

Incomplete Utterance Rewriting as Semantic Segmentation

Qian Liu^{†*}, Bei Chen[§], Jian-Guang Lou[§], Bin Zhou[†], Dongmei Zhang[§]

[†]School of Computer Science and Engineering, Beihang University, China

[§]Microsoft Research, Beijing, China

[†]{qian.liu, zhoubin}@buaa.edu.cn; [§]{beichen, jlou, dongmeiz}@microsoft.com

Abstract

Recent years the task of incomplete utterance rewriting has raised a large attention. Previous works usually shape it as a machine translation task and employ sequence to sequence based architecture with copy mechanism. In this paper, we present a novel and extensive approach, which formulates it as a semantic segmentation task. Instead of generating from scratch, such a formulation introduces edit operations and shapes the problem as prediction of a word-level edit matrix. Benefiting from being able to capture both local and global information, our approach achieves state-of-the-art performance on several public datasets. Furthermore, our approach is four times faster than the standard approach in inference.

1 Introduction

A dramatic progress has been achieved in single-turn dialogue modeling such as open-domain response generation (Shang et al., 2015), question answering (Rajpurkar et al., 2016), etc. By contrast, multi-turn dialogue modeling is still in its infancy, as users tend to use incomplete utterances which usually omit or refer back to entities or concepts appeared in the dialogue context, namely ellipsis and coreference. According to previous studies, ellipsis and coreference exist in more than 70% of the utterances (Su et al., 2019), for which a dialogue system must be equipped with the ability of understanding them. To tackle the problem, early works include learning a hierarchical representation (Serban et al., 2017; Zhang et al., 2018) and concatenating the dialogue utterances selectively (Yan et al., 2016). Recently, researchers focus on a more explicit and explainable solution: the task of **Incomplete Utterance Rewriting (IUR)**, also known as context rewriting (Kumar and Joshi, 2016; Su et al., 2019; Liu et al.,

Turn	Utterance (Translation)
x_1 (A)	北京今天天气如何 <i>How is the weather in Beijing today</i>
x_2 (B)	北京今天是阴天 <i>Beijing is cloudy today</i>
x_3 (A)	为什么总是这样 <i>Why is always this</i>
x_3^*	北京为什么总是阴天 <i>Why is Beijing always cloudy</i>

Table 1: An example dialogue between user A and B, including the context utterances (x_1 , x_2), the incomplete utterance (x_3) and the rewritten utterance (x_3^*).

2019a; Pan et al., 2019; Elgohary et al., 2019; Zhou et al., 2019). IUR aims to rewrite an incomplete utterance into an utterance which is semantically equivalent but self-contained to be understood without context. As shown in Table 1, the incomplete utterance x_3 not only omits the subject “北京”(Beijing), but also refers to the semantic of “阴天”(cloudy) via “这样”(this). By explicitly recovering the hidden semantics behind x_3 into x_3^* , IUR makes the downstream dialogue modeling more precise.

To deal with IUR, a natural idea is to transfer models from coreference resolution (Clark and Manning, 2016). However, this idea is not easy to realize, as ellipsis also accounts for a large proportion. Despite being different, coreference and ellipsis both can be resolved without introducing out-of-dialogue words in most cases. That is to say, words of the rewritten utterance are nearly from either the context utterances or the incomplete utterance. Observing it, most previous works employ the pointer network (Vinyals et al., 2015) or the sequence to sequence model with copy mechanism (Gu et al., 2016; See et al., 2017). However, they generate the rewritten utterance

*Work done during an internship at Microsoft Research.

from scratch, neglecting a key trait that the main structure of a rewritten utterance is always the same as the incomplete utterance. To highlight it, we imagine the rewritten utterance as the outcome after a series of edit operations (i.e. substitute and insert) on the incomplete utterance. Taking the example from Table 1, \mathbf{x}_3^* can be obtained by substituting “这样”(this) in \mathbf{x}_3 with “阴天”(cloudy) in \mathbf{x}_2 and inserting “北京”(Beijing) before “为什么”(Why), much easier than producing \mathbf{x}_3^* via decoding word by word. These edit operations are carried out between word pairs of the context utterances and the incomplete utterance, analogous to semantic segmentation (a well-known task in computer vision): Given relevance features between word pairs as an image, the model is to predict the edit type for each word pair as a pixel-level mask (elaborated in Section 3). Inspired by the above, in this paper, we propose a novel and extensive approach which formulates IUR as semantic segmentation¹. Our contributions are as follows:

- As far as we know, we are the first to present such a highly extensive approach which formulates the incomplete utterance rewriting as a semantic segmentation task.
- Benefiting from being able to capture both local and global information, our approach achieves state-of-the-art performance on several datasets across different domains and languages.
- Furthermore, our model predicts the edit operations in parallel, and thus obtains a much faster inference speed than traditional methods.

2 Related Work

The most related work to ours is the line of incomplete utterance rewriting. Recently, it has raised a large attention in several domains. In question answering, previous works include non-sentential utterance resolution using the sequence to sequence based architecture (Kumar and Joshi, 2016), incomplete follow-up question resolution via a retrieval sequence to sequence model (Kumar and Joshi, 2017) and sequence to sequence model with a copy mechanism (Elgohary et al., 2019; Quan et al., 2019). In conversational semantic parsing, Liu et al. (2019b) proposed a novel approach which considers the structures of questions, while

¹Our code is available at <https://github.com/microsoft/ContextualSP>.

