

Hi-Transformer: Hierarchical Interactive Transformer for Efficient and Effective Long Document Modeling

Chuhan Wu[†] Fangzhao Wu[‡] Tao Qi[†] Yongfeng Huang[†]

[†]Department of Electronic Engineering & BNRist, Tsinghua University, Beijing 100084, China

[‡]Microsoft Research Asia, Beijing 100080, China

{wuchuhan15, wufangzhao, taoqi.qt}@gmail.com
yfhuang@tsinghua.edu.cn

Abstract

Transformer is important for text modeling. However, it has difficulty in handling long documents due to the quadratic complexity with input text length. In order to handle this problem, we propose a hierarchical interactive Transformer (Hi-Transformer) for efficient and effective long document modeling. Hi-Transformer models documents in a hierarchical way, i.e., first learns sentence representations and then learns document representations. It can effectively reduce the complexity and meanwhile capture global document context in the modeling of each sentence. More specifically, we first use a sentence Transformer to learn the representations of each sentence. Then we use a document Transformer to model the global document context from these sentence representations. Next, we use another sentence Transformer to enhance sentence modeling using the global document context. Finally, we use hierarchical pooling method to obtain document embedding. Extensive experiments on three benchmark datasets validate the efficiency and effectiveness of Hi-Transformer in long document modeling.

1 Introduction

Transformer (Vaswani et al., 2017) is an effective architecture for text modeling, and has been an essential component in many state-of-the-art NLP models like BERT (Devlin et al., 2019; Radford et al., 2019; Yang et al., 2019; Wu et al., 2021). The standard Transformer needs to compute a dense self-attention matrix based on the interactions between each pair of tokens in text, where the computational complexity is proportional to the square of text length (Vaswani et al., 2017; Wu et al., 2020b). Thus, it is difficult for Transformer to model long documents efficiently (Child et al., 2019).

There are several methods to accelerate Transformer for long document modeling (Wu et al.,

2019; Kitaev et al., 2019; Wang et al., 2020; Qiu et al., 2020). One direction is using Transformer in a hierarchical manner to reduce sequence length, e.g., first learn sentence representations and then learn document representations from sentence representations (Zhang et al., 2019; Yang et al., 2020). However, the modeling of sentences is agnostic to the global document context, which may be suboptimal because the local context within sentence is usually insufficient. Another direction is using a sparse self-attention matrix instead of a dense one. For example, Beltagy et al. (2020) proposed to combine local self-attention with a dilated sliding window and sparse global attention. Zaheer et al. (2020) proposed to incorporate a random sparse attention mechanism to model the interactions between a random set of tokens. However, these methods cannot fully model the global context of document (Tay et al., 2020).

In this paper, we propose a hierarchical interactive Transformer (*Hi-Transformer*)¹ for efficient and effective long document modeling, which models documents in a hierarchical way to effectively reduce the complexity and at the same time can capture the global document context for sentence modeling. In *Hi-Transformer*, we first use a sentence Transformer to learn the representation of each sentence within a document. Next, we use a document Transformer to model the global document context from these sentence representations. Then, we use another sentence Transformer to further improve the modeling of each sentence with the help of the global document context. Finally, we use hierarchical pooling method to obtain the document representation. Extensive experiments are conducted on three benchmark datasets. The results show that *Hi-Transformer* is both efficient and effective in long document modeling.

¹<https://github.com/wuch15/HiTransformer>.




Figure 1: The architecture of *Hi-Transformer*.

2 Hi-Transformer

In this section, we introduce our hierarchical interactive Transformer (*Hi-Transformer*) approach for efficient and effective long document modeling. Its framework is shown in Fig. 1. It uses a hierarchical architecture that first models the contexts within a sentence, next models the document contexts by capturing the interactions between sentences, then employs the global document contexts to enhance sentence modeling, and finally uses hierarchical pooling techniques to obtain document embeddings. In this way, the input sequence length of each Transformer is much shorter than directly taking the word sequence in document as input, and the global contexts can be fully modeled. The details of *Hi-Transformer* are introduced as follows.

