

# SARG: A Novel Semi Autoregressive Generator for Multi-turn Incomplete Utterance Restoration

Mengzuo Huang<sup>1,2\*</sup>, Feng Li<sup>2\*</sup>, Wuhe Zou<sup>2</sup>, Weidong Zhang<sup>2†</sup>

<sup>1</sup> Dalian University of Technology

<sup>2</sup> Netease Games AI Lab, HangZhou, China

hmengzuo@mail.dlut.edu.cn

{lifeng06,zouwuhe,zhangweidong02}@corp.netease.com

## Abstract

Dialogue systems in open domain have achieved great success due to the easily obtained single-turn corpus and the development of deep learning, but the multi-turn scenario is still a challenge because of the frequent coreference and information omission. In this paper, we investigate the incomplete utterance restoration which has brought general improvement over multi-turn dialogue systems in recent studies. Meanwhile, jointly inspired by the autoregression for text generation and the sequence labeling for text editing, we propose a novel semi autoregressive generator (SARG) with the high efficiency and flexibility. Moreover, experiments on two benchmarks show that our proposed model significantly outperforms the state-of-the-art models in terms of quality and inference speed.

## Introduction

Dialogue systems in open-domain have attracted increasing attention (Li 2020; Huang, Zhu, and Gao 2020), and been widely utilized in real-world applications (Adiwardana et al. 2020; Gong et al. 2019; Hewitt and Beaver 2020). However, due to frequently occurred coreference and information omission, as shown in Figure 1, there still exists a major challenge: it is hard for machines to understand the real intention from the original utterance without the context. A series of models of retrieval-based and generative-based have been studied for multi-turn systems (Yan, Song, and Wu 2016; Zhang et al. 2018; Zhou et al. 2018; Wu et al. 2016), and they generally combine the context and the original utterance as input to retrieve or generate responses. However, these methods lack great generalizations since they have a strong reliance on the size of the multi-turn corpus.

Su et al. 2019 and Pan et al. 2019 propose their utterance restoration models, respectively, which are aimed at restoring the semantic information of the original utterance based on the history of the session from a different perspective. Restoration methods decouple multi-turn systems into the single-turn problems, which alleviate the dependence

\* Equal contribution. This work was conducted when Mengzuo Huang was interning at Netease Games AI Lab

† Corresponding author

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

	<b>Context 1</b>
Utterance 1 (Translation)	Human: 为什么? Human: <i>Why?</i>
Utterance 2	Chatbot: 这个你得问李淳风呀。 Chatbot: <i>You'll have to ask Li Chunfeng about that.</i>
Utterance 3	Human: 我去问他。 Human: <i>I'll ask him.</i>
Utterance 3'	Human: 我去问李淳风。 Human: <i>I'll ask Li Chunfeng.</i>
	<b>Context 2</b>
Utterance 1	Human: 你最喜欢什么电影? Human: <i>What movie do you like most?</i>
Utterance 2	Chatbot: 泰坦尼克。 Chatbot: <i>Titanic.</i>
Utterance 3	Human: 为什么呢? Human: <i>Why?</i>
Utterance 3'	Human: 为什么最喜欢泰坦尼克? Human: <i>Why do you like Titanic most?</i>

Figure 1: An example of utterance restoration in human-machine dialogue system. Utterance 3' is the restored sentence based on Utterance 3. Red means coreference and blue means omission.

on multi-turn dialogue corpus and also achieve leading performance. Specifically, Su et al. 2019 employ transformer-based Seq2Seq architecture and pointer network to rewrite the original utterance, and they split the whole session into history and original utterance for capturing different attentions. Pan et al. 2019 propose a cascade frame of “pick-and-combine” to restore the incomplete utterance from history. And both of them generate restored utterance from scratch in an autoregressive manner of Seq2Seq, which is highly time-consuming during inference.

Unlike some traditional end-to-end text generation task, where the apparent disparity exists between the sources and

targets, utterance restoration always has some considerable overlapping regions between inputs and outputs. Intuitively, some sequence labeling methods can be utilized to speed up the inference stage in this task, since Seq2Seq from scratch is time wasteful. Further, Malmi et al. 2019 introduce LaserTagger, a sequence labeling method, which casts text generation as a text editing task. However, the insertions of LaserTagger are restricted to a fixed phrase vocabulary that is derived from the training data. In multi-turn dialogue, some rare phrases are habitually omitted by the speaker without affecting the listening comprehension; as shown in Figure 1, “*Li Chunfeng*” is a rare phrase and omitted in Utterance 3 of Context 1. And LaserTagger can not solve such a coreference problem well, since the rare phrase is discarded when constructing the fixed phrase vocabulary.

As a first attempt to combine the sequence labeling and autoregression in utterance restoration, we propose a semi autoregressive generator (SARG), which can well tackle the challenges brought by highly time-consuming and discarded rare words or phrases. SARG retains the flexibility of autoregression and takes advantage of the fast inference speed of sequence labeling.

