

# Summarize, Outline, and Elaborate: Long-Text Generation via Hierarchical Supervision from Extractive Summaries

Xiaofei Sun<sup>1,2</sup>, Zijun Sun<sup>2</sup>, Yuxian Meng<sup>2</sup>, Jiwei Li<sup>1,2</sup> and Chun Fan<sup>3</sup>

<sup>1</sup>Zhejiang University, <sup>2</sup>Shannon.AI, <sup>3</sup>Computer Center, Peking University

<sup>3</sup>National Biomedical Imaging Center, Peking University, <sup>3</sup>Peng Cheng Laboratory  
xiaofei\_sun@zju.edu.cn, jiwei\_li@shannonai.com

## Abstract

The difficulty of generating coherent long texts lies in the fact that existing models overwhelmingly focus on predicting local words, and cannot make high level plans on what to generate or capture the high-level discourse dependencies between chunks of texts. Inspired by human writing processes, where a list of bullet points or a catalog is first outlined, and then each bullet point is expanded to form the whole article, we propose *SOE*, a pipelined system that involves of summarizing, outlining and elaborating for long text generation: the model first outlines the summaries for different segments of long texts, and then elaborates on each bullet point to generate the corresponding segment. To avoid the labor-intensive process of summary soliciting, we propose the *reconstruction* strategy, which extracts segment summaries in an unsupervised manner by selecting its most informative part to reconstruct the segment.

The proposed generation system comes with the following merits: (1) the summary provides high-level guidance for text generation and avoids the local minimum of individual word predictions; (2) the high-level discourse dependencies are captured in the conditional dependencies between summaries and are preserved during the summary expansion process and (3) additionally, we are able to consider significantly more contexts by representing contexts as concise summaries. Extensive experiments demonstrate that *SOE* produces long texts with significantly better quality, along with faster convergence speed.

## 1 Introduction

Despite that recent large-scale pretrained language models (PLMs) (Devlin et al., 2018; Liu et al., 2019; Yang et al., 2019) are able to produce high-quality passages that can be hardly recognized by humans (Zellers et al., 2019), most of the generated “good” texts are within very limited length, e.g. hun-

dreds of tokens for most cases (Yan et al., 2020), thus generating coherent long texts remains a challenge (Radford et al., 2019; Tan et al., 2020). The difficulty lies in the fact that existing models generate texts in a word-by-word manner: predicting each subsequent token given its preceding contexts. This word-by-word strategy overwhelmingly focuses on predicting local words, and cannot make high level plans on what to generate. The strategy results in the fact that long texts generated by current models are usually repetitive, generic and self-contradictory (Shen et al., 2019).

To address these issues, the coarse-to-fine generation strategy is proposed (Fan et al., 2018; Xu et al., 2018; Yao et al., 2019; Mao et al., 2019). In coarse-to-fine generation, a list of keywords or a short prompt is first generated, serving as a summary of the original text. The prompt is then fed to a seq2seq model as an input to output the complete text. The coarse-to-fine generation strategy significantly improves generation over the word-by-word strategy, but still suffers from the following shortcomings: (a) *limited capacity of the prompt*: a single keyword list or prompt does not have enough capacity to summarize all the text of long passages, since long texts are usually consists of several parts, each of which focuses on a specific aspect or topic (Zhou et al., 2018; Guan et al., 2019). The usage of the coarse-to-fine generation strategy is thus limited to texts that can be summarized by a single prompt (e.g., short stories). This explains why text length generated by the progressive generation model is still limited, e.g., the introduced writing prompts dataset in Fan et al. (2018) has an average length of stories around 735, and the average length of prompts is 28; (b) *ignorance of high-level discourse dependency*: the coarse-to-fine generation strategy does not capture discourse-level dependencies (Li and Jurafsky, 2016), which handle the high-level information flow and interactions between

segments of texts. The ignorance of discourse-level dependencies results in texts lacking for coherence.

Humans write in a hierarchical top-down manner: before writing a thousand-word-long essay, a human usually first prepares a list of bullet points or a catalogue, and then expands them to form the whole article. The sentence-level coherence between these bullet points is preserved when the bullet points are expanded, providing guarantees that the full text is coherent.

To mimic this top-down manner in human writing, we propose *SOE*, a pipelined system that involves of summarizing, outlining and expanding for long text generation: the model first outlines the summaries for different segments of long texts, which actually mimics the process of humans outlining bullet points; next, the model elaborates on each bullet point to generate the corresponding segment. The proposed strategy comes with the following merits: (a) Since each segment is associated with its own summary rather than the entire text sharing a single prompt, the capacity of summaries to reconstruct the full text can be guaranteed; (b) The conditional generation probability between summaries captures the high-level discourse dependencies, and these dependencies are preserved when they are expanded to segments. This naturally resolves the incapability of modeling discourse-level dependencies in the coarse-to-fine generation approach. (c) This model is able to consider significantly larger amount of contexts by representing chunks of contexts as concise summaries.

Empirically, we do not readily have summaries for segments in hand. The model thus needs to learn to summarize in an unsupervised manner. To this end, we propose the *reconstruction* strategy, which extracts segment summaries by selecting its most informative part to reconstruct the segment. Extensive experiments demonstrate that *SOE* produces long texts with significantly better quality than existing baselines.