2.1 Model Architecture

Hi-Transformer mainly contains three modules, i.e., sentence context modeling, document context modeling and global document context-enhanced sentence modeling. The sentence-level context is first

modeled by a sentence Transformer. Assume a document contains M sentences, and the words in the i -th sentence are denoted as $[w_{i,1}, w_{i,2}, \dots, w_{i,K}]$ (K is the sentence length). We insert a “[CLS]” token (denoted as w^s) after the end of each sentence. This token is used to convey the contextual information within this sentence. The sequence of words in each sentence is first converted into a word embedding sequence via a word and position embedding layer. Denote the word embedding sequence for the i -th sentence as $[e_{i,1}, e_{i,2}, \dots, e_{i,K}, e^s]$. Since sentence length is usually short, we apply a sentence Transformer to each sentence to fully model the interactions between the words within this sentence. It takes the word embedding sequence as the input, and outputs the contextual representations of words, which are denoted as $[h_{i,1}, h_{i,2}, \dots, h_{i,K}, h_i^s]$. Specially, the representation h_i^s of the “[CLS]” token is regarded as the sentence representation.

Next, the document-level context is modeled by a document Transformer from the representations of the sentences within this document. Denote the

embedding sequence of sentences in this document as $[\mathbf{h}_1^s, \mathbf{h}_2^s, \dots, \mathbf{h}_M^s]$. We add a sentence position embedding (denoted as \mathbf{p}_i for the i -th sentence) to the sentence representations to capture sentence orders. We then apply a document Transformer to these sentence representations to capture the global context of document, and further learn document context-aware sentence representations, which are denoted as $[\mathbf{r}_1^s, \mathbf{r}_2^s, \dots, \mathbf{r}_M^s]$.

Then, we use the document context-aware sentence representations to further improve the sentence context modeling by propagating the global document context to each sentence. Motivated by (Guo et al., 2019), we apply another sentence Transformer to the hidden word representations and the document-aware sentence representation for each sentence. It outputs a document context-aware word representation sequence for each sentence, which is denoted as $[\mathbf{d}_{i,1}, \mathbf{d}_{i,2}, \dots, \mathbf{d}_{i,K}, \mathbf{d}_i^s]$. In this way, the contextual representations of words can benefit from both local sentence context and global document context.

By stacking multiple layers of *Hi-Transformer*, the contexts within a document can be fully modeled. Finally, we use hierarchical pooling (Wu et al., 2020a) techniques to obtain the document embedding. We first aggregate the document context-aware word representations in each sentence into a global context-aware sentence embedding \mathbf{s}_i , and then aggregate the global context-aware embeddings of sentence within a document into a unified document embedding \mathbf{d} , which is further used for downstream tasks.

2.2 Efficiency Analysis

In this section, we provide some discussions on the computational complexity of *Hi-Transformer*. In sentence context modeling and document context propagation, the total computational complexity is $O(M \cdot K^2 \cdot d)$, where M is sentence number with a document, K is sentence length, and d is the hidden dimension. In document context modeling, the computational complexity is $O(M^2 \cdot d)$. Thus, the total computational cost is $O(M \cdot K^2 \cdot d + M^2 \cdot d)$.² Compared with the standard Transformer whose computational complexity is $O(M^2 \cdot K^2 \cdot d)$, *Hi-Transformer* is much more efficient.

²Note that *Hi-Transformer* can be combined with other existing techniques of efficient Transformer to further improve the efficiency for long document modeling.

3 Experiments

3.1 Datasets and Experimental Settings

Our experiments are conducted on three benchmark document modeling datasets. The first one is Amazon Electronics (He and McAuley, 2016) (denoted as Amazon), which is for product review rating prediction.³ The second one is IMDB (Diao et al., 2014), a widely used dataset for movie review rating prediction.⁴ The third one is the MIND dataset (Wu et al., 2020c), which is a large-scale dataset for news intelligence.⁵ We use the content based news topic classification task on this dataset. The detailed dataset statistics are shown in Table 1.

In our experiments, we use the 300-dimensional pre-trained Glove (Pennington et al., 2014) embeddings for initializing word embeddings. We use two *Hi-Transformers* layers in our approach and two Transformer layers in other baseline methods.⁶ We use attentive pooling (Yang et al., 2016) to implement the hierarchical pooling module. The hidden dimension is set to 256, i.e., 8 self-attention heads in total and the output dimension of each head is 32. Due to the limitation of GPU memory, the input sequence lengths of vanilla Transformer and its variants for long documents are 512 and 2048, respectively. The dropout (Srivastava et al., 2014) ratio is 0.2. The optimizer is Adam (Bengio and LeCun, 2015), and the learning rate is $1e-4$. The maximum training epoch is 3. The models are implemented using the Keras library with Tensorflow backend. The GPU we used is GeForce GTX 1080 Ti with a memory of 11 GB. We use accuracy and macro-F scores as the performance metrics. We repeat each experiment 5 times and report both average results and standard deviations.

