

# Few-Shot Table-to-Text Generation with Prefix-Controlled Generator

Yutao Luo , Menghua Lu , Gongshen Liu\*, Shilin Wang

Shanghai Jiao Tong University

{luoyt1996, 610228633, lgshen, wsl}@sjtu.edu.cn

## Abstract

Neural table-to-text generation approaches are data-hungry, limiting their adaptation for low-resource real-world applications. Previous works mostly resort to Pre-trained Language Models (PLMs) to generate fluent summaries of a table. However, they often contain hallucinated contents due to the uncontrolled nature of PLMs. Moreover, the topological differences between tables and sequences are rarely studied. Last but not least, fine-tuning on PLMs with a handful of instances may lead to overfitting and catastrophic forgetting. To alleviate these problems, we propose a prompt-based approach, Prefix-Controlled Generator (i.e., PCG), for few-shot table-to-text generation. We prepend a task-specific prefix for a PLM to make the table structure better fit the pre-trained input. In addition, we generate an input-specific prefix to control the factual contents and word order of the generated text. Both automatic and human evaluations on different domains (humans, books and songs) of the Wikibio dataset show substantial improvements over baseline approaches.

## 1 Introduction

Table-to-text generation is a significant branch of Natural Language Generation (NLG), aiming at generating descriptive text given an input table. There is a wide range of application scenarios for automatic table-to-text generation, such as sport news generation (Wiseman et al., 2017), story generation (Liu et al., 2020), weather forecasting report (Liang et al., 2009), and open-domain question answering (Chen et al., 2021).

Recent years have witnessed the great development of pre-trained language models (PLMs) (Devlin et al., 2019; Radford et al., 2019; Lewis et al., 2020), which achieve state-of-the-art performance on many text generation tasks, such as neural machine translation, document summarization, etc.

Unlike these tasks, table-to-text generation faces the lack of labeled data. Due to the development of data science, many statistical tables are generated in our daily life, but they scarcely have corresponding natural language descriptions, which limits the real-world application of data-hungry pre-trained models. To address this problem, researchers investigate workarounds in the few-shot setting. Chen et al. (2020), Gong et al. (2020) and Su et al. (2021a) leverage pre-trained linguistic knowledge of neural language models, then fine-tune them in target domains with limited labeled data. This “*pre-train and fine-tune*” paradigm performs relatively well in generating descriptive text from tables. Recently, another paradigm named “*pre-train and prompt*” has been proposed in order to adapt PLMs to downstream tasks without fine-tuning, which is more suitable for few-shot and zero-shot scenarios. Li and Liang (2021) prepends prompt tokens to adapt table-to-text generation to sequential generation task, and freezes PLMs’ weights to fully leverage their prior knowledge learned in the pre-training stage.

Despite their contributions, however, two main challenges for table-to-text generation remain to be explored, namely (1) the **topological structure difference** between tables and sequential inputs and (2) model’s ability to **select and rearrange factual content** from tables.

In order to address the aforementioned problems, we follow the “*pre-train and prompt*” paradigm and propose Prefix-Controlled Generator (i.e., PCG), an end-to-end generation framework along with two kinds of prefix tokens. Specifically, we prepend a task-specific (i.e., static) prompt and an input-specific (i.e., dynamic) prompt to the tabular input. The task-specific prompt aims to bridge the topological structure gap between a table and a word sequence, while the input-specific prompt aims to plan the factual content and the slot order of a table. Both prefixes are optimized during

\*Corresponding author.

Attribute(K)	Value(V)
name	edinho júnior
fullname	edon júnior viegas amaral
birth_date	7 march 1994
birth_place	faro , Portugal
currentclub	farese
clubnumber	21
position	forward

**Gold:** edon júnior viegas amaral , known as edinho júnior ( born 7 march 1994 ) is a portuguese **footballer** who plays as a forward for farese .

**Prefix-Tuning:** edon júnior ( born 7 march 1994 in faro , portugal ) is a portuguese **football midfielder** .

**BART:** edino júnior ( born 7 march 1994 ) is a portuguese **footballer** who plays for farese as a forward .

**Switch+PLM:** edon júnior viegas amaral or edinho júnior ( born 7 march 1994 in faro ) is a **portugal football player** and he plays for farese .

**Ours:** edinho júnior viegas amaral ( born 7 march 1994 ) is a portuguese **footballer** who plays as a forward for farese .

Figure 1: An example from *Wikibio Humans* domain and the generated descriptions via various approaches. Words in blue, red, orange and yellow indicate factual contents, wrong generation, hallucinated contents and inferred contents respectively. BART represents BART-large (Lewis et al., 2020). Switch+PLM represents Chen et al. (2020)’s approach with BART-large.

the training phase with the PLM remaining frozen, making our approach parameter-efficient – we only save one copy of the PLM while training in three different domains.

We basically follow the idea of prefix-tuning (Li and Liang, 2021) to design the task-specific prefix, except for some modifications. Firstly, due to the importance of a proper initialization of prefix tokens, we use task-relevant words (e.g., “*Summarize the following table:*”, or “*TL;DR:*”) as the initial prefix to better linearize the tabular input. Secondly, He et al. (2022) proves that the length of prefix tokens and the design of adding additional parameters solely on the attention module are two bottlenecks of prefix-tuning. Inspired by their work, we add Scaled Parallel Adapters (He et al., 2022) in parallel with both the attention layer and the feed-forward layer to improve the bottleneck of prefix-tuning.

For the input-specific prefix, we expect it can hint to the model which key-value pairs should be selected and in what order they should be arranged. Therefore, we propose a content planner to select the keys that appear in the gold summary and sort them according to the order of occurrence in the summary. For example, given a table in Fig. 1, we expect the content planner to generate a word sequence “*fullname name birth\_date birth\_place position currentclub*” that indicates all the keys and their occurrence order whose values appear in the gold summary. The word sequence will be used as hard prompts to feed into the PLM.

We evaluate our model on multi-domain table-to-text dataset (Chen et al., 2020). We show that our model outperforms previous state-of-the-art methods on both automatic evaluation metrics (§5.4) and human evaluation metrics (§5.5). We also con-

duct ablation studies to verify the effectiveness of the two kinds of prefixes (§5.6).

In a nutshell, our contributions are as follows:

1. We propose a Prompt-Controlled Generator that attends to the task-specific prefix to bridge the topological structure gap between tables and sequences, and the input-specific prefix to select factual contents from the tables and reorder them.
2. We propose a simple yet effective content planner to generate the input-specific prefix as the hard prompt of the PLM.
3. We conduct experiments on different domains of the Wikibio dataset to prove the effectiveness of our approach.