## 2 Related Work

### 2.1 Generating Long Texts

There are two trends for generating long text: This first trend of work tackles the problem from the model perspective. New model structures (Kitaev et al., 2020; Child et al., 2019; Dai et al., 2019; Ye et al., 2019; Guo et al., 2019; Sukhbaatar et al.,

2019; Correia et al., 2019; Beltagy et al., 2020; Zaheer et al., 2020) are designed to give the model the ability to ingest more contexts given limited memories or computing power. For example, Transformer-XL (Dai et al., 2019), a modifier to Transformers (Vaswani et al., 2017) uses a segment-level recurrence mechanism to enable learning long-term dependencies; Child et al. (2019); Correia et al. (2019); Kitaev et al. (2020); Beltagy et al. (2020); Zaheer et al. (2020) proposed to sparsify transformers by focusing only on a fraction of attention connections; Tay et al. (2020) replaced the dot-product self-attention with learned synthetic attention weights.

The second trend of researches focus on developing new generation strategies. Efforts have been devoted to the idea of planning-then-generation or coarse-to-fine generation (Sha et al., 2017; Gehrmann et al., 2018a; Wiseman et al., 2019; Hua and Wang, 2019; Shen et al., 2020; Fu et al., 2020), which greatly inspires this work. In coarse-to-fine generation, a list of keywords or a short sentence is first generated, providing guidance to generate the full text. A recent work from Tan et al. (2020) takes a multi-step strategy, which progressively refines the generated incomplete text until reaching a specified stage. Similar ideas have also been applied to text summarization, where Gehrmann et al. (2018b) proposed a bottom-up method that first identifies phrases within a document that are likely included in its summary. Our work is also inspired by the strategy of hierarchical generation (Li et al., 2015b), which consider text units with bigger granularity: Li et al. (2015b) proposed hierarchical LSTMs that arrange tokens, sentences and paragraphs in a hierarchical structure, with different levels of LSTMs capturing compositionality, and Shen et al. (2019) used multi-level structures to learn a VAE model for generating long coherent text.

### 2.2 Extractive Summarization

Extractive summarization refers to the problem of selecting part of the input text as its summary. A fundamental problem in extractive summarization is to score constituent texts units (e.g., phrases, sentences or paragraphs) and select highly-ranked one(s) as the summary. Haghighi and Vanderwende (2009) used word frequencies in the input text to assign scores to words, which are then in turn used to score sentences. Higher-ranked sentences are

selected as the summary of the input text. Liu et al. (2018) presented a two-stage extractive-abstractive framework, which first coarsely identifies salient information, followed by a generation model used to refine it. Neural models have been widely used for scoring Cao et al. (2015); Ren et al. (2017). Liu and Lapata (2019) finetuned BERT (Devlin et al., 2018) to score each sentence for extractive summarization; Zhang et al. (2019) computed token similarity in each sentence using BERT contextual embeddings to serve as an automatic evaluation metric for text generation.

### 3 Background

**Language Modeling** refers to the process of calculating the probability  $p(\mathbf{y})$  of a sequence  $\mathbf{y} = [y_1, \dots, y_T]$ , where each  $y_i$  denotes a constituent token of  $p(\mathbf{y})$ . The probability can be computed by decomposing the joint distribution  $p(\mathbf{y})$  into a product of conditional distributions over tokens:

$$p(\mathbf{y}) = \prod_{t=1}^T p(y_t | \mathbf{y}_{<t}) \quad (1)$$

where  $\mathbf{y}_{<i} = [y_1, \dots, y_{i-1}]$  is the partial sequence of tokens generated previously. During training, the model is optimized to minimize the negative log-likelihood (NLL)  $-\sum_{\mathbf{y} \in \mathcal{D}} \log p(\mathbf{y})$ . During inference, the model decodes a token at each time step  $t$  according to  $p(y_t | \mathbf{y}_{<t})$  based on the softmax functions  $y_t \propto \text{softmax}(\mathbf{W}^{\text{out}} \mathbf{h}_t)$  where  $\mathbf{W}^{\text{out}} \in \mathbb{R}^{d \times |V|}$  is the output word embedding matrix and  $\mathbf{h}_t$  is the hidden state at time-step  $t$ .

**Sequence-to-Sequence (Seq2Seq) Generation** models generate a target sequence  $\mathbf{y}$  conditioning on a given source sequence  $\mathbf{x}$ , which differs from language models (LMs) in terms of whether or not conditioning on another input sequence. Similar to LMs, the probability of the target sequence can be typically factorized as:

$$p(\mathbf{y} | \mathbf{x}) = \prod_{i=1}^T p(y_i | \mathbf{y}_{<i}, \mathbf{x}) \quad (2)$$

Seq2seq models are also optimized to minimize the NLL  $-\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \log p(\mathbf{y} | \mathbf{x})$ . In the rest of this paper, we unify the notation of  $p(\mathbf{y})$  and  $p(\mathbf{y} | \mathbf{x})$  by setting  $\mathbf{x} = \emptyset$  for LMs. Different architectures have been proposed to model  $p(y_t | \mathbf{y}_{<t}, \mathbf{x})$ , including transformers (Vaswani et al., 2017), LSTMs (Luong et al., 2015) and CNNs (Dauphin et al.,




Figure 1: An overview of the proposed method. Given preceding tokens  $\mathbf{y}^{<i}$ , we first sequentially generate summaries  $s^{j \sim k}$  for each snippet. Next we expand each summary  $s$  to form the full text  $\mathbf{y}^{j \sim k}$ .