3.2 Performance Evaluation

We compare *Hi-Transformer* with several baselines, including: (1) *Transformer* (Vaswani et al., 2017), the vanilla Transformer architecture; (2) *Longformer* (Beltagy et al., 2020), a variant of Transformer with local and global attention for long documents; (3) *BigBird* (Zaheer et al., 2020), extending *Longformer* with random attention; (4) *HI-BERT* (Zhang et al., 2019), using Transformers

³<https://jmcauley.ucsd.edu/data/amazon/>

⁴<https://github.com/nihalb/JMARS>

⁵<https://msnews.github.io/>

⁶We also tried more Transformer layers for baseline methods but do not observe significant performance improvement in our experiments.

Dataset	#Train	#Val	#Test	Avg. #word	Avg. #sent	#Class
Amazon	40.0k	5.0k	5.0k	133.38	6.17	5
IMDB	108.5k	13.6k	13.6k	385.70	15.29	10
MIND	128.8k	16.1k	16.1k	505.46	25.14	18

Table 1: Statistics of datasets.

Methods	Amazon		IMDB		MIND	
	Accuracy	Macro-F	Accuracy	Macro-F	Accuracy	Macro-F
Transformer	65.23±0.38	42.23±0.37	51.98±0.48	42.76±0.49	80.96±0.22	59.97±0.24
Longformer	65.35±0.44	42.45±0.41	52.33±0.40	43.51±0.42	81.42±0.25	62.68±0.26
BigBird	66.05±0.48	42.89±0.46	52.87±0.51	43.79±0.50	81.81±0.29	63.44±0.31
HI-BERT	66.56±0.32	42.65±0.34	52.96±0.46	43.84±0.46	81.89±0.23	63.63±0.20
Hi-Transformer	67.24±0.35	43.69±0.32	53.78±0.49	44.54±0.47	82.51±0.25	64.22±0.22

Table 2: The results of different methods on different datasets.

Method	Complexity
Transformer	$O(M^2 \cdot K^2 \cdot d)$
Longformer	$O(T \cdot M \cdot K \cdot d)$
BigBird	$O(T \cdot M \cdot K \cdot d)$
HI-BERT	$O(M \cdot K^2 \cdot d + M^2 \cdot d)$
Hi-Transformer	$O(M \cdot K^2 \cdot d + M^2 \cdot d)$

Table 3: Complexity of different methods. K is sentence length, M is the number of sentences in a document, T is the number of positions for sparse attention, and d is the hidden dimension.

at both word and sentence levels. The results of these methods on the three datasets are shown in Table 2. We find that Transformers designed for long documents like *Hi-Transformer* and *BigBird* outperform the vanilla Transformer. This is because vanilla Transformer cannot handle long sequence due to the restriction of computation resources, and truncating the input sequence leads to the loss of much useful contextual information. In addition, *Hi-Transformer* and *HI-BERT* outperform *Longformer* and *BigBird*. This is because the sparse attention mechanism used in *Longformer* and *BigBird* cannot fully model the global contexts within a document. Besides, *Hi-Transformer* achieves the best performance, and the t-test results show the improvements over baselines are significant. This is because *Hi-Transformer* can incorporate global document contexts to enhance sentence modeling.

We also compare the computational complexity of these methods in Table 3. The complexity of *Hi-Transformer* is much less than the vanilla Transformer and is comparable with other Transformer variants designed for long documents. These re-

sults indicate the efficiency and effectiveness of *Hi-Transformer*.


3.3 Model Effectiveness

Nest, we verify the effectiveness of the global document contexts for enhancing sentence modeling in *Hi-Transformer*. We compare *Hi-Transformer* and its variants without global document contexts in Fig. 2. We find the performance consistently declines when the global document contexts are not encoded into sentence representations. This is because the local contexts within a single sentence may be insufficient for accurate sentence modeling, and global contexts in the entire document can provide rich complementary information for sentence understanding. Thus, propagating the document contexts to enhance sentence modeling can improve long document modeling.


3.4 Influence of Text Length

Then, we study the influence of text length on the model performance and computational cost. Since the documents in the MIND dataset are longest, we conduct experiments on MIND to compare the model performance as well as the training time per layer of *Transformer* and *Hi-Transformer* under different input text length⁷, and the results are shown in Fig. 3. We find the performance of both methods improves when longer text sequences are used. This is intuitive because more information can be incorporated when longer text is input to the model for document modeling. However, the computational cost of *Transformer* grows very fast,


⁷The maximum length of *Transformer* is 512 due to GPU memory limitation.