## 2 Related Work

### 2.1 Few-shot Table-to-text generation

Table-to-text generation has aroused much interest in recent years. Most of the existing studies resort to the end-to-end framework to generate fluent and faithful natural language descriptions given tables. Ma et al. (2019) firstly studied table-to-text generation under the low-resource constraint, and separated the generation process into two stages – key fact prediction and surface realization. With the advances of PLMs, many researchers fine-tune pre-trained GPT-2 (Radford et al., 2019) or BART (Lewis et al., 2020) to augment the scarce training data, which can better assist few-shot table-to-text generation. Chen et al. (2020) used copy mechanism (See et al., 2017) to improve the fidelity of sentences generated by GPT-2 by choosing to copy words from tabular input. Gong et al. (2020) adopted a unified GPT-2 model for table structure

reconstruction and generation. Zhao et al. (2021) proposed a token-level attention and a slot-level attention to exploit natural linguistic and table structural information. All these works utilized tabular input for free text generation, neglecting the importance of content planning for text fidelity. Su et al. (2021a) introduced an information retrieval (IR) system to select prototype sentences similar to the gold summary from large unlabeled parallel corpus, then use them as the auxiliary content plan for tabular input to generate natural language description. However, the IR system might see all gold summaries in the Wikipedia corpus, which violates the true few-shot setting. Different from the above studies, we focus on how to select factual contents via content planning, introducing a slot-aligned content planner.

## 2.2 Prompt Tuning for Generation

Prompt tuning is a nascent approach for natural language generation (NLG), first proposed by GPT-3 (Brown et al., 2020), introducing in-context learning for few-shot domain adaptation. Prefix-tuning (Li and Liang, 2021) prepended a sequence of continuous vectors to all examples of the downstream tasks. These vectors, which are adjusted as additional key-value pairs, steer the frozen PLMs by augmenting the left context at every Transformer layer. He et al. (2022) classified prefix-tuning as a parameter-efficient tuning approach similar to adapter (Houlsby et al., 2019) and made improvements on its bottlenecks. Clive et al. (2021) extended prefix-tuning to the input-specific prefix (e.g., topic of the datapoint, target output length) to have a finer-grained control for downstream generation tasks. Different from their work, we use the input-specific prompt not to guide the generated text in a certain style, but to improve the fidelity of generated text and the correctness of word order via content planning.

## 2.3 Controllable Text Generation

Controllable text generation (CTG) is a supplementary field for prompt-based generation, aiming at incorporating guidance signals into generative models. Control signals include text style (Keskar et al., 2019), grammar (Lyu et al., 2021), length (Kikuchi et al., 2016), etc. Recent CTG approaches involve generative adversarial networks (Yu et al., 2017), refactoring a PLM (Chan et al., 2021), fine-tuning adapted modules (Zeldes et al., 2020), prompt learning (Yu et al., 2021) and dif-

fusion model (Li et al., 2022). However, these approaches requires large amount of training data, which does not match our few-shot setting. For table-to-text generation, Su et al. (2021b) proposed a content planner to assist data-to-text generation, which inspired us to pre-plan the order and occurrence of the tabular input for improving the controllability of the generated text.

## 3 Problem Definition

Given a table  $T$  with  $n$  key-value pairs  $\{\mathbf{K}_i : \mathbf{V}_i\}_{i=1}^n$ , where  $\mathbf{K}_i = \{k_1^i, k_2^i, \dots, k_m^i\}$  and  $\mathbf{V}_i = \{v_1^i, v_2^i, \dots, v_{m'}^i\}$  refer to the key and the value of the  $i$ -th table slot respectively, we aim to generate a fluent and faithful natural language description of the table in a low-resource constraint. Note that  $\mathbf{K}_i$  and  $\mathbf{V}_i$  represent sequences of  $m$  and  $m'$  words respectively.

## 4 Methodology

We first provide intuition of using a task-specific prefix and an input-specific prefix for few-shot table-to-text generation (§4.1). Fig. 2 depicts the overall architecture of our method. As shown in the figure, given the input table, the content planner selects the factual contents and reorders them to form a dynamic prompt (§4.2). After that, a static prompt is designed and fed to the generative PLM along with the dynamic prompt (§4.3).

### 4.1 Intuition

The intuition of introducing prompt to few-shot table-to-text generation is that prompt-tuning effectively solves the catastrophic forgetting problem. Since table-to-text generation requires the language understanding ability of the table content, we hope to fine-tune downstream tasks while retaining the prior knowledge of PLMs, which is exactly what prompt-tuning does. Unlike model fine-tuning, which might be over-parameterized, prompt tuning only adjusts a few parameters and is less prone to over-fitting.

Observing Gong et al. (2020)’s experimental results, we find that table format transformation plays a vital role in improving the generation process, so we focus on bridging the topological structure gap between tables and sequential inputs. The first thought is that we can flatten a table into a word sequence using template (Gong et al., 2020). For example, given a table shown in Fig. 1, we serialize the key-value pair  $\{name: edinho junior\}$

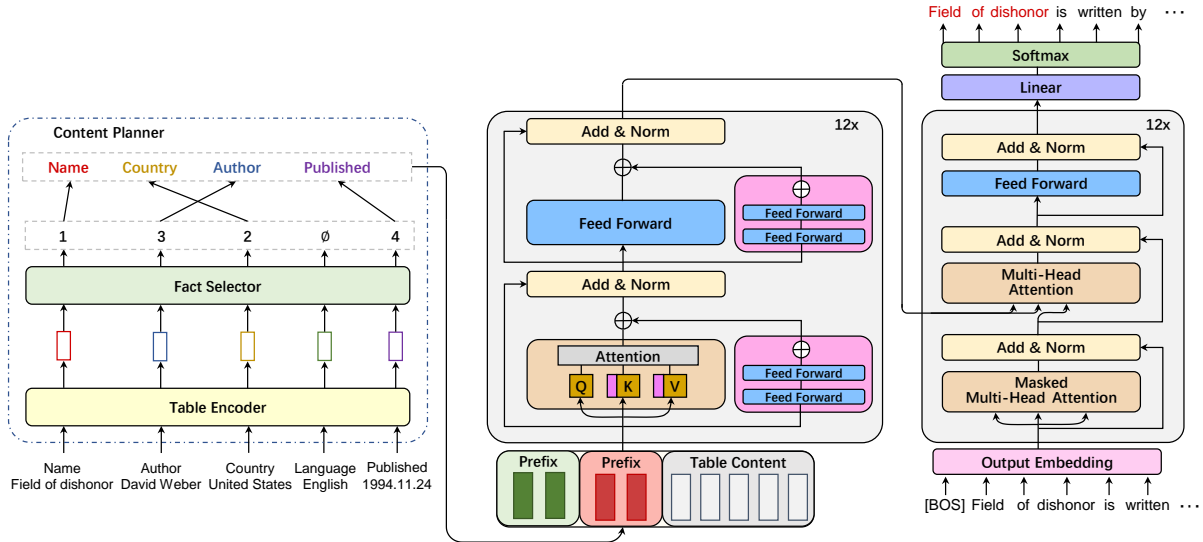


Figure 2: The overall architecture of the proposed method, which can be divided into Content Planner and Prompt-Controlled Generator. Tokens in red and yellow indicate these words are consistent with the value corresponding to the key “name” and “country” respectively.