2017). At test time, sequences are usually generated using beam search, or its variants to promote diversity (Vijayakumar et al., 2016).

## 4 Model Details for SOE

### 4.1 Notations

A long sequence of tokens  $\mathbf{Y} = \{\mathbf{y}^1, \mathbf{y}^2, \dots, \mathbf{y}^K\}$  is first sliced into a series of snippets  $\mathbf{y}^i$ 's, where  $K$  denotes the number of constituent snippets. Here we use the bold font  $\mathbf{y}$  to denote snippets, and the normal font  $y$  to denote tokens. The number of tokens  $N$  within each snippet is a hyperparameter. We also use superscript  $i$  to denote the index of a snippet, and subscript  $l$  to denote the index of a token. Each  $\mathbf{y}^i$  consists a sequence of tokens  $\mathbf{y}^i = \{y_1^i, \dots, y_{n_i}^i\}$ , where  $n_i$  denotes the length of  $\mathbf{y}^i$ . Our goal is to generate a subset of  $\mathbf{Y}$ , denoted by  $\mathbf{y}^{j \sim k} = \{\mathbf{y}^j, \mathbf{y}^{j+1}, \dots, \mathbf{y}^k\}$  given its preceding snippets, denoted by  $p(\mathbf{y}^{j \sim k} | \mathbf{y}^{<j})$ . Each snippet  $\mathbf{y}^i$  is associated with a short summary  $s^i = \{s_1^i, s_2^i, \dots, s_{m_i}^i\}$ , where  $s_l^i$  denotes tokens and  $m_i$  is the number of tokens in  $s^i$ .

### 4.2 Pipeline Overview

Instead of generating all constituent words in  $\mathbf{Y}$  one by one, we adopt a hierarchical strategy. The process of generating  $\mathbf{y}^{j \sim k}$  is decoupled into the following two stages.

**Stage1: Outlining Segment Summaries** : we sequentially generate the summary  $s^i$  for each snippet conditioning on the summaries for previous snippets. This mimics the process of catalog generation when humans write.

**Stage2: Expanding Summaries to Texts**: we expand each summary  $s^i$  to the full segment by sequentially generating its constituent words.

An overview of the proposed method is shown in

Figure 1.

### 4.3 Extracting Golden Summaries

At the training time, we need to learn to generate summaries. But this is not straightforward because the golden summary  $s^i$  for the snippet  $y^i$  is not readily at hand. Manually soliciting summaries like Fan et al. (2018) is both costly and slow. We thus propose to take the idea of unsupervised extractive summarization, and for each snippet  $y^i$ , we extract its summary  $s^i$  unsupervisedly, and use the extracted  $s^i$  as the golden summary for learning.

We investigate Random, TF-IDF (Ramos, 2003), TextRank (Mihalcea and Tarau, 2004) and Reconstruction methods to access the importance of selecting summary sentences, the first three of which are similar to Liu et al. (2018). More details of these three methods can be found at A.

**Reconstruction** A summary should be more informative than non-summary sentences, that is, a summary should have the most ability to reconstruct the full text. To measure the degree of a sentence’s *reconstruction* ability, we use a seq2seq model to predict the original given text the summary sentence, the probability of which is regarded as the reconstruction score. Suppose there are  $n$  sentences in  $y^i$  (e.g.,  $n = 4$ ), and  $y^i = \{y^{i,1}, y^{i,2}, y^{i,3}, y^{i,4}\}$ , and  $y^{i,j}$  denotes the  $j$ -th sentence in  $y^i$ . The reconstruction score for  $y^{i,j}$ , denoted by  $\text{Score}(y^{i,j})$  is given as follows:

$$\text{Score}(y^{i,j}) = \frac{1}{|y^i|} \log p(y^i | y^{i,j}) \quad (3)$$

To obtain  $p(y^i | y^{i,j})$ , we train another seq2seq model, where the input is  $y^{i,j}$  for each  $j$ , and the output is  $y^i$  by sequentially predicting tokens in  $y^i$ . Given the trained model, we rank all sentences in  $y^i$  and use the one with the highest score as the golden summary  $s^i$ .

### 4.4 Outlining Segment Summaries

In the summary generation stage, we cannot observe  $y^{j \sim k}$ , and our goal is to sequentially generate  $s^{j \sim k}$  given  $y^{<j}$ :

$$p(s^{j \sim k} | y^{<j}) = \prod_{i \in [j,k]} p(s^i | y^{<i}, s^{<i}) \quad (4)$$

The generation of summary  $s^i$  can be factorized into sequentially generating the constituent word

within it:

$$p(s^i | y^{<i}, s^{<i}) = \prod_{l \in [1, m_i]} p(s_l^i | s_{<l}^i, y^{<i}, s^{<i}) \quad (5)$$

This process ends until generating a special end-of-sequence token  $\langle \text{EOS} \rangle$  or reaching a specified summary length  $m$ . We use the Transformer-base (Vaswani et al., 2017) architecture as the backbone. For considering more contexts, we adopt the segment-level recurrence strategy, similar to Dai et al. (2019), where the hidden states computed for far away snippets are fixed and cached to be reused for the next new snippet. Gradients are not propagated to these far away snippets for memory and computation efficiency. This strategy allows the model to exploit information in history to the largest extent.