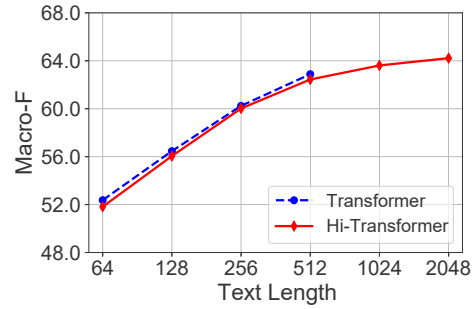
(a) Amazon.




(a) Accuracy.



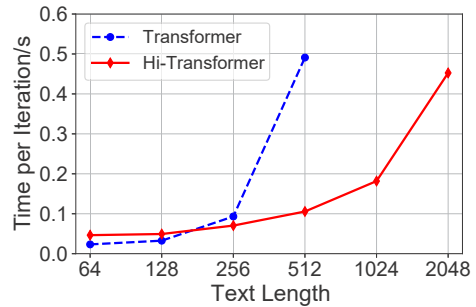
(b) IMDB.



(b) Macro-F.



(c) MIND.



(c) Training time per layer.

Figure 2: Effectiveness of global document context propagation in *Hi-Transformer*.

which limits its maximal input text length. Different from *Transformer*, *Hi-Transformer* is much more efficient and meanwhile can achieve better performance with longer sequence length. These results further verify the efficiency and effectiveness of *Hi-Transformer* in long document modeling.

4 Conclusion

In this paper, we propose a *Hi-Transformer* approach for both efficient and effective long document modeling. It incorporates a hierarchical architecture that first learns sentence representations and then learns document representations. It can effectively reduce the computational complexity and meanwhile be aware of the global document

Figure 3: Influence of input text length on performance and training time on the MIND dataset.

contexts in sentence modeling to help understand document content accurately. Extensive experiments on three benchmark datasets validate the efficiency and effectiveness of *Hi-Transformer* in long document modeling.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant numbers U1936216, U1936208, U1836204, and U1705261. We are grateful to Xing Xie, Shaoyu Zhou, Dan Shen, and Zhisong Wang for their insightful comments and suggestions on this work.

References

- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Yoshua Bengio and Yann LeCun. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186.
- Qiming Diao, Minghui Qiu, Chao-Yuan Wu, Alexander J Smola, Jing Jiang, and Chong Wang. 2014. Jointly modeling aspects, ratings and sentiments for movie recommendation (jmars). In *KDD*, pages 193–202.
- Qipeng Guo, Xipeng Qiu, Pengfei Liu, Yunfan Shao, Xiangyang Xue, and Zheng Zhang. 2019. Star-transformer. In *NAACL-HLT*, pages 1315–1325.
- Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *WWW*, pages 507–517.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2019. Reformer: The efficient transformer. In *ICLR*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Jiezhong Qiu, Hao Ma, Omer Levy, Wen-tau Yih, Sinong Wang, and Jie Tang. 2020. Blockwise self-attention for long document understanding. In *EMNLP: Findings*, pages 2555–2565.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *JMLR*, 15(1):1929–1958.
- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2020. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Sinong Wang, Belinda Li, Madian Khabza, Han Fang, and Hao Ma. 2020. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*.
- Chuhan Wu, Fangzhao Wu, and Yongfeng Huang. 2021. DA-transformer: Distance-aware transformer. In *NAACL-HLT*, pages 2059–2068.
- Chuhan Wu, Fangzhao Wu, Tao Qi, Xiaohui Cui, and Yongfeng Huang. 2020a. Attentive pooling with learnable norms for text representation. In *ACL*, pages 2961–2970.
- Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020b. Improving attention mechanism with query-value interaction. *arXiv preprint arXiv:2010.03766*.
- Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, et al. 2020c. Mind: A large-scale dataset for news recommendation. In *ACL*, pages 3597–3606.
- Zhanghao Wu, Zhijian Liu, Ji Lin, Yujun Lin, and Song Han. 2019. Lite transformer with long-short range attention. In *ICLR*.
- Liu Yang, Mingyang Zhang, Cheng Li, Michael Bendersky, and Marc Najork. 2020. Beyond 512 tokens: Siamese multi-depth transformer-based hierarchical encoder for long-form document matching. In *CIKM*, pages 1725–1734.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*, pages 5753–5763.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *NAACL-HLT*, pages 1480–1489.
- Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *arXiv preprint arXiv:2007.14062*.
- Xingxing Zhang, Furu Wei, and Ming Zhou. 2019. Hi-bert: Document level pre-training of hierarchical bidirectional transformers for document summarization. In *ACL*, pages 5059–5069.