ($k=2$)	$\times \lceil n/2 \rceil$	8.41	2.41	5.87	2.70	6.20
MLM ($k=3$)	$\times \lceil n/3 \rceil$	11.58	2.60	5.95	2.87	6.41
MLM ($k=\frac{n}{3}$)	$\times 3$	-	2.75	6.71	2.98	6.93
MLM ($k=\frac{n}{2}$)	$\times 2$	-	2.80	6.80	3.01	6.99
SLM	$\times 3$	**3.85**	**2.23**	**5.54**	**2.49**	**5.72**

Table 5: Comparisons of sampling different k tokens for prediction on MLM. We choose ASR reranking task on LibriSpeech dataset to evaluate the results and also report PPL on a subset of sentences with same length ($n = 20$).

5.2 Varying Numbers of Forward Passes in MLM

As mentioned above, MLM needs to forward n passes as each time only mask one token. So what will happen if we allow each pass to mask more tokens? Therefore, we design experiments that enforces MLM to forward k tokens for prediction so that it only needs $\lceil n/k \rceil$ passes, and investigate the effect of different k. For a certain k, we randomly split the sentence as $\lceil n/k \rceil$ sets, and each time

masks one subset for prediction. The comparisons are listed in Table 5. We find that using larger k will severely harm the performance, even if it can reduce the number of inference passes. When k is set as $\frac{n}{3}$, which is equal to the cost of our SLM, it can hardly give any improvements over the baseline. We think that masking more tokens in the sentence will make it more difficult to estimate the token probability at a time. These comparisons also highlight the efficiency and effectiveness of our SLM for sentence scoring.

5.3 Sentence Scoring Quality at Each Position

To better analyze the effectiveness of bidirectional context to estimate token probability, we count the cross-entropy loss of each position in sentences for each LM. Specifically, we sample a subset of sentences as $\mathcal{S}$, and each sentence has the same token number n (here n is set as 20). For each position i, we count the average cross-entropy over the sampled subset $\mathcal{S}$ based on the output probability of each LM. The results are displayed in Figure 3. We can find that: 1) For CLM, the cross-entropy loss is higher at the first several positions and gradually decreases for the subsequent positions but is still higher than MLM and SLM, which indicates that only using undirectional information is not enough to measure the sentence score precisely. 2) SLM can almost obtain similar loss as MLM at each position. Considering SLM just needs a single pass while MLM needs n passes, this phenomenon further validates the superiority and efficiency of SLM in scoring sentences.

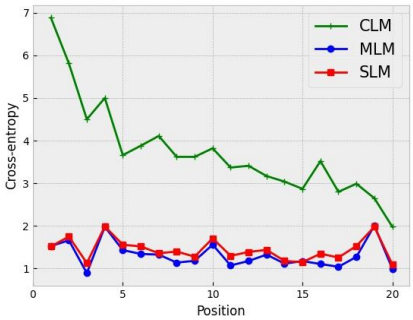

Figure 3: The average cross-entropy loss of each LM at each position. MLM uses n passes to predict each position.

6 Conclusion

In this paper, we propose Transcormer, a Transformer with a novel sliding language modeling for sentence scoring. Specifically, our SLM is able to produce the probability of each token over the whole sentence within a single forward pass, and utilizes bidirectional context for prediction, and thus inherents the advantages of CLM and MLM and avoids their deficiencies. To the best of our knowledge, the proposed Transcormer is the first pre-trained language model tailored for sentence scoring. Experimental results on multiple datasets demonstrate the effectiveness of our Transcormer in computing sentence score for reranking tasks.