### 4.5 Expanding Summaries to Texts

Next, we expand each summary  $s^i$  to the full text for each segment by sequentially generating its constituent words

$$p(y^i | y^{<i}, s^i) = \prod_{l \in [1, n_i]} p(y_l^i | y_{<l}^i, s^i, y^{<i}) \quad (6)$$

which has the same termination conditions as in the summarization generation.

### 4.6 Training and Inference

**Training** For summary generation, the transformer model takes  $[y^{<i}; s^{<i}]$  as the input and is optimized by minimizing the NLL loss  $-\log p(s^i | y^{<i}; s^{<i})$ . Due to the memory limitation, we limit  $y^{<i}$  to preceding 384 tokens, and  $s^{<i}$  to 128 tokens at training. It is worth noting that the 384 tokens of  $y^{<i}$  mostly come from the segment right before, i.e.,  $y^{i-1}$ , while  $s^{<i}$  comes from multiple preceding segments since the summary is more concise.

For the summary expanding stage, the transformer model takes  $[y^{<i}; s^i]$  as input and is optimized by minimizing the NLL loss  $-\log p(\hat{y}^i | y^{<i}; s^i)$ . The two models, i.e., the summary generation and the summary expansion model share parameters, with a task-specific token appended to the start to notify the model on what to generate, summaries or segments.

**Inference** At test time, we first use beam search with beam size 5 to generate summaries. Given the generated summary, beam search is used again to

generate the corresponding segment. We consider more contexts at test time, where  $\mathbf{y}^{<i}$  is limited to 1,156 tokens and  $\mathbf{s}^{<i}$  is limited to 512 tokens.

Additionally, we augment the vanilla beam search with the strategy of mutual information reranking (Li et al., 2015a; Fang et al., 2015). The key point of mutual information is to, instead of merely handling the uni-directional dependency from the source to target based on the forward probability  $\log p(\text{target}|\text{source})$ , it models the mutual dependency between the source and target in sequence-to-sequence generation, i.e., the combination of the forward probability  $\log p(\text{target}|\text{source})$  and the backward probability  $\log p(\text{source}|\text{target})$ . Specifically in our case, during summary generation,  $\mathbf{s}^i$  is generated as follows:

$$\mathbf{s}^i = \arg \max_{\mathbf{s}^i} [\log p(\mathbf{s}^i | \mathbf{y}^{<i}, \mathbf{s}^{<i}) + \log p(\mathbf{s}^{i-1} | \mathbf{y}^{<i-1}, \mathbf{s}^i)] \quad (7)$$

where  $p(\mathbf{s}^{i-1} | \mathbf{y}^{<i}, \mathbf{s}^i)$  is the backward probability of predicting the preceding summary  $\mathbf{s}^{i-1}$  given  $\mathbf{s}^i$ . Since direct decoding from Eq.7 is infeasible, we follow the practical solution in Li et al. (2015a), where we first generate an  $N$ -best list based on the forward probability  $p(\mathbf{s}^i | \mathbf{y}^{<i}, \mathbf{s}^{<i})$ ,<sup>1</sup> and then rerank the  $N$ -best list by combining the forward probability and the backward probability.

Similar strategy can also be applied to the summary expanding stage, where  $\mathbf{y}^i$  is obtained as follows:

$$\mathbf{y}^i = \arg \max_{\mathbf{y}^i} [\log p(\mathbf{y}^i | \mathbf{y}^{<i}, \mathbf{s}^i) + \log p(\mathbf{y}^{i-1} | \mathbf{y}^i)] \quad (8)$$

The backward probability  $p(\mathbf{y}^{i-1} | \mathbf{y}^i)$  predicts the preceding segment given the current segment. Again, beam search is combined with reranking to approximately find the optimal result.

#### 4.7 Slicing Texts based on Coherence Scores

One more thing we need to care about is how to slice the text into segments. The simplest way is to slice the full text equally. But this is sub-optimal since the break point could be in the middle of two closely related sentences and one segment might contain multiple aspects.

We thus propose a slicing strategy based on sentence-level coherent scores. Using the Next

<sup>1</sup>We simplify  $p(\mathbf{s}^{i-1} | \mathbf{y}^{<i-1}, \mathbf{s}^i)$  as  $p(\mathbf{s}^{i-1} | \mathbf{s}^i)$ , where we train a seq2seq model to predict the preceding summary given the current summary.

Model	WikiText-103		BookCorpus	
	PPL↓	# Para	PPL↓	# Para
<i>Base</i>				
<i>Vanilla</i>	25.0	130M	29.0	130M
<i>WritingPrompts-Keyword</i>	23.8	135M	28.3	135M
<i>WritingPrompts-Sentence</i>	24.1	135M	28.6	135M
<i>Progressive WritingPrompts</i>	23.3	150M	27.7	150M
<i>SOE</i>	<b>22.2</b>	132M	<b>25.7</b>	132M
<i>Large</i>				
<i>Vanilla</i>	20.0	220M	24.8	220M
<i>SOE</i>	<b>17.4</b>	224M	<b>22.5</b>	224M

Table 1: Perplexity of different models on WikiText-103 and BookCorpus. *Vanilla* stands for our implementation of Transformer-XL (Dai et al., 2019).