as “name is edinho júnior;”, then concatenate all key-value pairs to form a sentence, that is, “name is edinho júnior; fullname is edon júnior viegas amaral; birth\_date is 7 march, 1994; ...”. Considering that the template-generated sentence is still somewhat different from the pre-training input, we want to find a way that adapts it to a natural sentence. Intuitively, we can add some prompt tokens like “summarize the following table:” to make the template-generated sentence an incidental component of the whole input. In this way, “summarize the following table” becomes the major component of the sentence, which is more similar to the sequential form of the pre-training input. In addition, many language models now have prefix LM pre-training tasks, which makes our sentences more consistent with the pre-training input.

We also seek to minimize the generated hallucinated content. Considering that some table slots are redundant, we intuitively want to hint the model what are the factual contents. To be consistent with the above table linearization approach, we follow the idea of controllable generation, providing a hard prompt as the guidance signal for each example to control both the table content to be selected and the word order.

## 4.2 Content Planner

Content Planner aims to generate input-specific prompts that guide the generation process in terms of factual contents and words order, which is shown

in Fig. 2 (left). Content Planner contains two modules, namely Table Encoder and Fact Selector. Since we study table-to-text generation under a strict few-shot constraint, we strive for simplicity of Content Planner. Therefore, we use a bi-directional LSTM and a linear-chain Conditional Random Fields (CRF) (Lafferty et al., 2001) for Table Encoder and Fact Selector respectively, which are learned given a handful of training instances.

Table Encoder takes all key-value pairs  $\{\mathbf{K}_i : \mathbf{V}_i\}_{i=1}^n$  from table  $T$  as the input, and produces a hidden representation  $\mathbf{h}_i \in \mathbb{R}^{d_e}$  for each table slot, where  $d_e$  is the hidden dimension. Specifically, for each table slot that contains a key-value pair  $\{\mathbf{K}_i : \mathbf{V}_i\}$ , we embed  $\mathbf{K}_i$  and  $\mathbf{V}_i$  by:

$$\mathbf{e}_i = \lambda \frac{1}{m} \sum_{j=1}^m \mathbf{E}(k_j^i) + (1 - \lambda) \frac{1}{m'} \sum_{j=1}^{m'} \mathbf{E}(v_j^i), \quad (1)$$

where  $\mathbf{e}_i$  denotes the embedding of the  $i$ -th slot,  $\mathbf{E}$  is the embedding lookup table and  $\lambda$  is a hyper-parameter that controls the ratio of key embedding and value embedding in  $\mathbf{e}_i$ . We use pre-trained Roberta (Liu et al., 2019) embedding to initialize  $\mathbf{E}$ . After that, we feed all embeddings  $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$  to the BiLSTM encoder to obtain  $\{\overrightarrow{\mathbf{h}}_1, \overrightarrow{\mathbf{h}}_2, \dots, \overrightarrow{\mathbf{h}}_n\}$  and  $\{\overleftarrow{\mathbf{h}}_1, \overleftarrow{\mathbf{h}}_2, \dots, \overleftarrow{\mathbf{h}}_n\}$  in the left-to-right and right-to-left directions respectively. The calculation of each direction uses a distinct set of parameters. The final hidden states  $\mathbf{h}_i$  can be obtained by:

$$\mathbf{h}_i = [\overrightarrow{\mathbf{h}}_i; \overleftarrow{\mathbf{h}}_i]. \quad (2)$$

Fact Selector selects key-value pairs that occur in the ground-true table summary, and rearranges them according to the order of occurrence in the summary. In practice, we use a standard CRF layer with a feed-forward layer as our Fact Selector to compute the global optimal sequence. On top of the hidden states  $\mathbf{H}_c = \{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n\}$ , the probability distributions of the label sequence  $\mathbf{y} = \{l_1, l_2, \dots, l_n\}$  is computed by:

$$P(\mathbf{y}|\mathbf{H}_c) = \frac{\exp(\sum_i (\mathbf{W}_{CRF}^{l_i} \mathbf{h}_i + \mathbf{M}_{l_{i-1}, l_i}))}{\sum_{\mathbf{y}'} \exp(\sum_i (\mathbf{W}_{CRF}^{l'_i} \mathbf{h}_i + \mathbf{M}_{l'_{i-1}, l'_i})}. \quad (3)$$

Here  $\mathbf{y}'$  represents an arbitrary label sequence,  $\mathbf{W}_{CRF}^{l_i}$  denotes the parameters specific to  $l_i$ , and  $\mathbf{M}_{l_{i-1}, l_i}$  denotes the transition score from  $l_{i-1}$  to  $l_i$ . The learning objective is defined as:

$$\mathcal{L}_{CRF} = -\log P(\mathbf{y}|\mathbf{H}_c). \quad (4)$$

Content Planner is trained independently with Prompt-Controlled Generator. The labeled key-value pair order is extracted from the ground-true summary by finding keys<sup>1</sup> and sorting them according to their positions. During inference, we use first-order Viterbi algorithm to decode the best label sequence  $\tilde{\mathbf{y}} = \text{argmax}_{\mathbf{y}'} P(\mathbf{y}'|\mathbf{H}_c)$ . Take Fig. 2 as an example, Content Planner generates a label sequence “1,3,2,∅,4”. The first label “1” denotes “Name” should appear in the front of the content plan, while the fourth label “∅” denotes “Language” does not occur in the gold summary. According to the label sequence, we rearrange all keys to form a content plan  $\mathbf{c}$ , which in Fig. 2 is “Name Country Author Published”.

### 4.3 Prompt-Controlled Generator

Prompt-Controlled Generator aims to generate fluent and faithful descriptions given the tabular input and the content plan. Our approach is model-agnostic, thus the generator could be any pre-trained generation model. Here we use BART-large (Lewis et al., 2020) as the basic generator for their best overall performances, and propose two kinds of prefixes that are prepended to the input of BART encoder, namely task-specific prompt  $\mathbf{p}_s$  and input-specific prompt  $\mathbf{c}$ . The latter (i.e., content plan) serves as the guiding signal of Prompt-Controlled Generator.

<sup>1</sup>Some keys such as “nationality” are fuzzy-matched.

