

Big Bidirectional Insertion Representations for Documents

Lala Li

Google Research, Brain Team
lala@google.com

William Chan

Google Research, Brain Team
williamchan@google.com

Abstract

The Insertion Transformer is well suited for long form text generation due to its parallel generation capabilities, requiring $O(\log_2 n)$ generation steps to generate n tokens. However, modeling long sequences is difficult, as there is more ambiguity captured in the attention mechanism. This work proposes the Big Bidirectional Insertion Representations for Documents (Big BIRD), an insertion-based model for document-level translation tasks. We scale up the insertion-based models to long form documents. Our key contribution is introducing sentence alignment via sentence-positional embeddings between the source and target document. We show an improvement of +4.3 BLEU on the WMT'19 English→German document-level translation task compared with the Insertion Transformer baseline.

1 Introduction

Recently, insertion-based models (Stern et al., 2019; Welleck et al., 2019; Gu et al., 2019; Chan et al., 2019) have been introduced for text generation. Unlike traditional autoregressive left-to-right models (Cho et al., 2014; Sutskever et al., 2014; Vaswani et al., 2017), insertion-based models are not restricted to generating text sequences in a serial left-to-right manner, but these models are endowed with the capabilities of parallel generation. More specifically, Stern et al. (2019); Chan et al. (2019) showed that we can teach neural nets to generate text to follow a balanced binary tree order. An autoregressive left-to-right model would require $O(n)$ generation steps to generate n tokens, whereas the Insertion Transformer (Stern et al., 2019) and KERMIT (Chan et al., 2019) following a balanced binary tree policy requires only $O(\log_2 n)$ generation steps to generate n tokens. This is especially important for long-

form text generation, for example, Document-Level Machine Translation.

Document-Level Machine Translation is becoming an increasingly important task. Recent research suggests we are nearing human-level parity for sentence-level translation in certain domains (Hassan et al., 2018), however, we lag significantly behind in document-level translation (Lubli et al., 2018). Various papers have proposed incorporating context for document-level translation (Junczys-Dowmunt, 2019), which has been shown to improve translation quality. There are two primary methods to include context in a document-level machine translation model compared to a sentence-level translation model.

1. **Source Contextualization.** We can include source context, wherein when we generate the target sentence, we can condition on the corresponding source sentence and its neighbours, or even the whole source document. This allows the target sentence to be contextualized to the source document.
2. **Target Contextualization.** We can include target context, wherein when we generate the target sentence, we can condition on all the target tokens generated thus far in the whole document. This allows the target sentence to be contextualized to other target sentences.

Target contextualization is especially difficult in an autoregressive left-to-right model (i.e., Transformer (Vaswani et al., 2017)), the model must generate the whole document in linear fashion, which would be prohibitively expensive costing $O(n)$ iterations to generate n tokens. Additionally, the model is unable to model bidirectional context, since the text is always generated in a left-to-right manner. Some prior work have focused on utilizing block coordinate descent like algorithms during inference (Maruf and Haffari, 2018), however

this adds complexity and additional runtime cost during inference.

Insertion-based models, for example, the Insertion Transformer (Stern et al., 2019) is one potential solution. The Insertion Transformer can generate text following a balanced binary tree order. It requires $O(\log_2 n)$ iterations to generate n tokens, offering significant inference time advantages over a serial generation model. The source document is naturally fully conditioned on, which provides full source contextualization. Additionally, the generation order offers bidirectional contextualization, permitting target contextualization that is not solely on a left-to-right basis.

In this paper, we present Big Bidirectional Insertion Representations for Documents (Big BIRD). We address the limitations of scaling up the Insertion Transformer to document-level machine translation. We present a model that can model long-form documents with thousands of tokens in a fully contextualized manner.

2 Big BIRD

In this section, we present Big Bidirectional Representations for Documents (Big BIRD). Big BIRD is an extension of the Insertion Transformer (Stern et al., 2019), scaling up from sentences to documents. The key contributions are 1) extending the context window size to cover a document, and 2) informing the model of sentence positional information, which are aligned between source and target sentences.

Insertion Transformer. In the Insertion Transformer (Stern et al., 2019), sequences are generated via insertion operations. In the context of Machine Translation, there is a source canvas x and a target canvas y , where the target canvas is updated at each iteration via inserting one token at each plausible location. At time t during training, a hypothesis target canvas \hat{y}_t must be a subsequence of the final output. For example, if the final output is $[A, B, C, D, E]$, then $\hat{y}_t = [B, D]$ would be a valid intermediate canvas, in which case the model would be taught to predict $[A, C, E]$. The model is taught to insert multiple tokens at incomplete slots, or predict end-of-slot for completed slots. The intermediate canvases are uniformly sampled from the ground truth target sequence. During inference, the target canvas starts empty, and tokens will be inserted iteratively until the model predicts to insert empty tokens everywhere,

or the sequence has exceeded the specified maximum length.