Sentence Prediction (NSP) from BERT (Devlin et al., 2018), we are able to measure the coherence score between two consecutive sentences with index  $i$  and  $i + 1$ , denoted by  $\text{Score}(i, i + 1)$ . Given a full text  $\mathbf{y} = \{y(1), y(2), \dots, y(T)\}$ , let  $T$  denote the number sentences in  $\mathbf{y}$ , and  $y(i)$  denote the  $i$ th sentence. Given a fixed value  $K$  for the number sliced segments,  $\mathbf{y}$  will be sliced into  $K$  segments, i.e.,  $\mathbf{y}^1, \mathbf{y}^2, \dots, \mathbf{y}^K$ , where each  $\mathbf{y}^k$  consists of a group of consecutive sentences from  $\mathbf{y}$ . Let  $G_k$  denote the list of indexes of sentence in original  $\mathbf{y}$ , where  $G_k[1]$  denotes the index of the first sentence in  $G_k$ ,  $G_k[2]$  denotes the second sentence, etc. Let  $R_k = |G_k|$  denote the number of sentences in  $G_k$ .

We wish to maximize the coherence scores between two consecutive sentences within the same segment and minimize the score between two consecutive sentences belonging to different segments, giving the following objective to optimize:

$$L = \sum_{k=1}^K \sum_{i \in [1, R_k - 1]} \text{Score}(G(k)[i], G(k)[i + 1]) - \sum_{k=1}^{K-1} \text{Score}(G[k][R_k], G[k][1]) \quad (9)$$

where  $\text{Score}(G[k][R_k], G[k][1])$  the coherence score between the ending sentence of a segment and the starting sentence of the next segment. Given  $\text{Score}(i, j)$ , Eq.9 can be readily solved using linear programming.

## 5 Experiments

In this section, we present experiment results. For different methods to generate summaries, we find that the performance of *Reconstruction* consistently outperforms the rest in our preliminary results. We

Model	MSJ $\uparrow$			Diversity $\uparrow$		Adversarial Success $\uparrow$	S-Level Coherence $\uparrow$
	MSJ-2	MSJ-3	MSJ-4	D-1	D-2	Adversarial Success	NSP
	<i>Base</i>						
<i>Vanilla</i>	62.6	41.5	16.9	7.4	19.8	0.037	0.812
<i>WritingPrompts-Keyword</i>	63.0	42.2	17.5	8.9	22.0	0.057	0.836
<i>WritingPrompts-Sentence</i>	63.1	42.2	17.7	8.5	21.0	0.046	0.834
<i>Progressive WritingPrompts</i>	63.9	42.5	18.0	10.7	25.9	0.055	0.854
<i>SOE</i>	64.8	43.9	19.4	16.4	34.3	0.072	0.870
<i>SOE+MI</i>	<b>65.2</b>	<b>44.4</b>	<b>20.0</b>	<b>20.6</b>	<b>40.8</b>	<b>0.103</b>	<b>0.881</b>

Table 2: Results of different models in terms of diversity, adversarial success, MSJ and sentence-level coherence on the BookCorpus corpora. *Vanilla* stands for our implementation of Transformer-XL (Dai et al., 2019). “D- $n$ ” stands for “Distinct- $n$  ( $n = 1, 2$ )”, and *MI* stands the results for mutual information reranking.

Model	Distinct- $n$ $\uparrow$	
	Distinct-1 $\uparrow$	Distinct-2 $\uparrow$
	<i>Large</i>	
<i>Vanilla</i>	11.7	25.5
<i>SOE</i>	24.1	45.0
<i>SOE+MI</i>	29.3	48.8

Table 3: Results of different models with large volumes in terms of diversity on the BookCorpus dataset.

thus only report results from *Reconstruction* in the section. We will get back to analysis on different summary generation methods in the ablation study section.

## 5.1 Datasets and Evaluation Metrics

We need a corpus of contiguous and long text to test SOE. We use two word-level datasets, WikiText-103 (Merity et al., 2016) and the BookCorpus dataset (Zhu et al., 2015). WikiText-103 contains 103M training words from 28K articles, with an average length of 3.6K words per article. WikiText-103 can be used to test the ability of modeling long-term dependencies. The BookCorpus dataset is a more suitable dataset for our purpose, with much longer and more contiguous texts. It contains a total number of roughly 1 billion words and 74 million sentences from 11k books, with an average length of 89K words for each book. The average number of words per sentence is 13. For both datasets, we predict the last 2,000 tokens at test time.

We use Perplexity (PPL), Diversity (Distinct- $n$ ) (Li et al., 2016), Adversarial Success (Kannan and Vinyals, 2017; Li et al., 2017), MS-Jaccard (MSJ) (Montahaei et al., 2019) and Sentence-Level Coherence (Tan et al., 2020) as evaluation metrics.

- **Perplexity (PPL)** Perplexity measures how fluent a piece of generated text could be (Dai et al., 2019). We use PPL as the basic evalua-

tion metric in our experiments.