First, we employ a tagger to predict the editing labels, which involves three main operations: KEEP a token, DELETE a token, CHANGE a token with other phrases. Then, instead of adding phrases from a pre-defined phrase vocabulary, we utilize an autoregressive decoder based on LSTM with copy mechanism for generating the added phrases. Moreover, inspired by the great success of the pretrained transformer models (Vaswani et al. 2017; Radford et al. 2018), we also design an encoder based on BERT (Devlin et al. 2018) to obtain the contextual encodings. Finally, we perform experiments on two benchmarks: the *Restoration-200k* (Pan et al. 2019) and *CANARD* (Elgohary, Peskov, and Boyd-Graber 2019), the SARG shows superiorities on the automatic evaluation, the human evaluation, and the inference speed respectively. In summary, our contributions are:

- SARG is a creative fusion of sequence labeling and autoregressive generation, which is suitable for utterance restoration task;
- SARG solves the restoration problem by a joint way and can easily load the pretrained BERT weights for the overall model;
- SARG obtains a competitive performance and faster inference speed.

## Related Work

### Multi-turn Dialogue Systems

Recently, building a chatbot with data-driven approaches in open-domain has drawn significant attention (Ritter, Cherry, and Dolan 2011; Ji, Lu, and Li 2014; Athreya, Ngonga Ngomo, and Usbeck 2018). Most of works on conversational systems can be divided into retrieval-based methods (Ji, Lu, and Li 2014; Yan, Song, and Wu 2016; Zhou et al. 2016; Wu, Wang, and Xue 2016; Wu et al. 2016; Zhou et al. 2018; Zhang et al. 2018) and generation-based methods (Serban et al. 2016; Xing et al. 2016; Serban et al.

2017; Zhao, Xu, and Wu 2020; Lin et al. 2020). Though the above methods are enlightening, there is a lack of high-quality multi-turn dialogue data to train them.

In multi-turn dialogue systems, existing methods are still far from satisfactory compared to the single-turn ones, since the coreference and information omission frequently occur in our daily conversation, which makes machines hard to understand the real intention (Su et al. 2019). Recent studies suggest simplifying the multi-turn dialogue modeling into a single-turn problem by restoring the incomplete utterance (Su et al. 2019; Pan et al. 2019). Su et al. 2019 rewrite the utterance based on transformer-based Seq2Seq and pointer network from context with two-channel attentions. Pan et al. 2019 propose a cascaded “pick-and-combine” model to restore the incomplete utterance from its context. Moreover, Pan et al. 2019 release the high quality datasets *Restoration-200k* for the study of incomplete utterance restoration in open-domain dialogue systems.

### Sentence Rewriting

Sentence rewriting is a general task which has high overlap between input text and output text, such as: text summarization(See, Liu, and Manning 2017; Chen and Bansal 2018; Cao et al. 2018), text simplification(Wubben, van den Bosch, and Kraehmer 2012; Zhang and Lapata 2017; Wubben, van den Bosch, and Kraehmer 2012), grammatical error correction(Ng et al. 2014; Ge, Wei, and Zhou 2018; Chollampatt and Ng 2018; Zhao et al. 2019) and sentence fusion (Thadani and McKeown 2013; Lebanoff et al. 2019), ect. Seq2Seq model, which provides a powerful framework for learning to translate source texts into target texts, is the main approach for sentence rewriting. However, conventional Seq2Seq approaches require large amounts of training data and take low-efficiency on inference.

Malmi et al. 2019 propose a sequence labeling approach for sentence rewriting that casts text generation as a text editing task. And the method is fast enough at inference time with performance comparable to the state-of-the-art Seq2Seq models. However, it can’t be applied to our incomplete utterance restoration well, due to some limitations of inflexibility.

To make full use of the flexibility of autoregressive models and the efficiency of sequence labeling models, we combine the autoregressive generation and the sequence labeling for the trade-off between inference time and model flexibility.

## Methodology

In this section, we demonstrate our proposed SARG for the multi-turn incomplete utterance restoration. The restoration problem can be denoted as  $f(H, U) = R$ , where  $H = \{w_1^h, w_2^h, \dots, w_m^h\}$  is the history of dialogue (context),  $U = \{w_1^u, w_2^u, \dots, w_n^u\}$  is the original utterance (source) to be rewritten and  $R$  is the restored utterance (target). The overall architecture of SARG is shown in Figure 2. Instead of generating the restored utterance from scratch as traditional Seq2Seq, we first determine the editing operation sequence across the original utterance; then generate the potential phrases according to the operation sequence; finally