The task-specific prompt is designed to bridge the topological structure gap between tables and sequences. A first thought is that we can linearize the table via template (see §4.1), then prepend discrete prompt words “summarize the following table:” to the template-generated sequence to make the tabular input more consistent with the pre-training input. Nevertheless, discrete optimization needs enormous computing power and human crafts. Instead of using discrete prompt, we follow prefix-tuning (Li and Liang, 2021) to optimize a sequence of continuous prefix tokens while keeping the PLM frozen. However, the prefix length and acting on the attention layer bound the presentation ability of the prefix (He et al., 2022). Considering these bottlenecks, we additionally parallel two Scaled Parallel Adapters to the attention layer and the feed-forward layer respectively, then perform scaled addition for these Adapters.

Next, we will introduce our modifications to BART encoder. Let us denote the template-generated sentence as  $\mathbf{s} = \{s_1, s_2, \dots, s_L\}$  and content plan as  $\mathbf{c} = \{c_1, c_2, \dots, c_{L_c}\}$ , where  $L$  and  $L_c$  are the lengths. The prefix length is denoted by  $L_p$ . We concatenate the content plan and the template-generated sentence (denoted by  $[\mathbf{c} : \mathbf{s}]$  where  $[\cdot : \cdot]$  is the concatenation operator) to feed into BART encoder. In the multi-head self-attention layer, we first compute the queries  $\mathbf{Q} \in \mathbb{R}^{(L+L_c) \times d}$ , keys  $\mathbf{K} \in \mathbb{R}^{(L+L_c) \times d}$  and values  $\mathbf{V} \in \mathbb{R}^{(L+L_c) \times d}$  via Equation (5):

$$\mathbf{Q} = \mathbf{x}\mathbf{W}_q, \mathbf{K} = \mathbf{x}\mathbf{W}_k, \mathbf{V} = \mathbf{x}\mathbf{W}_v, \quad (5)$$

where  $d$  denotes the hidden dimension of BART,  $\mathbf{W}_q$ ,  $\mathbf{W}_k$  and  $\mathbf{W}_v$  are trainable parameters.  $\mathbf{x}$  denotes  $\mathbf{E}_b([\mathbf{c} : \mathbf{s}])$  when the first layer is being computed, the output of the previous BART layer otherwise.  $\mathbf{E}_b$  denotes the embedding lookup table of BART. Then the attention score is computed via Equation (6):

$$\text{head} = [\text{head}^1 : \text{head}^2 : \dots : \text{head}^{N_h}], \quad (6)$$

where  $N_h$  denotes the number of heads.  $\text{head}^i$  is computed via Equation (7):

$$\begin{aligned} \text{head}^i &= \text{Attn}(\mathbf{Q}^i, \mathbf{K}^i, \mathbf{V}^i) \\ &= \text{softmax}\left(\frac{\mathbf{Q}^i [\mathbf{P}_k^i : \mathbf{K}^i]^T}{\sqrt{d_k}}\right) [\mathbf{P}_v^i : \mathbf{V}^i], \end{aligned} \quad (7)$$

where  $d_k = \frac{d}{N_h}$  denotes the hidden dimension of

Domain	Humans				Books				Songs			
Training set size	50	100	200	500	50	100	200	500	50	100	200	500
<b>Switch+GPT-2(R)</b>	25.7	29.5	36.1	41.7	34.3	36.2	37.9	40.3	36.1	37.2	39.4	42.2
<b>TableGPT(R)</b>	29.8	34.5	40.6	45.6	35.1	<b>37.3</b>	38.5	41.6	36.7	37.8	39.3	42.3
<b>Bart-large</b>	37.6	39.3	41.2	44.3	34.2	37.1	<b>39.8</b>	42.9	37.7	38.9	40.1	43.9
<b>AMG(R)</b>	-	-	-	49.0	-	-	-	43.9	-	-	-	<b>45.1</b>
<b>Hard-prompt+GPT-2</b>	22.8	28.1	29.7	30.8	25.8	27.9	28.8	32.1	26.6	30.0	30.1	32.1
<b>Prefix-Tuning+GPT-2</b>	25.6	30.3	33.4	37.3	34.9	36.2	36.3	37.3	32.5	33.0	35.1	36.1
<b>Prefix-Tuning+T5</b>	34.5	39.9	41.6	44.1	35.5	<b>37.3</b>	39.6	41.2	37.5	38.5	40.0	41.1
<b>Switch+BART(PT)</b>	36.8	41.8	44.0	48.1	33.6	35.0	38.3	43.4	<b>40.9</b>	<b>41.7</b>	42.1	43.2
<b>Ours</b>	<b>39.9</b>	<b>43.3</b>	<b>45.8</b>	<b>49.4</b>	<b>36.6</b>	36.9	39.0	<b>45.6</b>	38.0	<b>41.7</b>	<b>42.5</b>	44.5

Table 1: BLEU results on three domains of the Wikibio test set. Each (R) is reported by the related paper.

Domain	Humans				Books				Songs			
Training set size	50	100	200	500	50	100	200	500	50	100	200	500
<b>Switch+GPT-2(R)</b>	30.6	34.6	40.5	45.6	42.7	42.8	43.4	44.9	40.2	41.7	44.0	44.8
<b>Bart-large(R)</b>	37.8	41.4	47.4	45.5	41.7	43.4	43.7	48.1	41.7	42.4	44.1	46.0
<b>AMG(R)</b>	43.6	47.7	50.1	<b>51.9</b>	43.4	46.0	<b>47.5</b>	48.6	42.0	43.3	45.9	<b>46.9</b>
<b>Prefix-Tuning+GPT-2</b>	32.7	35.9	36.6	38.7	29.8	31.8	31.7	32.7	31.7	33.3	32.3	31.5
<b>Prefix-Tuning+T5</b>	39.3	40.6	41.8	42.1	32.8	34.8	36.0	36.8	34.4	36.1	36.0	34.6
<b>Switch+BART(PT)</b>	35.2	41.7	45.1	50.5	33.0	37.2	41.2	46.4	36.7	39.4	42.0	45.9
<b>Ours</b>	<b>46.7</b>	<b>48.3</b>	<b>50.4</b>	51.8	<b>46.3</b>	<b>46.2</b>	<b>47.5</b>	<b>49.3</b>	<b>44.8</b>	<b>45.7</b>	<b>46.9</b>	46.0

Table 2: PARENT-F results on three domains of the Wikibio test set. All (R) are reported by Zhao et al. (2021).

each head, and  $\mathbf{P}_k \in \mathbb{R}^{L_p \times d}$ ,  $\mathbf{P}_v \in \mathbb{R}^{L_p \times d}$  denote two sets of prefix vectors.  $\mathbf{Q}^i \in \mathbb{R}^{(L+L_c) \times d_k}$ ,  $\mathbf{K}^i \in \mathbb{R}^{(L+L_c) \times d_k}$ ,  $\mathbf{V}^i \in \mathbb{R}^{(L+L_c) \times d_k}$ ,  $\mathbf{P}_k^i \in \mathbb{R}^{(L_p) \times d_k}$  and  $\mathbf{P}_v^i \in \mathbb{R}^{(L_p) \times d_k}$  denote a block of  $\mathbf{Q}$ ,  $\mathbf{K}$ ,  $\mathbf{V}$ ,  $\mathbf{P}_k$ , and  $\mathbf{P}_v$  respectively.