- **Diversity** Perplexity cannot measure how diverse the generated text is. We thus use the scaled number of distinct unigrams (Distinct-1) and bigrams (Distinct-2) to demonstrate the degree of diversity (Li et al., 2016) for generated texts.
- **Adversarial Success** Inspired by adversarial evaluations (Bowman et al., 2016; Kannan and Vinyals, 2017; Li et al., 2017), we use the adversarial success metric, which is defined as the fraction of a model successfully fooling a trained evaluator to believe that machine-generated texts are from humans. The evaluator is a binary classification model. At the training time, it takes as inputs machine-generated texts and original texts, and are trained to discriminate them. At test time, adversarial success is the value  $1 - acc$ , where  $acc$  denotes the accuracy of the trained evaluator predicting machine-generated texts as machine-generated. Higher values of adversarial success denotes better text quality.
- **MS-Jaccard (MSJ)** MSJ measures the similarity of the  $n$ -gram frequencies between the generated texts and the golden texts (Montahaei et al., 2019). We report MSJ-2, -3 and -4.
- **Sentence-Level Coherence** PPL, MSJ and diversity scores do not reflect the sentence-level coherence of generated texts. We adopt the strategy in Tan et al. (2020) where Next Sentence Prediction (NSP) from pretrained BERT model (Devlin et al., 2018) is used as a metric to measure the coherence between each sentence and its next sentence. We report average NSP scores for all consecutive sentence pairs within the generated text.

## 5.2 Baselines

In this paper, we use Transformer-XL (Dai et al., 2019), WritingPrompts (Fan et al., 2018), and Progressive WritingPrompts (Tan et al., 2020) as baselines. More details of the baseline models can be found at Appendix B.

For all models, we use Adam (Kingma and Ba, 2014) with learning rate of  $1e-4$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , rate warmup over the first 10,000 steps, and linear decay of the learning rate. We use a dropout rate of 0.1 on all layers.

## 5.3 Results

Table 1 shows the results of perplexity for different models on the WikiText-103 and BookCorpus datasets. On both datasets, SOE achieves the lowest PPL compared to baselines Transformer-XL (Dai et al., 2019), WritingPrompts (Fan et al., 2018) and Progressive (Tan et al., 2020). In particular, for WikiText-103, we gain a PPL decrease -2.8, -1.6 and -1.1 against our implemented Transformer-XL, WritingPrompts-Sentence and Progressive, while having the same or even fewer parameters. Similar trend can be observed on BookCorpus.

Table 2 and Table 3 show the results for MSJ, diversity, adversarial success and sentence-level coherence scores. As can be seen, WritingPrompt based models generally outperform the Transformer-XL model, which adopts the word-by-word generation strategy. This validates the superiority of two-step generation strategy over the naive word-by-word generation strategy for long-text generation. The progressive WritingPrompt model, which involves multi-step of generation and expanding, outperforms the one-step the WritingPrompt-keyword and WritingPrompt-sentence model, which is in accord with our expectation. SOE achieves significantly better results compared to Vanilla, WritingPrompts and Progressive models in terms of all evaluation metrics, showing that the proposed method can produce more fluent, coherent and diverse texts. The consistent performance boosts on all metrics demonstrate the importance of modeling discourse-level dependencies and necessity of summary expanding strategy for long-text generation.

Additionally, we observe additional performance boosts by mutual information (MI), especially for diversity and adversarial success. This is in accord with our expectation: since mutual information is




Figure 2: PPL on the BookCorpus dataset w.r.t. different segment lengths.

able to build bidirectional dependencies between the source and the target, models enhanced with mutual information can generate better summaries, and the phenomenon of generic and repetitive generation can be alleviated (Li et al., 2016), leading to more diverse results.

## 6 Ablation Studies

### 6.1 The Effect of Segment Length

The size of the segment can be neither too big nor too small: extremely long segments, might contain too many aspects or topics for the summary to summarize, in which case the model will degenerate into the WritingPrompts model (Fan et al., 2018). For too short segments, the summary cannot provide high-level guidance. We thus need to find the sweet spot for the segment length. Figure 2 shows results on the BookCorpus dataset. It is clear from the figure that too short segments and too long segments both lead to inferior performances.

### 6.2 The Effect of Summary Generation Strategies

It is worthwhile to explore how different summary extraction methods affect the final performances. To this end, we conduct experiments on the BookCorpus dataset, using different summary extraction methods, i.e., Random, TextRank, TF-IDF and Reconstruction. Table 4 shows the results. We first compare the ppl for summary generation, where the reconstruction model achieves the lowest ppl and thus produces summaries that are the easiest to predict given preceding contexts. It is also interesting to see that across all summary generation strategies, ppl for summarization generation is significantly larger than text prediction, which is reasonable since (1) generating summaries for the

Method	Summary PPL↓	Text PPL↓	MJ-4 ↑
<i>Vanilla</i>	-	29.0	16.9
<i>Random</i>	40.1	30.2	15.5
<i>TextRank</i>	30.7	26.2	17.8
<i>TF-IDF</i>	33.0	26.9	17.3
<i>Reconstruction</i>	<b>30.4</b>	<b>25.7</b>	<b>19.4</b>

Table 4: Performances of different summary extraction methods described in Section 4.3. *Vanilla* is the plain model that generates tokens one by one without summaries.

upcoming segment requires more generalization abilities; and (2) there are more diverse options for what the next segment should talk about than the local choices for what the next sentence should talk about. For the final text-generation ppl, reconstruction achieves the best results, in terms of PPL, MJ-3 and MJ-4. TextRank and TF-IDF are better than Vanilla. Interestingly, the strategy of using random sentences as summaries performs worse than without summaries, which can be explained by providing no guidance is better than incorrect guidance.