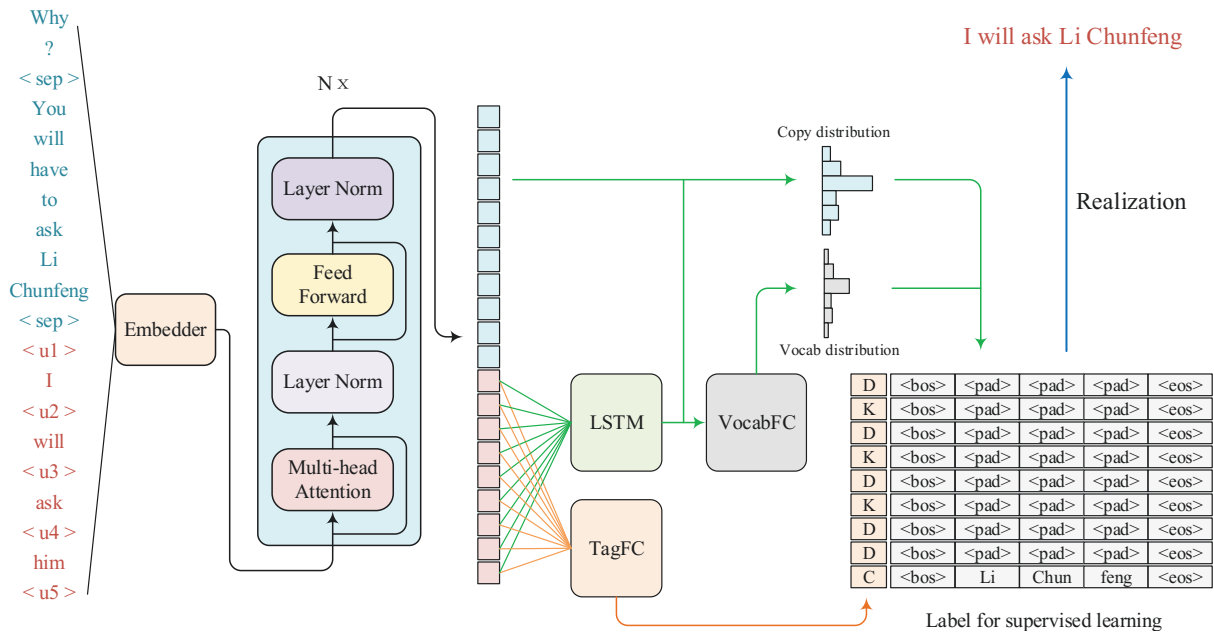


Figure 2: The overall architecture of the proposed SARG. In the constructed label, “D” means the DELETE operation, “K” means the KEEP, “C” means the CHANGE and the phrase of “Li Chunfeng” is the added phrase for this CHANGE operation. In the input, the blue words are the history of the session, the red words are the original utterance and the  $\langle u_i \rangle$  is the dummy token. In the dataflow, the black means encoding, the orange means tagging, the green means decoding, and the blue means the realization.

convert the operation sequence and the generated phrases to text. The detailed descriptions are as follows.<sup>1</sup>

### Tagging Operations

First of all, meaningless dummy tokens are inserted between every two tokens in the original utterance, as shown in first column of Figure 2. We can directly add the phrases in the gaps between every two tokens by the insertion of dummy tokens, which eliminates the ambiguity of possible editing operations to some extent. Moreover, we recommend that the original tokens can only be kept or deleted, and the dummy tokens can only be deleted or changed by other phrases.

Formally, three editing operations are defined in this work: KEEP, DELETE and CHANGE. Intuitively, KEEP means that the token remains in the restored utterance, DELETE means that the token is undesired, and CHANGE A means that the token should be replaced by the informative phrase A.

The following steps are employed to construct the supervised labels: (1) first compute the *longest common subsequence* (LCS) between original and restored utterance; (2) then greedily attempt to align the original utterance, restored utterance and the LCS; (3) finally replace the undesired tokens in original utterance with the added tokens in restored utterance. The detailed descriptions are demonstrated in Algorithm 1, and the constructed labels can be referenced in

<sup>1</sup>By convention, the bold letters represent the vectors, the capital letters represent the matrices and others represent the scalars.

---

#### Algorithm 1: Convert the target to label

---

**Input:**  $S$ : the original utterance  
 $T$ : the restored utterance  
**Output:**  $L$ : the supervised label

- 1 Insert dummy tokens in  $S$
- 2  $L[i] = \text{DELETE}, \forall i = 1, 2, \dots, 2n + 1$
- 3  $j = 0; k = 0; A = []$
- 4 Compute the *longest common subsequence*  $K$  between  $S$  and  $T$
- 5 **for**  $i \in [1, 2n + 1]$  **do**
- 6     **if**  $S[i] = K[k]$  **then**
- 7          $L[i] = \text{KEEP}$
- 8         **while**  $T[j] \neq K[k]$  **do**
- 9              $A = A + T[j]$
- 10             $j = j + 1$
- 11         **end**
- 12          $k = k + 1$
- 13         **if**  $A \neq \emptyset$  **then**
- 14              $L[i - 1] = \text{CHANGE } A$
- 15              $A = []$
- 16         **end**
- 17     **end**
- 18 **end**
- 19 **if**  $T[j:] \neq \emptyset$  **then**
- 20      $A = T[j:]$
- 21      $L[-1] = \text{CHANGE } A$
- 22 **end**
- 23 **return**  $L$

---

	Added Phrase	Restored Utterance
Avg. length	3.1	12.4

Table 1: Comparison of the average length between the added phrase and the restored utterance on *Restoration-200k*.

Figure 2. Specifically, the first column of labels is used to supervise the tagger and other columns are used for the decoder.

Moreover, the comparison of average length between added phrase and the restored utterance is listed in Table 1, which indicates that SARG saves at least three-quarters of the time for decoding compared to those complete autoregressive model.

## Encoder

Since pretrained transformers (Vaswani et al. 2017) have been shown to be beneficial in many downstream NLP tasks (Radford et al. 2018; Devlin et al. 2018), in this work, we utilize the standard transformer blocks as the backbone of the encoder, like the black lines in Figure 2.

In the embedding module, we concatenate the history  $H$  and the original utterance  $U$  (involved dummy tokens) as the input sequence  $W = \{w_1, w_2, \dots, w_k\}$ , then embed them into continuous space by looking up the following embedding tables:

- **Word Embedding:** the word embedding table is built on a pre-defined wordpiece vocabulary from pretrained transformers.
- **Position Embedding:** the position embedding table is also initialized by pretrained transformers.
- **Turn Embedding:** turn embedding is used to indicate which turn each token belongs to. The looking-up table is randomly initialized.