In parallel with the multi-head self-attention layer, a Scaled Parallel Adapter is added:

$$head' = \mathbf{x} + s \cdot \text{ReLU}(\mathbf{x}\mathbf{W}_{down})\mathbf{W}_{up}, \quad (8)$$

$$\mathbf{H}_{attn} = head + head', \quad (9)$$

where  $\mathbf{W}_{down} \in \mathbb{R}^{d \times r}$  and  $\mathbf{W}_{up} \in \mathbb{R}^{r \times d}$  are down-projection and up-projection,  $r$  is the bottleneck dimension and  $\mathbf{x}$  denotes the same vector as in Equation (5).  $s \geq 1$  is a trainable scaling hyper-parameter. We use  $\mathbf{H}_{attn}$  to replace the original attention output to conduct residual connection and layer normalization. Similarly, we insert another Scaled Parallel Adapter in parallel with the Feed Forward layer to enhance its representation:

$$\begin{aligned} \mathbf{o} = & \text{ReLU}(\mathbf{x}\mathbf{W}_1 + \mathbf{b}_1)\mathbf{W}_2 + \mathbf{b}_2 \\ & + s \cdot \text{ReLU}(\mathbf{x}\mathbf{W}'_{down})\mathbf{W}'_{up}, \end{aligned} \quad (10)$$

where  $\mathbf{x}$  and  $\mathbf{o}$  denote the input and output of the Feed Forward layer respectively.  $\mathbf{W}_1$ ,  $\mathbf{b}_1$ ,  $\mathbf{W}_2$ ,  $\mathbf{b}_2$ ,  $\mathbf{W}'_{down}$ ,  $\mathbf{W}'_{up}$  are trainable parameters.

We conduct residual connection and layer normalization over  $\mathbf{o}$  to get the hidden states of BART encoder  $\mathbf{H}_{enc}$ , then feed  $\mathbf{H}_{enc}$  along with decoder input to a normal BART decoder for sentence generation. The decoder input is the right-shifted gold summary in the training phase, and a simple “[BOS]” in the inference phase to generate tokens autoregressively. Given the gold summary  $\mathbf{g}$ , the learning objective is the cross-entropy loss, defined as:

$$\mathcal{L}_{LM} = - \sum_{i=1}^{|\mathbf{g}|} \log P_{dec}(\mathbf{g}_i | \mathbf{g}_{<i}; \mathbf{H}_{enc}). \quad (11)$$

## 5 Experiments

### 5.1 Datasets and Hyper-Parameters

Following Chen et al. (2020), we evaluate our method on three different domains (i.e., *Humans*, *Books* and *Songs*) of the Wikibio dataset, denoted

as Wiki-Humans, Wiki-Songs and Wiki-Books respectively. For all three domains, we conduct experiments in few-shot settings by varying the training set size to 50, 100, 200 and 500. The validation size is set to 1000, and the remaining instances are used for testing, which counts 13587, 5252 and 11879 for humans, books and songs respectively.

We use BART-large as our basic generator using transformers library (Wolf et al., 2020), which shares 12 layers and 16 heads for both encoder and decoder. We set the hidden and the embedding dimension of Content Planner to 768 (Roberta-base embedding dimension), and the key-value ensemble ratio  $\lambda$  is set to 0.7. The learning rates of Content Planner and Prefix-Controlled Generator are set to  $2e-4$  and  $1e-5$  respectively, both optimized by AdamW (Loshchilov and Hutter, 2017). We train our PCG for 200 epochs, with a batch size of 10 on one NVIDIA GeForce RTX 3090 GPU. Prefix length  $L_p$  is set to 30, and the bottleneck size of Scaled Parallel Adapter  $r$  is set to 512.

## 5.2 Baseline Models

We compare previous state-of-the-art few-shot table-to-text generation approaches, serving as baseline models:

- (i) **Switch+PLM** (Chen et al., 2020): The first work that introduces PLMs to the few-shot NLG task. They propose a switch policy to choose whether to copy words from the table or to generate from GPT-2. We also implement a variant using BART-large to replace GPT-2 and tuning the BART-large model with our task-specific prompt (denoted as **Switch+BART(PT)**).
- (ii) **TableGPT** (Gong et al., 2020): A further study based on Switch+PLM that leverages GPT-2’s prior knowledge, while enhancing generation fidelity with two auxiliary tasks.
- (iii) **AMG** (Zhao et al., 2021): A pre-train and fine-tune approach with a multi-grain attention to both tokens and slots, and introduces memory mechanism to back-track the allocation of table slots.
- (iv) **BART-large** (Lewis et al., 2020): A powerful PLM for conditional generation, which is proved effective in the few-shot scenario (Zhao et al., 2021). We fine-tune it on our few-shot datasets to report its performance.

Model	50		100		200		500	
	acc.	BLEU	acc.	BLEU	acc.	BLEU	acc.	BLEU
Roberta-base	0.32	14.6	0.33	14.2	0.39	21.3	0.56	32.5
ContentPlanner	0.53	30.4	0.56	32.6	0.59	35.4	0.64	37.5

Table 3: Results on Content Planner. acc. and BLEU denote test accuracy and BLEU-2 respectively.

- (v) **Hard-prompt+GPT-2**: Our earlier attempt on the few-shot table-to-text generation task, which uses actual tokens such as "Summarize the following table:" as the prompt words, then feed the transformed tabular input into GPT-2 to fine-tune on few-shot table-to-text generation task.
- (vi) **Prefix-Tuning** (Li and Liang, 2021): A novel prompt-based approach that prepends a continuous prefix and freezes the PLMs to retain their prior knowledge. We follow Ding et al. (2021)’s implementation, using GPT-2 and T5 (Raffel et al., 2019) as the base model.

Among above baseline approaches, **Prefix-Tuning** and **Switch+BART(PT)** follow a “pre-train and prompt-tuning” paradigm (keep LM’s parameters frozen), while **Hard-prompt+GPT-2** uses prompt for model fine-tuning. All the other baselines are following the standard “pre-train and fine-tune” paradigm.