### 6.3 The Effect of Coherence-based Text Slicing

We replace the coherence-based text slicing strategy with the naive equal slicing strategy, and see how this will negatively affect the performance. On the BookCorpus dataset, we observe an increase of summary generation ppl from 30.4 to 30.9, and an +0.7 increase of PPL from 25.7 to 26.4 in token generation, which demonstrates the importance of slicing text into coherent segments for generation. But it is also worth-noting that, even with the native equal slicing strategy, SOE still performs significantly better than other baseline models.

### 6.4 Decoupling The Effects of Summaries

The positive effects from summaries are two-fold: (1) it provides high-level guidance for segment generation; and (2) with far-away segments being concisely represented by summaries, it gives the model the ability to consider longer contexts. To quantitatively measure the influences from both aspects, we conduct the following experiments: at test time, for the computation  $p(s^i|y^{<i}; s^{<i})$  and  $p(y^i|y^{<i}; s^i)$ , the model can only access summaries for segments that are used as contexts. In other words, only summaries within the 1,156 tokens of preceding contexts can be fed as inputs. This is different from the original version of SOE, in which s can extend




Figure 3: Convergence speed for different models.

to preceding contexts until the limit of 512 tokens is reached. We did not retrain the model, but add this limitation at test time. On the BookCorpus dataset, this leads to an increase of 0.8 in PPL (25.7 vs 26.5), and a decrease of 0.5 and 0.8 in MJ-3 (43.5 vs 43.9) and MJ-4 (18.6 vs 19.4).

### 6.5 Simplifying $p(s^i|y^{<i}; s^{<i})$

Here we explore different simplifications for  $p(s^i|y^{<i}; s^{<i})$ . For  $p(s^i|y^{<i}; s^{<i})$ , the current summary is generated based on both previous summaries and segment tokens. We can simplify it as  $p(s^i|s^{<i})$ , where previous segment tokens are not fed as inputs to predict the summary, which will significantly decrease computing complexity. On the BookCorpus dataset, we observe an increase of PPL in summary generation from 30.4 to 31.2, which subsequently leads to an +0.9 increase of PPL from 25.7 to 26.6 in token generation.

### 6.6 Convergence Speed

At last, we investigate how quickly different models converge. Results are shown in Figure 3. With the guidance of extracted summaries, SOE has a conspicuously faster convergence speed, where at about 200K training steps it has approximately reached the best result while the other two models — Vanilla and WritingPrompts — do not converge until 1000K training steps. The WritingPrompts model converges faster than then Vanilla because of the high-level guidance from prompts.

## 7 Conclusion

In this paper, we propose a two-step hierarchical generation strategy for long-text generation: the model first generates the summary for each segment conditioning on previous summaries, and next, each summary is expanded to form the full



text segment. The proposed strategy provides high-level guidance for local text generation, and enables high-level discourse dependencies to be captured. Extensive experiments demonstrate that SOE produces long texts with significantly better quality,

## Acknowledgement

We would like to thank anonymous reviewers for their comments and suggestions. This work is supported by the Key R & D Projects of the Ministry of Science and Technology (2020YFC0832500), Program of Zhejiang Province Science and Technology (2022C01044).

## References

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer.
- Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space.
- Ziqiang Cao, Furu Wei, Li Dong, Sujian Li, and Ming Zhou. 2015. Ranking with recursive neural networks and its application to multi-document summarization. In *AAAI*, pages 2153–2159. Citeseer.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*.
- Gonçalo M. Correia, Vlad Niculae, and André F. T. Martins. 2019. Adaptively sparse transformers. 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 2174–2184, Hong Kong, China. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. Language modeling with gated convolutional networks. In *International conference on machine learning*, pages 933–941.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*.
- Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. 2015. From captions to visual concepts and back. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1473–1482.
- Yao Fu, Yansong Feng, and John P. Cunningham. 2020. Paraphrase generation with latent bag of words.
- Sebastian Gehrmann, Falcon Dai, Henry Elder, and Alexander Rush. 2018a. End-to-end content and plan selection for data-to-text generation. In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 46–56, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018b. Bottom-up abstractive summarization. *arXiv preprint arXiv:1808.10792*.
- Jian Guan, Yansen Wang, and Minlie Huang. 2019. Story ending generation with incremental encoding and commonsense knowledge. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6473–6480.
- Qipeng Guo, Xipeng Qiu, Pengfei Liu, Yunfan Shao, Xiangyang Xue, and Zheng Zhang. 2019. Star-transformer. 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 1315–1325, Minneapolis, Minnesota. Association for Computational Linguistics.
- Aria Haghighi and Lucy Vanderwende. 2009. Exploring content models for multi-document summarization. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 362–370.
- Xinyu Hua and Lu Wang. 2019. Sentence-level content planning and style specification for neural text generation.
- Anjali Kannan and Oriol Vinyals. 2017. Adversarial evaluation of dialogue models. *arXiv preprint arXiv:1701.08198*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015a. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A diversity-promoting objective