For each token  $w_i$ , we sum and normalize (Ba, Kiros, and Hinton 2016) the above three embeddings, then acquire the input embedding:

$$E_i^{(0)} = \text{LN}(WE(w_i) + PE(w_i) + TE(w_i)), \quad (1)$$

where  $WE$  is the word embedding,  $PE$  is the position embedding and  $TE$  is the turn embedding. Once the input embedding is acquired, we feed such representation into the  $L$  stacked transformer blocks:

$$E^{(l)} = \text{TransformerBlock}(E^{(l-1)}). \quad (2)$$

At last, we obtain the final encodings  $E^{(L)}$ , which can be further divided into two parts according to the partitions of history and original utterance:

$$E_h = \{h_1, h_2, \dots, h_m\}, \quad (3)$$

$$E_u = \{u_1, u_2, \dots, u_{2n+1}\}, \quad (4)$$

where  $E_h$  is the encodings of history and  $E_u$  is the encodings of original utterance. There are  $n + 1$  dummy tokens in the original utterance, which collect the information from those original tokens by the self-attention.

## Tagger

Tagger takes the encodings  $E_u$  as the input and predicts the editing labels on each token in original utterance. As shown in Figure 2, the orange lines stand for the dataflow of tagger. In our setting, a single linear transformation layer with softmax activation function is employed for projecting the encoding to the space of editing labels, the formula is as follows:

$$p(y_i|u_i) = \text{softmax}(W_t \cdot u_i + b_t), \quad (5)$$

where  $W_t$  and  $b_t$  are parameters to be learned, and the following  $W$  and  $b$  are all learnable. Finally, the loss provided by the tagger is defined as negative log-likelihood:

$$\text{loss}_{tag} = - \sum_i \log p(y_i|u_i), \quad (6)$$

where  $i$  is corresponding to the index of token in original utterance.

## Decoder

Different from the general autoregressive decoder that performs decoding from scratch, in our setting, the decoder, as green lines in Figure 2, works in parallel on the tokens which get CHANGE operations in tagger. Specifically, the decoder is only one and shared by these tokens.

For the consideration of efficiency, we employ one layer of unidirectional LSTM (Hochreiter and Schmidhuber 1997) as the backbone of our decoder. For each token in original utterance, the related initial state  $s_0$  is initialized with the according hidden representation <sup>2</sup>:

$$s_0 = u_i \in E_u. \quad (7)$$

Then the autoregressive generation is described as follows:

$$s_t = \text{LSTM}(WE(x_t), s_{t-1}), \quad (8)$$

where  $x_t$  is the output of decoder in the previous step, and the  $x_1$  is initialized by a special start token.

Moreover, in order to dynamically choose copying from the history or sampling from the overall vocabulary, we introduce the recurrent attention and coverage mechanism as in pointer-generator network (See, Liu, and Manning 2017). At each decoding step, we utilize the output  $s_t$  to collect information from the encodings of history  $E_h$ . The detailed calculations are as follows:

$$e_j^t = v^T \tanh(W_s s_t + W_h h_j + w_c c_j^t + b_{attn}), \quad (9)$$

$$a^t = \text{softmax}(e^t), \quad (10)$$

where  $j$  is corresponding to the index of token in the history,  $t$  is corresponding to the decoding steps and the  $c^t$  is the coverage vector in  $t$ -th step. Specifically, the coverage vector is initialized by zero at the beginning of decoding and accumulated as follow:

$$c_j^t = \sum_{t'=0}^{t-1} a_j^{t'}. \quad (11)$$

<sup>2</sup>It is a remarkable fact that there is a one-to-one correspondence between the hidden representation  $u_i$  and the state  $s_0$ , however, we omit the subscript  $i$  in  $s_0$  for the convenient expression.

Model	$f_1$	$f_2$	$f_3$	BLEU-1	BLEU-2	ROUGE-1	ROUGE-2
CopyNet	50.3	41.1	34.9	84.7	81.7	89.0	80.9
T-Ptr- $\lambda$	51.0	40.4	33.3	90.3	87.4	90.1	83.0
PAC <sup>‡</sup>	<b>63.7</b>	49.7	40.4	89.9	86.3	91.6	82.8
Seq2Seq-Uni <sup>‡</sup>	56.8	46.4	39.8	90.8	88.3	91.4	85.0
SARG <sup>‡</sup>	62.4	<b>52.5</b>	<b>46.3</b>	<b>92.2</b>	<b>89.6</b>	<b>92.1</b>	<b>86.0</b>

Table 2: The main results on *Restoration-200k* of our method and other SOTA methods. The models with “<sup>‡</sup>” means that pre-trained weights like BERT are utilized. Except to SARG, other models employ the 5-beam-search in their decoding procedure. SARG employs the greedy search in decoding step.