## 5.3 Results of Content Planner

We first report the experimental results of Content Planner. Intuitively, we use accuracy to evaluate the percentage of words that are both correct and in the right position. Following Zhao et al. (2020), we also use BLEU-2 to evaluate the correctness of the words occurring in the content plan. We train Content Planner in 200 epochs with 50/100/200/500 training instances respectively and compare it with *RobertaforSequenceClassification* from transformers library. The results are shown in Table 3. We show that in the few-shot setting, Bi-LSTM+CRF performs better than fine-tuning Roberta in both word co-occurrence and positional correctness.

## 5.4 Automatic Evaluation

We conduct automatic evaluations on various domains of the Wikibio dataset to prove the effectiveness of our method. We select two kinds of evaluation metrics – **BLEU** (Papineni et al., 2002) for evaluating overlap between the generated sentence and the gold description, and **PARENT** (Dhingra

et al., 2019) for evaluating both the matching between the generated sentence and the reference and the fidelity of the generated sentence to the original table. Here we use F1 score of PARENT, denoted as PARENT-F.

Regarding the overlapping-based metrics BLEU, we show that our method has the best overall performance compared with all other baselines. Specifically, our approach improves 1.8%/0.4% BLEU score on Wiki-Humans/Wiki-Songs compared with the second best model with 200 training instances. On Wiki-Books, we improve 1.7% BLEU score than AMG with 500 training instances. The results show that our method can produce fluent descriptions. We attribute this to the task-specific prefix that better linearizes the tabular input by comparing fine-tuning BART-large (see §5.6).

Regarding the fidelity-based metrics PARENT, our method has better performances over AMG especially in extremely low-resource scenarios, while outperforming other baseline models. Our method performs 1.6% PARENT-F better than AMG on average in 9 terms and 0.5% PARENT-F worse on average in 2 terms. Reviewing their approach, AMG uses the Wikibio dataset, which is very similar to the few-shot datasets, for task adaptive pre-training. In a real-world low-resource scenario, however, it’s less likely to obtain a large unlabeled corpus related to the target domain. Moreover, our approach is parameter-efficient and storage-saving. Therefore, we provide a more lightweight alternative with better generation fidelity than AMG.

We also implement a variant of Chen et al. (2020)’s work with some modifications. We replace the GPT-2 with BART-large, and use prompt-tuning instead of fine-tuning to generate sentences. Therefore, the encoder in **Switch+BART** is consistent with our BART encoder in Fig. 2. **Switch+BART(PT)** achieves the second best performance in text fluency evaluation, obtaining the highest BLEU score on 2 terms. However, it’s bad at keeping faithful to the original table when the training set size is small, which is contrary to the motivation of copy mechanism. A reasonable explanation is that the objectives of Prompt-tuning and fine-tuning are contradictory. Prompt-tuning expects that the continuous prefix can transfer to downstream tasks, while fine-tuning Pointer Generator (See et al., 2017) aims to learn to copy words and to decide whether to copy or to generate. This contradiction makes Pointer Generator unable to

get effective training, especially when lacking training instances. We print the selected words when the model switches to “copy” state, finding that the words are far from the tokens that should be copied. We also see from Table 2 that changing copy mechanism to the input-specific prefix significantly improves the text fidelity.

## 5.5 Human Evaluation

We randomly select 100 generated sentences (training set size is set to 500) and corresponding tables and references from the test set, then present them to three voluntary human evaluators. All volunteers are postgraduate students with extensive research experience in document summarization and natural language generation. Inspired by Chen et al. (2020), we assure that each sentence is evaluated according to its (1) *faithfulness to the table and the reference* and (2) *language fluency*. To evaluate the effectiveness of our Content Planner, we also evaluate the generated sentences according to their (3) *words order correctness*. In the first task, all evaluators count the number of facts  $n_{co}$  that co-occur in the table slot and the reference<sup>2</sup>, and the number of facts  $n_{hal}$  that contradict with/ miss from the table (i.e., hallucinated contents). The percentage of factual content is then computed through

$$f_p = \frac{\sum_{s \in \mathcal{S}} n_{co}^s}{\sum_{s \in \mathcal{S}} (n_{co}^s + n_{hal}^s)},$$

where  $\mathcal{S}$  denotes the select corpus. In the second task, we ask each evaluator to compare sentences in a sentence set (descriptions of an instance generated from various methods), then rank them based on their fluency and grammatical correctness. The ranking then is normalized to 0-1, the smaller the better. Finally, we average the normalized ranking of the 100 sentences to get  $r_{avg}$ . In the third task, all volunteers are asked to count the words order correctness. For example, given a ground-true content plan “Name Published Genre Author” and hypothesis “A push and a shove is a 2007 novel by Christopher Kelly.”, volunteers count the correct key pair order in the hypothesis, such as “Published” is in front of “Author”. “Genre” is not in the hypothesis, thus all its key pair order (“Name Genre”, “Published Genre”, “Genre Author”) are wrong. The words order accuracy  $acc_{wo}$  is averaged over all generated hypothesis. Human evaluation results are shown in Table 5. We compute the final score via  $f_p - r_{avg} + acc_{wo}$  to

<sup>2</sup>Here we do not define co-occurrence as exact-matching or fuzzy-matching, instead we ask volunteers to decide co-occurrence based on human knowledge.



Model	Humans	Books	Songs
PCG	<b>43.3/48.3</b>	36.9/46.2	<b>41.7/45.7</b>
PCG w/o c	43.2/47.2	37.3/44.8	41.5/44.5
PCG w/o c&SPA	41.9/46.0	<b>37.4/44.5</b>	39.6/44.3
PCG w/o p <sub>s</sub> &c&SPA	39.3/41.4	37.1/43.4	38.9/42.4

Table 4: Ablation study results on two kinds of prefixes. p<sub>s</sub>, c and SPA denote the task-specific prefix, the input-specific prefix and Scaled Parallel Adapter respectively. In each entry, a/b denotes the BLEU/PARENT-F score.

Model	$f_p$	$r_{avg}$	$acc_{wo}$	overall
<b>Switch+GPT-2</b>	0.62	0.58	0.71	0.75
<b>Prefix-tuning+T5</b>	0.73	0.28	0.79	1.24
<b>PCG</b>	<b>0.75</b>	<b>0.20</b>	<b>0.84</b>	<b>1.39</b>

Table 5: Human evaluation results.

measure the models’ performance, the larger the better.

## 5.6 Ablation Study

We conduct ablation studies to evaluate the improvement brought by the two kinds of prefixes we proposed. We experiment on all three datasets with the training set size of 100. The automatic results are shown in Table 4. Observing the results, we conclude that both task-specific prefix and input-specific prefix improve the fidelity of the generated sentences, while input-specific prefix contributes little to the text fluency. These conclusions are consistent with our intuitions, given that input-specific prefix aims to improve the faithfulness by planning the content. Through ablation, we show that prepending a continuous prefix to the encoder input performs better than fine-tuning BART in the few-shot scenario. In addition, adding Scaled Parallel Adapters to enhance the representation ability of prompt vectors has also proved to be effective.