Larger Context and Sentence-Positional Embeddings. Longer sequences lead to more uncertainties for the Insertion Transformer. For example, if a token in \hat{y}_t appears in multiple sentences in the final output, there is ambiguity to the model which sentence it belongs to (and therefore where to attend to on both the source and target canvases). While there is location information endowed in the Transformer model, we hypothesize that token level positional information is insufficient (especially since we have limited training data). We believe that endowing the model with sentence-level positional information (i.e., which sentence each token belongs to) may help significantly disambiguate in such situations and help the model build a more robust attention mechanism.

Based on this motivation and assuming that the datasets have not only parallel documents, but also sentence alignment between source and target documents (which is true for WMT’19 document-level translation), we use sentence-positional embeddings on both the source and target sequences as shown in Figure 1. The intention is to endow the model with this prior knowledge on sentence alignment between the source and target, and thus more easily attend to the appropriate sentences based on sentence locality. More specifically, on the source side, we do not use any sentence separator tokens; on the target side, we start each sentence with a sentence separator. During inference we initialize the output hypothesis with empty $\langle s \rangle$ sentence separator tokens, where the number of $\langle s \rangle$ equals to the number of source sentences, which is equal to the number of target sentences to be generated. These $\langle s \rangle$ tokens serve as sentence anchor points, and have sentence-positional information. Figure 1 visualizes the model.

In this work we increased the context window size to cover multiple sentences or a short document. Note that there is only a limit on the maximum number of tokens in the entire sequence; there is no limit on the length of a single sentence, or the total number of sentences in the sequence.

3 Experiments

We experiment with the WMT’19 English→German document-level translation task (Barrault et al., 2019). The training dataset consists of parallel document-level data (Eu-

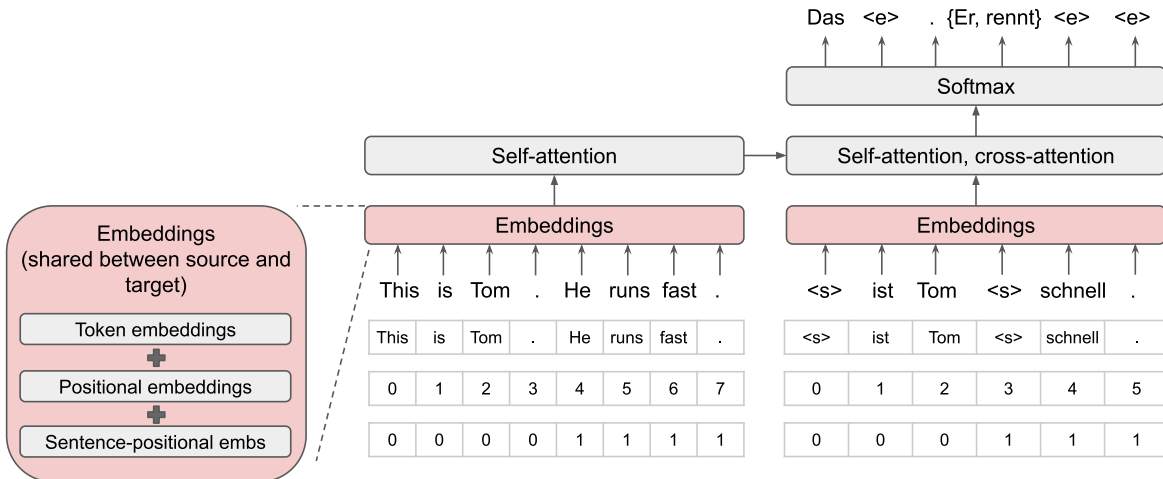


Figure 1: Big Bidirectional Insertion Representations for Documents

roparl, Rapid, News-Commentary) and parallel sentence-level data (WikiTitles, Common Crawl, Paracrawl). The test set is newstest2019. The document-level portion contains 68.4k parallel documents, or a total of 7.7M parallel sentences; while the sentence-level portion has 19.9M parallel sentences. We generated a vocabulary of 32k subwords from the training data using the SentencePiece tokenizer (Kudo and Richardson, 2018).

The Big BIRD model is as described in Section 2, and the baseline Insertion Transformer model has exactly the same configurations except without sentence-positional embeddings. To be explicit, our baseline Insertion Transformer model is also given the prior knowledge of number of source sentences in the document. The target canvas is initialized target with $\langle s \rangle$ sentence separator tokens, where the number of $\langle s \rangle$ tokens is equal to the number of sentences in the document. All our models follow the same architecture as the Transformer Base model in (Vaswani et al., 2017), and a context window of 1536 tokens during training (determined based on the longest document in the test set). All models were trained with the SM3 optimizer (Anil et al., 2019) with momentum 0.9, learning rate 0.1, and a quadratic learning rate warm-up schedule with 10k warm-up steps. The learning rate were chosen after some preliminary comparison runs between Adam and SM3. We opted to use the SM3 optimizer over Adam due to its more memory efficient properties, thus allowing us to use larger minibatches. Training was around 800k steps at batch size 512.