Once the normalized weights  $\mathbf{a}^t$  are obtained, we can calculate the results of attention:

$$\mathbf{s}_t^* = \sum_j \mathbf{a}_j^t \cdot \mathbf{h}_j. \quad (12)$$

Then, the  $\mathbf{s}_t^*$  is forwarded into the subsequent modules for acquiring the predicted word:

$$g = \sigma(\mathbf{w}_s^T \mathbf{s}_t^* + \mathbf{w}_s^T \mathbf{s}_t + \mathbf{w}_x^T \text{WE}(\mathbf{x}_t) + b_g), \quad (13)$$

$$\mathbf{p}_{vocab} = \text{softmax}(W_v \cdot \mathbf{s}_t^* + b_v), \quad (14)$$

$$\mathbf{p}(\mathbf{x}_{t+1}) = g \cdot \mathbf{p}_{vocab} + (1 - g) \sum_{j:w_j=\mathbf{x}_t} \mathbf{a}_j^t, \quad (15)$$

where  $\sigma$  is the sigmoid function to output a value between 0 and 1, the  $g$  is the gate to make a trade-off between copying and generating, the  $\mathbf{p}(\mathbf{x}_{t+1})$  is the final probability distribution of generated word. Moreover, the coverage loss is introduced to penalize repeatedly attending:

$$\text{covloss}_t = \sum_j \min(\mathbf{a}_j^t, \mathbf{c}_j^t). \quad (16)$$

Finally, the loss of the decoder is the weighted sum of negative log-likelihood and the coverage loss:

$$\text{loss}_{dec} = \sum_i \sum_t -\log \mathbf{p}(\mathbf{x}_t^i) + \lambda \text{covloss}_t^i, \quad (17)$$

where  $i$  is corresponding to the index of token in original utterance,  $\lambda$  is the hyperparameter for adjusting the weight.

## Joint Training

The model is optimized jointly. Once the loss of tagger and decoder are obtained, we sum and backward propagate the total loss as below:

$$\text{loss} = \alpha \text{loss}_{tag} + \text{loss}_{dec}, \quad (18)$$

where  $\alpha$  is also the hyperparameter for adjusting the weight.

## Realization

In the realization, we convert the predicted editing labels and the generated phrases to a complete utterance. In detail, we remain the `KEEP` denoted token and remove the `DELETE` token in the original utterance (involved dummy tokens), and replace the token, assigned by `CHANGE A`, with the generated phrase `A`.

	train	dev	test
<i>Restoration-200k</i>	194k	5k	5k
<i>CANARD</i>	32k	4k	6k

Table 3: The count of conversations in different datasets.

## Experiments

In this section, we first detail the experimental settings and the compared methods; then the main results and ablation study are described; finally, we report the human evaluation results and additional analysis based on some cases. Our experiments are conducted on *Restoration-200K* (Pan et al. 2019) and *CANARD* (Elgohary, Peskov, and Boyd-Graber 2019). The statistics of the datasets are shown in Table 3.

### Experiment Settings

We initialize SARG with RoBERTa-wwm-ext (Cui et al. 2019) for *Restoration-200k* and bert-base-uncased (Devlin et al. 2018) for *CANARD*, the hidden size is set to 768, the number of attention heads to 12, the number of attention layers to 12. Adam optimizer is utilized, the loss of tagger weighted to  $\alpha = 3$ , coverage loss weighted to  $\lambda = 1$  and the initial learning rate is  $5e-5$ . The above hyperparameters are all tuned on the standard validation data, and our code is available on <https://github.com/NetEase-GameAI/SARG>.

The according automatic evaluation metrics are utilized as in previous works (Pan et al. 2019; Elgohary, Peskov, and Boyd-Graber 2019), which contain BLEU, ROUGE, and restoration score.

### Compared Methods

We compare the performance of our proposed SARG with the following methods:

- **CopyNet**: in this baseline, LSTM-based Seq2Seq model with attention and a copy mechanism is employed.
- **PAC** (Pan et al. 2019): this model restores the incomplete utterance in a cascade way: firstly, select the remained words by finetuning BERT, and then roughly concatenate the selected words, history, original utterance and feed them into a standard pointer-generator network.
- **T-Ptr- $\lambda$** <sup>3</sup> (Su et al. 2019): this model solves such restoration task in an end-to-end way. It employs six layers of

<sup>3</sup>We re-implement the transformer-based method and evaluate on the same blind test set for the fair comparison.

	$f_1$	$f_2$	$f_3$	BLEU-1	BLEU-2	ROUGE-1	ROUGE-2
SARG	62.4	52.5	46.3	92.2	89.6	92.1	86.0
w/o WEIGHT	52.8	41.1	33.8	89.2	86.7	89.9	83.6
w/o COPY	55.6	38.9	32.8	89.4	85.6	89.9	81.7
w/o GEN	56.2	48.0	42.9	90.4	88.2	91.4	85.6

Table 4: Ablation study of proposed model on the *Restoration-200K*. The beam size is fixed to 1.

transformer blocks as encoder and another six layers of transformer blocks as pointer decoder. Moreover, to emphasize the difference between history and utterance, it takes two individual channels in the encoder-decoder attention.

- **Seq2Seq-Uni**: we construct this baseline by employing the unified transformer blocks (Dong et al. 2019) as the backbone of Seq2Seq, so that we can load the pretrained transformers easily.

## Main Results

The main results on *Restoration-200k* are as shown in Table 2. Focusing on the automatic metrics, we observe that SARG achieves the best results on 6 of 7 automatic metrics. The superiority of SARG is as expected, on one hand, the words of original utterance can be easily kept or deleted by sequence labeling, on the other hand, the rest of words can be easily copied from history or generated from vocabulary. And we also find PAC is 1.2 higher than SARG on restoration  $f_1$  score but 3.0 and 6.1 lower on  $f_2$  and  $f_3$  separately. In fact,  $f_1$  pays more attention to those tokens restored from history than others from the original utterance. In other words, though PAC can recall appropriate restored tokens from history, it may not place these restored tokens in their right positions well. We also exemplify such problem in the case study. Additionally, we compare the results of the beam-search with those of the greedy-each, which we find that the beam-search brings pretty significant improvements on those complete autoregressive models, but less obvious on our model. It means that, SARG is less dependent on the beam-search and can be more time-efficient in the inference phase.