## 6 Conclusion

In this paper, we propose Prompt-Controlled Generator, using two kinds of prompts to address current challenges in few-shot table-to-text generation. The task-specific prefix aims to bridge the topological structure gap between tables and sequences, which is learned via freezing the PLM and tuning the continuous prompt vectors. The input-specific prefix is designed to guide the generation process in terms of factual content and word order. We propose Content Planner to generate the input-specific prefix. Experiments on Wiki-Humans, Wiki-Books and Wiki-Songs datasets prove the effectiveness of

our method from the aspects of generation fluency and text fidelity to the table.

## Acknowledgements

We thank the anonymous reviewers for their constructive comments. This work was supported by the Joint Funds of the National Natural Science Foundation of China (Grant No. U21B2020). Gongshen Liu is the corresponding author.

## References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Alvin Chan, Yew-Soon Ong, Bill Pung, Aston Zhang, and Jie Fu. 2021. [Cocon: A self-supervised approach for controlled text generation](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Wenhu Chen, Ming-Wei Chang, Eva Schlinger, William Yang Wang, and William W. Cohen. 2021. [Open question answering over tables and text](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Zhiyu Chen, Harini Eavani, Wenhu Chen, Yinyin Liu, and William Yang Wang. 2020. [Few-shot NLG with pre-trained language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 183–190. Association for Computational Linguistics.
- Jordan Clive, Kris Cao, and Marek Rei. 2021. [Control prefixes for text generation](#). *CoRR*, abs/2110.08329.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

- Bhuwan Dhingra, Manaal Faruqui, Ankur P. Parikh, Ming-Wei Chang, Dipanjan Das, and William W. Cohen. 2019. [Handling divergent reference texts when evaluating table-to-text generation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4884–4895. Association for Computational Linguistics.
- Ning Ding, Shengding Hu, Weilin Zhao, Yulin Chen, Zhiyuan Liu, Hai-Tao Zheng, and Maosong Sun. 2021. [Openprompt: An open-source framework for prompt-learning](#). *arXiv preprint arXiv:2111.01998*.
- Heng Gong, Yawei Sun, Xiaocheng Feng, Bing Qin, Wei Bi, Xiaojiang Liu, and Ting Liu. 2020. [Tablegpt: Few-shot table-to-text generation with table structure reconstruction and content matching](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 1978–1988. International Committee on Computational Linguistics.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2022. [Towards a unified view of parameter-efficient transfer learning](#). In *International Conference on Learning Representations*.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Nitish Shirish Keskar, Bryan McCann, Lav Varshney, Caiming Xiong, and Richard Socher. 2019. [CTRL - A Conditional Transformer Language Model for Controllable Generation](#). *arXiv preprint arXiv:1909.05858*.
- Yuta Kikuchi, Graham Neubig, Ryohei Sasano, Hiroya Takamura, and Manabu Okumura. 2016. [Controlling output length in neural encoder-decoders](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1328–1338, Austin, Texas. Association for Computational Linguistics.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. [Conditional random fields: Probabilistic models for segmenting and labeling sequence data](#). In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001), Williams College, Williamstown, MA, USA, June 28 - July 1, 2001*, pages 282–289. Morgan Kaufmann.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7871–7880. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 4582–4597. Association for Computational Linguistics.
- Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. 2022. [Diffusion-lm improves controllable text generation](#). *CoRR*, abs/2205.14217.
- Percy Liang, Michael I. Jordan, and Dan Klein. 2009. [Learning semantic correspondences with less supervision](#). In *ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2-7 August 2009, Singapore*, pages 91–99. The Association for Computer Linguistics.
- Danyang Liu, Juntao Li, Meng-Hsuan Yu, Ziming Huang, Gongshen Liu, Dongyan Zhao, and Rui Yan. 2020. [A character-centric neural model for automated story generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(02):1725–1732.
- 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](#). *CoRR*, abs/1907.11692.
- Ilya Loshchilov and Frank Hutter. 2017. [Decoupled weight decay regularization](#). *arXiv preprint arXiv:1711.05101*.
- Yiwei Lyu, Paul Pu Liang, Hai Pham, Eduard Hovy, Barnabás Póczos, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2021. [StylePTB: A compositional benchmark for fine-grained controllable text style transfer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2116–2138, Online. Association for Computational Linguistics.
- Shuming Ma, Pengcheng Yang, Tianyu Liu, Peng Li, Jie Zhou, and Xu Sun. 2019. [Key fact as pivot: A two-stage model for low resource table-to-text generation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2047–2057. Association for Computational Linguistics.

- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [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.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *CoRR*, abs/1910.10683.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1073–1083. Association for Computational Linguistics.
- Yixuan Su, Zaiqiao Meng, Simon Baker, and Nigel Collier. 2021a. [Few-shot table-to-text generation with prototype memory](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 910–917. Association for Computational Linguistics.
- Yixuan Su, David Vandyke, Sihui Wang, Yimai Fang, and Nigel Collier. 2021b. [Plan-then-generate: Controlled data-to-text generation via planning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 895–909. Association for Computational Linguistics.
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2017. [Challenges in data-to-document generation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2253–2263. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Dian Yu, Zhou Yu, and Kenji Sagae. 2021. [Attribute alignment: Controlling text generation from pre-trained language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2251–2268, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. [Seqgan: Sequence generative adversarial nets with policy gradient](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 2852–2858. AAAI Press.
- Yoel Zeldes, Dan Padnos, Or Sharir, and Barak Peleg. 2020. [Technical report: Auxiliary tuning and its application to conditional text generation](#). *CoRR*, abs/2006.16823.
- Chao Zhao, Marilyn Walker, and Snigdha Chaturvedi. 2020. [Bridging the structural gap between encoding and decoding for data-to-text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2481–2491, Online. Association for Computational Linguistics.
- Wenting Zhao, Ye Liu, Yao Wan, and Philip S. Yu. 2021. [Attend, memorize and generate: Towards faithful table-to-text generation in few shots](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 4106–4117. Association for Computational Linguistics.