- function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Jiwei Li and Dan Jurafsky. 2016. Neural net models for open-domain discourse coherence. *arXiv preprint arXiv:1606.01545*.
- Jiwei Li, Minh-Thang Luong, and Dan Jurafsky. 2015b. A hierarchical neural autoencoder for paragraphs and documents. *arXiv preprint arXiv:1506.01057*.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial learning for neural dialogue generation. *arXiv preprint arXiv:1701.06547*.
- Peter J Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198*.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Huanru Henry Mao, Bodhisattwa Prasad Majumder, Julian McAuley, and Garrison W Cottrell. 2019. Improving neural story generation by targeted common sense grounding. *arXiv preprint arXiv:1908.09451*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 404–411.
- Ehsan Montahaee, Danial Alihosseini, and Mahdiah Soleymani Baghshah. 2019. Jointly measuring diversity and quality in text generation models. *arXiv preprint arXiv:1904.03971*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8).
- J. Ramos. 2003. Using tf-idf to determine word relevance in document queries.
- Pengjie Ren, Zhumin Chen, Zhaochun Ren, Furu Wei, Jun Ma, and Maarten de Rijke. 2017. Leveraging contextual sentence relations for extractive summarization using a neural attention model. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 95–104.
- Lei Sha, Lili Mou, Tianyu Liu, Pascal Poupart, Sujian Li, Baobao Chang, and Zhifang Sui. 2017. Order-planning neural text generation from structured data. *arXiv preprint arXiv:1709.00155*.
- Dinghan Shen, Asli Celikyilmaz, Yizhe Zhang, Liqun Chen, Xin Wang, Jianfeng Gao, and Lawrence Carin. 2019. Towards generating long and coherent text with multi-level latent variable models. *arXiv preprint arXiv:1902.00154*.
- Tianxiao Shen, Victor Quach, Regina Barzilay, and Tommi Jaakkola. 2020. Blank language models. *arXiv preprint arXiv:2002.03079*.
- Sainbayar Sukhbaatar, Edouard Grave, Piotr Bojanowski, and Armand Joulin. 2019. Adaptive attention span in transformers. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 331–335, Florence, Italy. Association for Computational Linguistics.
- Bowen Tan, Zichao Yang, Maruan Al-Shedivat, Eric P. Xing, and Zhiting Hu. 2020. Progressive generation of long text.
- Yi Tay, Dara Bahri, Donald Metzler, Da-Cheng Juan, Zhe Zhao, and Che Zheng. 2020. Synthesizer: Rethinking self-attention in transformer models. *arXiv preprint arXiv:2005.00743*.
- 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*.
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2019. Learning neural templates for text generation.
- Jingjing Xu, Xuancheng Ren, Yi Zhang, Qi Zeng, Xiaoyan Cai, and Xu Sun. 2018. A skeleton-based model for promoting coherence among sentences in narrative story generation. *arXiv preprint arXiv:1808.06945*.
- Yu Yan, Weizhen Qi, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. Prophetnet: Predicting future n-gram for sequence-to-sequence pre-training.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for lan-

guage understanding. In *Advances in neural information processing systems*, pages 5753–5763.

Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. Plan-and-write: Towards better automatic storytelling.

Zihao Ye, Qipeng Guo, Quan Gan, Xipeng Qiu, and Zheng Zhang. 2019. Bp-transformer: Modelling long-range context via binary partitioning. *arXiv preprint arXiv:1911.04070*.

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*.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. In *Advances in Neural Information Processing Systems*, pages 9054–9065.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Deyu Zhou, Linsen Guo, and Yulan He. 2018. Neural storyline extraction model for storyline generation from news articles. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1727–1736, New Orleans, Louisiana. Association for Computational Linguistics.

Yukun Zhu, Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books.

## A Methods for Selecting Summary Sentences

We investigate Random, TF-IDF (Ramos, 2003), TextRank (Mihalcea and Tarau, 2004) methods to access the importance of selecting summary sentences.

**Random** For comparing purposes, we use a random sentence as the summary.

**TF-IDF** We take the sentence with the highest average TF-IDF score (Ramos, 2003) as the golden summary  $s^i$ . A word is assigned a score by TF-IDF that scales proportionally to the number of times the word appears in the document and is offset by the number of documents in the corpus that contain the word, which can be expressed as  $N_w \cdot \log(\frac{N_d}{N_{dw}})$ , where  $N_w$  is the word count,  $N_d$  is

the total number of documents and  $N_{dw}$  is the total number of documents containing the word.

**TextRank** TextRank (Mihalcea and Tarau, 2004) is a weighted graph with tokens as nodes and the similarity between nodes as edges. We use BERT (Devlin et al., 2018) to compute the similarities between sentences and then rank them based on the TextRank algorithm.

## B Model Baselines

In this paper, we use Transformer-XL, WritingPrompts, and Progressive WritingPrompts as baselines.

**Transformer-XL** Transformers with segment-level recurrence strategy (Dai et al., 2019) naturally constitutes a baseline. The model sequentially generates texts in a word-by-word fashion.

**WritingPrompts** first predicts a list of keywords or a single prompt, and then generates the full text given the prompt (Fan et al., 2018). Different from Fan et al. (2018), where golden prompts for stories are available, we do not readily have the golden prompts. We thus use the extractive strategies described in Section 4.3, i.e, the TF-IDF method to pick the keyword list as the prompt (denoted by *WritingPrompts-keyword*) and the reconstruction method to select the highest ranking sentence as the prompt (denoted by *WritingPrompts-sentence*).

**Progressive WritingPrompts** The progressive strategy proposed in Tan et al. (2020) which involves multiple stages of prompt generation. Each stage produces a more fine-grained sequence than the stage that comes before, and is used as the input to generate the prompt for the next stage. We follow the protocols in Tan et al. (2020) and use the TF-IDF score to obtain golden prompts for each stage. The number of stages is set to 4.