	Dev	Test
CopyNet	51.37	49.67
T-Ptr- $\lambda$	46.26	45.37
Seq2Seq-Uni	52.71	45.31
SARG	<b>56.93</b>	<b>54.80</b>
Human Rewrites	59.92	

Table 5: The main results on *CANARD* of our method and other SOTA methods. Table shows the BLEU scores of the listed models on development and test data.

Table 5 shows the main results on *CANARD* dataset. As can be seen, SARG achieve the best BLEU score<sup>4</sup> on the development and test data. It is 5.56 higher than the previous

<sup>4</sup>We use multi-bleu-detok.perl (Sennrich et al. 2017) as in (Elgohary, Peskov, and Boyd-Graber 2019)

best on the development data and 5.13 higher on the test data. Moreover, we also find that the result of our method is far from the level of human rewrites, which means that there is still a large room for the improvement of existing rewriting methods.

	Inference Time
T-Ptr- $\lambda$ (n_beam=1)	522 s
T-Ptr- $\lambda$ (n_beam=5)	602 s
Seq2Seq-Uni (n_beam=1)	321 s
Seq2Seq-Uni (n_beam=5)	467 s
SARG (n_beam=1)	50 s
SARG (n_beam=5)	70 s

Table 6: The inference time on *Restoration-200k*, which is evaluated on the same blind test set (5104 examples) with one Nvidia Tesla P40. We do not consider the inference speed of PAC, because the cascade way takes lower efficiency than other end-to-end methods

Through the Table 6, we can observe that, compared to those complete autoregressive methods, our semi autoregressive model takes less time for inference. SARG is near 10x times as fast as T-Ptr- $\lambda$  and 6x times as fast as Seq2Seq-Uni. Beam-search increases the burden on inference. It needs more time and more memory for maintaining the candidate beams. Generally, the incomplete utterance restoration is required to be time-efficient as the intermediate subtask of multi-turn dialogue task, and it is impractical to maintain plenty of beams in decoding. Therefore, our SARG may be a suitable choice with the less dependent on the beam-search.

## Ablation Study

In this subsection, we conduct a series of ablation studies to evaluate the effectiveness of different components in our proposed SARG, which includes pretrained weights (WEIGHT), copy mechanism (COPY), and generation from vocabulary (GEN), and the results are shown in Table 4.

As can be seen, GEN plays the least important role in our model. By contrast, the absence of COPY or WEIGHT may raise a substantial lack of performance. Following our previous experimental setting, the above two variant models both can not converge well. In fact, the model without COPY only selects words from the pre-defined overall vocabulary, and the decoder is more difficult to be trained well. Furthermore, without the WEIGHT, the model needs to update the overall weights from scratch, which incorporate the 768-dimensional embedding table and 12 transformer lay-

	Example 1	Example 2	Example 3
$A_1$	男185 Male, 185	天蝎座有喜欢的吗 Does anyone like Scorpio	比尔吉沃特有互送皮肤的吗 Who want to exchange skin in Bilgewater
$B_1$	老乡你帖子要沉了 Bro, your post is totally ignored	天蝎座的小伙子你喜欢么 Do you like Scorpio boy	直接买个蘑菇的不就行了 Just buy the Timo's directly
$A_2$	这不还有你么哈哈 You're still here haha	喜欢啊 Yes	抽奖比较欢乐 Lucky draw is interesting
$B_2$	你雇我给你吆喝啊 You can hire me to cry out for you	可是不信星座谈 I don't believe in constellations	我一般都是从我同学的号挨个送我 I usually send myself from classmate's account
$A_3$	哈哈好 OK	这东西只是娱乐罢 It is just entertainment	看来我只能拿朋友的号送自己了 I'll have to send myself from friends' account
Reference	哈哈给我吆喝好 OK, <b>cry out for me</b>	星座这东西只是娱乐罢 <b>Constellation</b> is just entertainment	看来我只能拿朋友的号送 <b>皮肤</b> 自己了 I'll have to send myself <b>skin</b> from friends' account
SARG	哈哈好雇你给我吆喝 OK, <b>hire you to cry out for me</b>	星座这东西只是娱乐罢 <b>Constellation</b> is just entertainment	看来我只能拿朋友的号送自己 <b>皮肤</b> 了 I'll have to send myself <b>skin</b> from friends' account
PAC	哈哈吆喝好 OK, <b>cry out</b>	这东西只是娱乐 <b>星座</b> 罢 It is just entertainment <b>constellation</b>	看来我只能拿朋友的号送自己了 I'll have to send myself from friends' account
T-Ptr- $\lambda$	哈哈好吆喝 OK, <b>cry out</b>	<b>不信星座</b> 这东西只是娱乐罢 <b>Not believing constellation</b> is just entertainment	看来我只能拿朋友的号送自己了 I'll have to send myself from friends' account
Seq2Seq-Uni	哈哈好 OK	<b>不信星座</b> 这东西只是娱乐罢 <b>Not believing constellation</b> is just entertainment	看来我只能拿朋友的号送自己了 I'll have to send myself from friends' account

Figure 3: Examples for incomplete utterance restoration.  $A_1$  to  $B_2$  is the history of conversation,  $A_3$  is the original utterance.

ers. Therefore, it is a considerable burden for the optimization, where the limited corpus is provided.