Table Content	Generated Description
<p><b>name:</b> cody zeller  <b>image:</b> cody zeller iu hoosiers standout in chicago img 5907.jpg  <b>position:</b> power forward/center  <b>height_in:</b> 9  <b>league:</b> nba  <b>number:</b> 40  <b>birth_date:</b> 5 october 1992  <b>high_school:</b> washington, indiana  <b>draft_year:</b> 2013  <b>draft_pick:</b> 4  <b>career_start:</b> 2013</p>	<p><b>article_title:</b> cody zeller  <b>height_ft:</b> 7  <b>weight_lb:</b> 240  <b>team:</b> charlotte bobcats/hornets  <b>nationality:</b> american  <b>birth_place:</b> washington, indiana  <b>college:</b> indiana (2011 - 2013)  <b>draft_round:</b> 1  <b>draft_team:</b> charlotte bobcats  <b>years:</b> 2013 start - present</p> <p><b>Gold:</b> cody allen zeller (born october 5, 1992) is an american professional basketball player who currently plays for the charlotte hornets of the national basketball association (nba).  <b>Plan:</b> <b>birth_date</b> <b>nationality</b> <b>draft_team</b> <b>league</b>  <b>Switch+BART:</b> cody zeller (born 5 october 1992 in washington, indiana) is an american basketball player for the charlotte hornets in the nba <b>as a power forward</b>.  <b>PCG:</b> cody zeller (born october 5, 1992) is an american professional basketball player who currently plays for the charlotte hornets of the national basketball association (nba).</p>
<p><b>name:</b> brandon pieters  <b>birth_date:</b> 22 april 1976  <b>height:</b> 1.93 0  <b>residence:</b> benoni , south africa  <b>tour:</b> sunshine tour  <b>sunwins:</b> 4</p>	<p><b>fullname:</b> brandon paul pieters  <b>birth_place:</b> germiston , south africa  <b>weight:</b> 106 kg lb st on  <b>yearpro:</b> 1994  <b>prowins:</b> 4  <b>article_title:</b> brandon pieters</p> <p><b>Gold:</b> brandon paul pieters (born 22 april 1976) is a south african professional golfer.  <b>Plan:</b> <b>full_name</b> <b>birth_date</b> <b>birth_place</b>  <b>Switch+BART:</b> brandon paul pieters (born 22 april 1976) is a south african <b>cricketer</b> who plays for <b>the central bank</b> of southern africa.  <b>PCG:</b> brandon paul pieters (born 22 april 1976) is a south african <b>wrestler</b>.</p>
<p><b>name:</b> david highbaugh smith  <b>alt:</b> a man with dark hair and a mustache wearing a dark coat and white shirt  <b>state:</b> kentucky  <b>predecessor:</b> john w. lewis  <b>birth_date:</b> 19 december 1854  <b>death_date:</b> 17 december 1928  <b>restingplace:</b> red hill cemetery  <b>profession:</b> lawyer</p>	<p><b>image:</b> david-highbaugh-smith.jpg  <b>district:</b> 4th  <b>successor:</b> ben johnson  <b>birth_place:</b> hart county , kentucky  <b>death_place:</b> hodgenville , kentucky  <b>party:</b> democrat  <b>article_title:</b> david highbaugh smith</p> <p><b>Gold:</b> david highbaugh smith (december 19, 1854 - december 17, 1928) was a u.s. representative from kentucky.  <b>Plan:</b> <b>name</b> <b>birth_date</b> <b>death_date</b> <b>state</b>  <b>Switch+BART:</b> davic highbaugh smith (1854 - 1928) was a u.s. lawyer from kentucky.  <b>PCG:</b> david highbaugh smith (december 19, 1854 - december 17, 1928) was an american representative from kentucky.</p>
<b>Table Content (1)</b>	<p><b>name:</b> james gilbert <b>birth_name:</b> cecil james gilbert <b>birth_date:</b> 5 may 1923 <b>birth_place:</b> edinburgh , scotland , united kingdom <b>occupation:</b> tv producer/director <b>television:</b> "the two ronnies" "last of the summer wine" "open all hours" <b>years_active:</b> 1959 – 2003 <b>article_title:</b> james gilbert (producer)</p>
<b>Generated Description</b>	cecil james gilbert (born 5 may 1923) is a scottish television producer and director.
<b>Table Content (2)</b>	<p><b>name:</b> petr faldyna <b>birth_date:</b> 11 july 1976 <b>birth_place:</b> frýdlant nad ostravicí , czechoslovakia <b>height:</b> 1.86 <b>position:</b> forward <b>currentclub:</b> fk senica <b>years:</b> 2001 2001-2002 2002-2003 2003-2004 2004-2006 2006 2006-2009 2009-2011 <b>clubs:</b> lerk prostějov hfk olomouc fk kunovice sfc opava fk kunovice sk České budějovice fc vysočina jihlava fk senica <b>caps:</b> 0 14 8 33 24 41 15 88 32 <b>goals:</b> 0 1 0 16 1 21 6 41 2 <b>pcupdate:</b> 2011-12-31 <b>article_title:</b> petr faldyna</p>
<b>Generated Description</b>	petr faldyna (born 11 july 1976) is a czech professional footballer who plays for fk senica as a forward.
<b>Table Content (3)</b>	<p><b>image_size:</b> 225px <b>position:</b> left wing <b>shoots:</b> left <b>height_ft:</b> 5 <b>height_in:</b> 7 <b>weight_lb:</b> 159 <b>status:</b> retired <b>former_teams:</b> modo hockey , skellefteå aik <b>birth_date:</b> 7 october 1975 <b>birth_place:</b> Örnsköldsvik , swe <b>draft:</b> 234th overall <b>draft_year:</b> 1996 <b>draft_team:</b> boston bruins <b>career_start:</b> 1993 <b>career_end:</b> 2012 <b>article_title:</b> anders söderberg</p>
<b>Generated Description</b>	anders söderberg (born october 7, 1975) is an american retired ice hockey player who played in the modo hockey league and the boston bruins of the nhl.

Table 6: Examples of generated results from Wiki-Humans test set. Words in red denotes wrong generation.

## A Appendix. Examples of generated sentences

In this section, we further provide generated examples of one baseline and our model. Table 6 upper displays the comparison between **Switch+BART(PT)** and our approach for generation quality. We show that our generation considers more on the content plan, while retaining the linguistic understanding of the table content. For example, in example 1, our method neglects contents that does not occur in the content plan, while the baseline method contains some tabular information that is not suitable to appear in the summary; In example 3, our method generate “*representative*” by understanding the table content such as “*predecessor*”, “*successor*” and “*party*”,

while **Switch+BART(PT)** simply copies the profession slot. Table 6 lower gives more examples of our generation. All sentences are generated by the model trained with 100 instances.

It’s worth mentioning that our generation quality is bounded by the adequacy of the table information. For example 2 in Table 6 upper, since the table does not provide information about “*golfer*”, it is difficult for the model to infer “*golfer*” from “*sunshine tour*” through few-shot training. To alleviate this problem, we are also investigating open-domain knowledge-graph enhanced table-to-text generation, which might further improve the generation quality.