Besides, we summarize some potential directions of our Transcormer and SLM as the future works:

- Currently our Transcormer is only conducted on the English domain under the base setting, due to limited computation. We expect to develop large-scale Transcormer and use different language domains or multilingual data for training in the future.
- We design sliding language modeling for sentence scoring, and our experiments are mainly on reranking task. However, based on the characteristics of our SLM, we believe our model can also be used for other scenarios (*e.g.*, correction, data selection), and we will explore the specific fine-tuning techniques when applying our SLM on different downstream tasks.
- Besides, our Transcormer mainly pre-trains SLM on an encoder framework. However, our SLM is not limited to model structure. For example, SLM can be easily extended to encoder-decoder framework based on paired data. Therefore, we also expect to explore the possibility of using SLM on different frameworks.
- Although our paper mainly focuses on text data, we want to highlight that SLM can also be extended to other different modalities with sequential characteristic (*e.g.*, image, speech and time series data). Consequently, how to apply SLM to other modalities is also a valuable topic in the future.

References

[1] Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. Facebook fair's WMT19 news translation task submission. In *Proceedings of the Fourth Conference on Machine Translation, WMT 2019, Florence, Italy, August 1-2, 2019 - Volume 2: Shared Task Papers, Day 1*, pages 314–319. Association for Computational Linguistics, 2019.

[2] Yingce Xia, Xu Tan, Fei Tian, Fei Gao, Di He, Weicong Chen, Yang Fan, Linyuan Gong, Yichong Leng, Renqian Luo, Yiren Wang, Lijun Wu, Jinhua Zhu, Tao Qin, and Tie-Yan Liu. Microsoft research asia's systems for WMT19. In *Proceedings of the Fourth Conference on Machine Translation, WMT 2019, Florence, Italy, August 1-2, 2019 - Volume 2: Shared Task Papers, Day 1*, pages 424–433. Association for Computational Linguistics, 2019.

[3] Kyra Yee, Yann N. Dauphin, and Michael Auli. Simple and effective noisy channel modeling 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 2019, Hong Kong, China, November 3-7, 2019*, pages 5695–5700. Association for Computational Linguistics, 2019.

[4] Yuchen Liu, Long Zhou, Yining Wang, Yang Zhao, Jiajun Zhang, and Chengqing Zong. A comparable study on model averaging, ensembling and reranking in NMT. In *Natural Language Processing and Chinese Computing - 7th CCF International Conference, NLPCC 2018, Hohhot, China, August 26-30, 2018, Proceedings, Part II*, volume 11109 of *Lecture Notes in Computer Science*, pages 299–308. Springer, 2018.

[5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186, 2019.

[6] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019.

[7] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A lite BERT for self-supervised learning of language representations. In *8th International Conference on Learning Representations, ICLR 2020*, Addis Ababa, Ethiopia, 2020. OpenReview.net.

[8] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018.

[9] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.

[10] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, pages 1–16, virtual, 2020. Curran Associates, Inc.

[11] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*, volume 32, pages 5753–5763, Vancouver Convention Centre, Vancouver, Canada, 2019.

[12] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. Mpnet: Masked and permuted pre-training for language understanding. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020*, virtual, 2020. Curran Associates Inc.

[13] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MASS: Masked sequence to sequence pre-training for language generation. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 5926–5936, Long Beach Convention Center, Long Beach, 2019. PMLR.

[14] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. Unified language model pre-training for natural language understanding and generation. In *NeurIPS*, pages 13042–13054, Vancouver Convention Centre, Vancouver, Canada, 2019. Curran Associates, Inc.

[15] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *ACL*, pages 7871–7880, Online, 2020. Association for Computational Linguistics.

[16] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.

[17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008, Long Beach Convention Center, Long Beach, 2017. Curran Associates, Inc.

[18] Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. Masked language model scoring. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2699–2712. Association for Computational Linguistics, 2020.

[19] Alex Wang and Kyunghyun Cho. BERT has a mouth, and it must speak: BERT as a markov random field language model. *CoRR*, abs/1902.04094, 2019.

[20] Joonbo Shin, Yoonhyung Lee, and Kyomin Jung. Effective sentence scoring method using BERT for speech recognition. In *Proceedings of The 11th Asian Conference on Machine Learning, ACML 2019, 17-19 November 2019, Nagoya, Japan*, volume 101 of *Proceedings of Machine Learning Research*, pages 1081–1093. PMLR, 2019.