And we also compare the output of tagger among the above listed models. An observation is that the tagger without WEIGHT is conservative on predicting the CHANGE operations; by contrast, the decoder without WEIGHT is less affected and has normal-appearing. Therefore, in some cases, even though the decoder produces the right restored words, the model still can not output the correct answers because the tagger does not produce the corresponding CHANGE operations.

## Human Evaluation

	Quality	Fluency
SARG	<b>2.70</b>	2.85
PAC	2.67	2.83
T-Ptr- $\lambda$	2.58	2.80
Seq2Seq-Uni	2.65	<b>2.87</b>

Table 7: Human evaluation of the restoration quality and language fluency on *Restoration-200k*. Both quality and fluency score adopt a 3-point scale.

In the phase of human evaluation, we employ three experienced workers to score the restoration quality and sentence fluency separately on 200 randomly selected samples. Specifically, each sample is scored by the three workers in turn, and the final quality or fluency scores are calculated by averaging the annotated results.

As shown in Table 7, SARG obtains the highest score in restoration quality among the compared methods, which is consistent with the results of automatic evaluation. However, in the aspect of fluency score, Seq2Seq-Uni achieves the best performance. Seq2Seq-Uni takes a way of complete autoregression and benefits from the pretrained weights, which can

complete the causal language modeling well.

## Case Study

In this subsection, we observe the prediction results among different models, and then select several representative examples to illustrate the superiority of our proposed model as Figure 3 shows.

As can be seen in Example 1, the first three models can restore the action “*cry out*”, and only SARG can restore the predicate “*hire you*”, which is important to understand the direction of the action.

In Example 2, all four models restore the keyword “*constellation*” correctly. However, for T-Ptr- $\lambda$  and Seq2Seq-Uni, undesired words “*not believing*” are also restored, which changes the intention of utterance. In PAC, we can find the keyword “*constellation*” is placed in a wrong position, which leads to the difficulty in understanding. Moreover, for the restoration scores, the wrong position problem has no effect on  $f_1$  but is negative for  $f_2$  and  $f_3$ . That is a possible reason, compared with SARG, PAC has higher  $f_1$  but lower  $f_2$  and  $f_3$  in the automatic evaluation.

Finally Example 3 demonstrates the ability of SARG to restore utterance from distant context. Specifically, the keyword “*skin*” appears in  $A_1$ , and the model is required to restore it after three utterances.

## Conclusion

In this paper, we propose a novel semi autoregressive generator for multi-turn incomplete utterance restoration. The proposed model takes in the high efficiency of inference time from sequence labeling and the flexibility of generation from autoregressive modeling. Experimental results on two benchmarks demonstrate that the proposed model is significantly superior to other state-of-the-art methods and an appropriate model of utterance restoration for boosting the multi-turn dialogue system.

## Acknowledgments

We thank Hongbo Zhang, Xiaolei Qin, Fuxiao Zhang and all the anonymous reviewers for their valuable comments.