During training, each batch consists of 256 sub-

Model	BLEU
Insertion Transformer	25.3
Big BIRD	29.6

Table 1: WMT19 English→German Document-Level Translation.

documents and 256 sentences. Sub-documents are continuous sentences dynamically sampled from a document. The lengths of sub-documents are uniformly sampled in $(0, 1536]$ tokens. The number of sampled sub-documents from each document is 1/10 of the number of sentences in the full document. Sentences directly come from sentence-level data. This 1:1 mixing of sub-documents and sentences results in training examples of vastly different lengths and therefore many masked positions, and we plan to improve it in the future by packing multiple sentences into one example.

We report sacreBLEU (Post, 2018) scores of the two models in Table 1. Our Big BIRD model outperforms the Insertion Transformer model by +4.3 BLEU.

When we inspected the outputs more closely for the two models, we uncovered an interesting phenomenon. The Insertion Transformer, even though its target canvas is also initialized with the correct number of sentence $\langle s \rangle$ separators, struggles to align source and target sentences. For example, it can map two source sentences into one sentence in the target, or vice versa. This is not always bad, as long as it captures the semantics accurately. However, there are cases when misalignment causes loss of coherency. Table 2 shows such an example where Big BIRD captures alignment

Source:

(...) Chelsea faces Videoton in the UEFA Europa League at 3 p.m. on Thursday in London.

Target:

(...) Chelsea trifft in der UEFA Europa League am Donnerstag um 15 Uhr in London auf Videoton.

Insertion Transformer:

(...) Chelsea Gesichter am Donnerstag um 15.00 Uhr in London. Chelsea Gesichter Videoton in der UEFA Europa League.

Translation: (Google Translate)

Chelsea faces on Thursday at 15.00 in London. Chelsea faces Videoton in UEFA Europa League.

Big BIRD:

(...) Chelsea sieht am Donnerstag um 15.00 Uhr in London Videoton in der UEFA Europa League.

Translation: (Google Translate)

Chelsea sees Videoton in UEFA Europa League on Thursday at 15.00 in London.

Table 2: An example where the Insertion Transformer gets confused with sentence alignment: it maps one sentence from the source into two sentences in the translation and loses semantic accuracy. When given sentence alignment explicitly, i.e. Big BIRD, it translates the sentence coherently.

better than the Insertion Transformer, and therefore its translation is more accurate and coherent.

4 Conclusion

In this paper, we presented Big BIRD, an adaptation of the Insertion Transformer to document-level translation. In addition to a large context window, Big BIRD also uses sentence-positional embeddings to directly capture sentence alignment between source and target documents. We show both quantitatively and qualitatively the promise of Big BIRD, with a +4.3 BLEU improvement over the baseline model and examples where Big BIRD achieves better translation quality via sentence alignment. We believe Big BIRD is a promising direction for document level understanding and generation.

References

- Rohan Anil, Vineet Gupta, Tomer Koren, and Yoram Singer. 2019. Memory-Efficient Adaptive Optimization for Large-Scale Learning. In *arXiv*.
- Loc Barrault, Ondrej Bojar, Marta R. Costa-juss, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. Findings of the 2019 Conference on Machine Translation. In *ACL*.
- William Chan, Nikita Kitaev, Kelvin Guu, Mitchell Stern, and Jakob Uszkoreit. 2019. KERMIT: Generative Insertion-Based Modeling for Sequences. In *arXiv*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP*.
- Jiatao Gu, Qi Liu, and Kyunghyun Cho. 2019. Insertion-based Decoding with Automatically Inferred Generation Order. In *arXiv*.
- Hany Hassan, Anthony Aue and Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, Shujie Liu, Tie-Yan Liu, Renqian Luo, Arul Menezes, Tao Qin, Frank Seide, Xu Tan, Fei Tian, Lijun Wu, Shuangzhi Wu, Yingce Xia, Dongdong Zhang, Zhirui Zhang, and Ming Zhou. 2018. Achieving Human Parity on Automatic Chinese to English News Translation. In *arXiv*.
- Marcin Junczys-Dowmunt. 2019. Microsoft Translator at WMT 2019: Towards Large-Scale Document-Level Neural Machine Translation. In *WMT*.
- 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.
- Samuel Lubli, Rico Sennrich, and Martin Volk. 2018. Has Machine Translation Achieved Human Parity? A Case for Document-level Evaluation. In *EMNLP*.

- Sameen Maruf and Gholamreza Haffari. 2018. Document Context Neural Machine Translation with Memory Networks. In *ACL*.
- Matt Post. 2018. A Call for Clarity in Reporting BLEU Scores. In *WMT*.
- Mitchell Stern, William Chan, Jamie Kiros, and Jakob Uszkoreit. 2019. Insertion Transformer: Flexible Sequence Generation via Insertion Operations. In *ICML*.
- Ilya Sutskever, Oriol Vinyals, and Quoc Le. 2014. Sequence to Sequence Learning with Neural Networks. In *NIPS*.
- 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 *NIPS*.
- Sean Welleck, Kianté Brantley, Hal Daume, and Kyunghyun Cho. 2019. Non-Monotonic Sequential Text Generation. In *ICML*.