[21] Shih-Hsuan Chiu and Berlin Chen. Innovative bert-based reranking language models for speech recognition. In *IEEE Spoken Language Technology Workshop, SLT 2021, Shenzhen, China, January 19-22, 2021*, pages 266–271. IEEE, 2021.

[22] Yue Weng, Sai Sumanth Miryala, Chandra Khatri, Runze Wang, Huaixiu Zheng, Piero Molino, Mahdi Namazifar, Alexandros Papangelis, Hugh Williams, Franziska Bell, and Gökhan Tür. Joint contextual modeling for ASR correction and language understanding. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 6349–6353. IEEE, 2020.

[23] Shankar Kumar and William J. Byrne. Minimum bayes-risk decoding for statistical machine translation. In *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, HLT-NAACL 2004, Boston, Massachusetts, USA, May 2-7, 2004*, pages 169–176, 2004.

[24] Raphael Shu and Hideki Nakayama. Later-stage minimum bayes-risk decoding for neural machine translation. *CoRR*, abs/1704.03169, 2017.

[25] Libin Shen, Anoop Sarkar, and Franz Josef Och. Discriminative reranking for machine translation. In *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, HLT-NAACL 2004, Boston, Massachusetts, USA, May 2-7, 2004*, pages 177–184, 2004.

[26] Franz Josef Och. Minimum error rate training in statistical machine translation. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics, 7-12 July 2003, Sapporo Convention Center, Sapporo, Japan*, pages 160–167, 2003.

[27] Yukun Ma, Erik Cambria, and Benjamin Bigot. ASR hypothesis reranking using prior-informed restricted boltzmann machine. In *Computational Linguistics and Intelligent Text Processing - 18th International Conference, CICLing 2017, Budapest, Hungary, April 17-23, 2017, Revised Selected Papers, Part I*, volume 10761, pages 503–514. Springer, 2017.

[28] Michael Auli and Jianfeng Gao. Decoder integration and expected BLEU training for recurrent neural network language models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 2: Short Papers*, pages 136–142. The Association for Computer Linguistics, 2014.

[29] Terumasa Ehara. In *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation: 5th Workshop on Asian Translation, WAT@PACLIC 2018, Hong Kong, December 1-3, 2018*. Association for Computational Linguistics, 2018.

[30] Yuguang Wang, Shanbo Cheng, Liyang Jiang, Jiajun Yang, Wei Chen, Muze Li, Lin Shi, Yanfeng Wang, and Hongtao Yang. Sogou neural machine translation systems for WMT17. In *Proceedings of the Second Conference on Machine Translation, WMT 2017, Copenhagen, Denmark, September 7-8, 2017*, pages 410–415. Association for Computational Linguistics, 2017.

[31] Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. Prophetnet: Predicting future n-gram for sequence-to-sequence pre-training. In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 2401–2410. Association for Computational Linguistics, 2020.

[32] Dongling Xiao, Han Zhang, Yu-Kun Li, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. ERNIE-GEN: an enhanced multi-flow pre-training and fine-tuning framework for natural language generation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 3997–4003. ijcai.org, 2020.

[33] Jungo Kasai, James Cross, Marjan Ghazvininejad, and Jiatao Gu. 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, 2020.

[34] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics, 2018.

[35] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997.

[36] Diederik P. Kingma and Jimmy Ba. 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*, 2015.

[37] Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. Report on the 11th IWSLT evaluation campaign. In *Proceedings of the 11th International Workshop on Spoken Language Translation: Evaluation Campaign*, pages 2–17, Lake Tahoe, California, dec 2014.

[38] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL, 2002.

[39] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2015, South Brisbane, Queensland, Australia, April 19-24, 2015*, pages 5206–5210. IEEE, 2015.

[40] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. Conformer: Convolution-augmented transformer for speech recognition. In *INTERSPEECH*, 2020.

[41] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. Specaugment: A simple data augmentation method for automatic speech recognition. In *INTERSPEECH*, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Please see Section refsubsec:nmt and Appendix.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Please see Section 4.1
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]