## References

- Adiwardana, D.; Luong, M.-T.; So, D. R.; Hall, J.; Fiedel, N.; Thoppilan, R.; Yang, Z.; Kulshreshtha, A.; Nemade, G.; Lu, Y.; et al. 2020. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977* .
- Athreya, R. G.; Ngonga Ngomo, A.-C.; and Usbeck, R. 2018. Enhancing Community Interactions with Data-Driven Chatbots—The DBpedia Chatbot. In *Companion Proceedings of the The Web Conference 2018*, 143–146.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450* .
- Cao, Z.; Li, W.; Li, S.; and Wei, F. 2018. Retrieve, rerank and rewrite: Soft template based neural summarization. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 152–161.
- Chen, Y.-C.; and Bansal, M. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. *arXiv preprint arXiv:1805.11080* .
- Chollampatt, S.; and Ng, H. T. 2018. Neural quality estimation of grammatical error correction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2528–2539.
- Cui, Y.; Che, W.; Liu, T.; Qin, B.; Yang, Z.; Wang, S.; and Hu, G. 2019. Pre-Training with Whole Word Masking for Chinese BERT. *arXiv preprint arXiv:1906.08101* .
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* .
- Dong, L.; Yang, N.; Wang, W.; Wei, F.; Liu, X.; Wang, Y.; Gao, J.; Zhou, M.; and Hon, H.-W. 2019. Unified language model pre-training for natural language understanding and generation. In *Advances in Neural Information Processing Systems*, 13063–13075.
- Elgohary, A.; Peskov, D.; and Boyd-Graber, J. 2019. Can You Unpack That? Learning to Rewrite Questions-in-Context. 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)*, 5918–5924. Hong Kong, China: Association for Computational Linguistics. doi:10.18653/v1/D19-1605. URL <https://www.aclweb.org/anthology/D19-1605>.
- Ge, T.; Wei, F.; and Zhou, M. 2018. Fluency boost learning and inference for neural grammatical error correction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1055–1065.
- Gong, X.; Kong, X.; Zhang, Z.; Tan, L.; Zhang, Z.; and Shao, B. 2019. Customer Service Automatic Answering System Based on Natural Language Processing. In *Proceedings of the 2019 International Symposium on Signal Processing Systems*, 115–120.
- Hewitt, T.; and Beaver, I. 2020. A Case Study of User Communication Styles with Customer Service Agents versus Intelligent Virtual Agents. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 79–85.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8): 1735–1780.
- Huang, M.; Zhu, X.; and Gao, J. 2020. Challenges in building intelligent open-domain dialog systems. *ACM Transactions on Information Systems (TOIS)* 38(3): 1–32.
- Ji, Z.; Lu, Z.; and Li, H. 2014. An information retrieval approach to short text conversation. *arXiv preprint arXiv:1408.6988* .
- Lebanoff, L.; Muchovej, J.; Dernoncourt, F.; Kim, D. S.; Kim, S.; Chang, W.; and Liu, F. 2019. Analyzing sentence fusion in abstractive summarization. *arXiv preprint arXiv:1910.00203* .
- Li, P. 2020. An Empirical Investigation of Pre-Trained Transformer Language Models for Open-Domain Dialogue Generation. *arXiv preprint arXiv:2003.04195* .
- Lin, F.; Zhang, C.; Liu, S.; and Ma, H. 2020. A Hierarchical Structured Multi-Head Attention Network for Multi-Turn Response Generation. *IEEE Access* 8: 46802–46810.
- Malmi, E.; Krause, S.; Rothe, S.; Mirylenka, D.; and Severyn, A. 2019. Encode, tag, realize: High-precision text editing. *arXiv preprint arXiv:1909.01187* .
- Ng, H. T.; Wu, S. M.; Briscoe, T.; Hadiwinoto, C.; Susanto, R. H.; and Bryant, C. 2014. The CoNLL-2014 shared task on grammatical error correction. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning: Shared Task*, 1–14.
- Pan, Z.; Bai, K.; Wang, Y.; Zhou, L.; and Liu, X. 2019. Improving open-domain dialogue systems via multi-turn incomplete utterance restoration. 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)*, 1824–1833.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training. URL <https://openai.com/blog/language-unsupervised/>.
- Ritter, A.; Cherry, C.; and Dolan, W. B. 2011. Data-driven response generation in social media. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, 583–593.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368* .
- Sennrich, R.; Firat, O.; Cho, K.; Birch, A.; Haddow, B.; Hirschler, J.; Junczys-Dowmunt, M.; Läubli, S.;



- Miceli Barone, A. V.; Mokry, J.; and Nădejde, M. 2017. Nematus: a Toolkit for Neural Machine Translation. In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, 65–68. Valencia, Spain: Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/E17-3017>.
- Serban, I. V.; Klinger, T.; Tesauro, G.; Talamadupula, K.; Zhou, B.; Bengio, Y.; and Courville, A. 2017. Multiresolution recurrent neural networks: An application to dialogue response generation. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A.; and Pineau, J. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Su, H.; Shen, X.; Zhang, R.; Sun, F.; Hu, P.; Niu, C.; and Zhou, J. 2019. Improving Multi-turn Dialogue Modelling with Utterance ReWriter. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 22–31.
- Thadani, K.; and McKeown, K. 2013. Supervised sentence fusion with single-stage inference. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, 1410–1418.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wu, B.; Wang, B.; and Xue, H. 2016. Ranking responses oriented to conversational relevance in chat-bots. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, 652–662.
- Wu, Y.; Wu, W.; Xing, C.; Zhou, M.; and Li, Z. 2016. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. *arXiv preprint arXiv:1612.01627*.
- Wubben, S.; van den Bosch, A.; and Krahmer, E. 2012. Sentence Simplification by Monolingual Machine Translation. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1015–1024.
- Xing, C.; Wu, W.; Wu, Y.; Liu, J.; Huang, Y.; Zhou, M.; and Ma, W.-Y. 2016. Topic augmented neural response generation with a joint attention mechanism. *arXiv preprint arXiv:1606.08340* 2(2).
- Yan, R.; Song, Y.; and Wu, H. 2016. Learning to respond with deep neural networks for retrieval-based human-computer conversation system. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, 55–64.
- Zhang, X.; and Lapata, M. 2017. Sentence Simplification with Deep Reinforcement Learning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 584–594.
- Zhang, Z.; Li, J.; Zhu, P.; Zhao, H.; and Liu, G. 2018. Modeling multi-turn conversation with deep utterance aggregation. *arXiv preprint arXiv:1806.09102*.
- Zhao, W.; Wang, L.; Shen, K.; Jia, R.; and Liu, J. 2019. Improving grammatical error correction via pre-training a copy-augmented architecture with unlabeled data. *arXiv preprint arXiv:1903.00138*.
- Zhao, Y.; Xu, C.; and Wu, W. 2020. Learning a Simple and Effective Model for Multi-turn Response Generation with Auxiliary Tasks. *arXiv preprint arXiv:2004.01972*.
- Zhou, X.; Dong, D.; Wu, H.; Zhao, S.; Yu, D.; Tian, H.; Liu, X.; and Yan, R. 2016. Multi-view response selection for human-computer conversation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 372–381.
- Zhou, X.; Li, L.; Dong, D.; Liu, Y.; Chen, Y.; Zhao, W. X.; Yu, D.; and Wu, H. 2018. Multi-turn response selection for chatbots with deep attention matching network. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1118–1127.