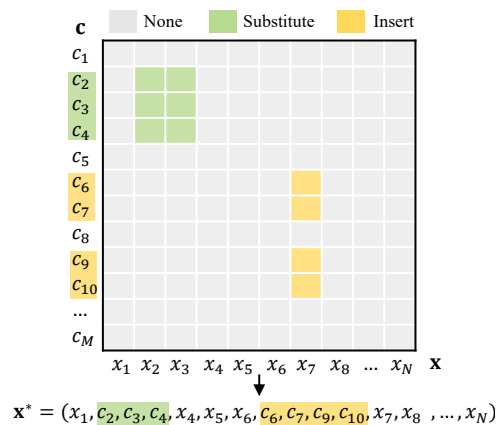


Figure 1: The illustration of the word-level edit matrix applied in our formulation. Each cell belongs to one of three edit types: None, Substitute and Insert.

Liu et al. (2019a) imposed an intermediate structure span and decomposed the incomplete utterance rewriting into two sub-tasks. In dialogue generation, Pan et al. (2019) presented a cascaded model which first picks words from the context via BERT, and then combines these words to generate the rewritten utterance, and Su et al. (2019) distinguished the weights of context utterances and the incomplete utterance using a hyper-parameter λ . Different from all of them, we formulate the task as a semantic segmentation task.

Our work is also closely related to coreference resolution. It is an active task that has been studied years, and deep learning based methods have achieved state-of-the-art performance via the paradigm of scoring span or mention pairs (Clark and Manning, 2015, 2016; Lee et al., 2017, 2018). Researchers also explored to use unsupervised contextualized representations to enhance the coreference resolution. Joshi et al. (2019) applied SpanBERT (Joshi et al., 2020) to enhance the span representation in coreference resolution, and Wu et al. (2020) formulated coreference resolution as query-based span prediction and employed SpanBERT to solve it as a machine reading task. The above works only focus on coreference resolution, while our work deals with coreference and ellipsis under a unified approach.

From the perspective of the methodology, our work is correlated with directions of edit-based text generation and semantic similarity measurement. Wu et al. (2019) proposed a prototype-then-edit paradigm for open-domain response generation, while Malmi et al. (2019) cast text generation as a text editing task and tackled it with a

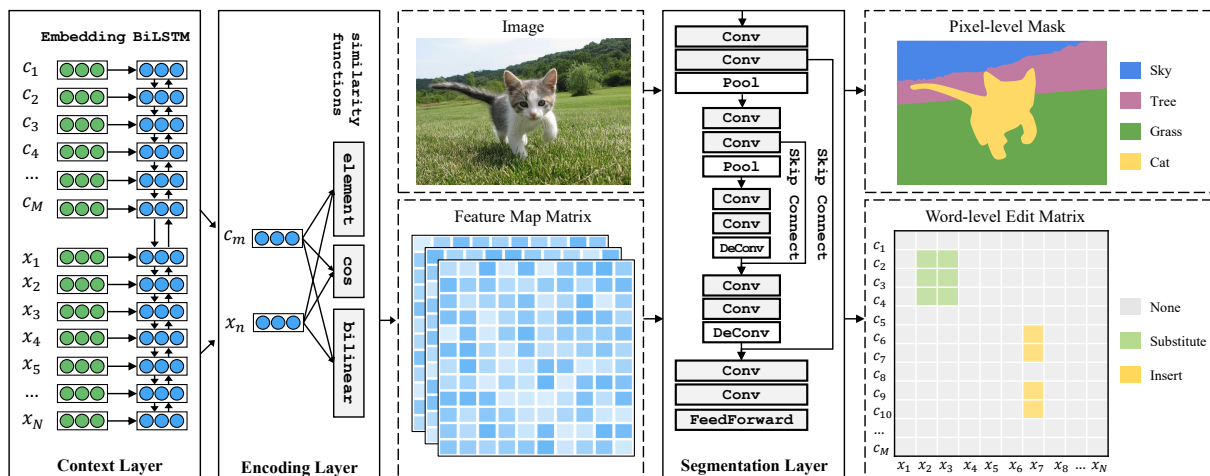


Figure 2: The illustration of the word-level edit matrix construction. The dashed boxes represent the intermediate results of our proposed model (**bottom**) and their counterparts in semantic segmentation (**above**). Inside the segmentation layer, a “Conv” module consists of a convolutional neural network, batch normalization and an activation function ReLU. “Pool” and “DeConv” are short for max pooling and deconvolution neural network respectively.

sequence tagging approach. Our work is different from theirs since we model the editing process between two sentences as a semantic segmentation task. As for semantic similarity measurement, similar to us, both He and Lin (2016) and Pang et al. (2016) used convolutional neural networks to capture similarities between sentences.

3 Incomplete Utterance Rewriting as Semantic Segmentation

In this section, we will have a glance at the fundamental idea behind our approach: incomplete utterance rewriting as semantic segmentation.

In a multi-turn dialogue, given the context utterances (x_1, \dots, x_{t-1}) and the incomplete utterance x_t , IUR is to rewrite x_t to x_t^* using contextual information, where x_t^* has the same meaning with x_t . The rewritten utterance x_t^* has self-contained semantics and can be understood solely. To produce x_t^* , our approach formulates the problem as a semantic segmentation task. Concretely, we concatenate all the context utterances to produce an M -length word sequence $\mathbf{c} = (c_1, c_2, \dots, c_M)$. To separate context utterances in different turns, we insert a special word [S] between each context utterance. Meanwhile, the incomplete utterance is denoted by $\mathbf{x} = (x_1, x_2, \dots, x_N)$. As mentioned, the rewritten utterance \mathbf{x}^* can be obtained by editing the incomplete utterance \mathbf{x} with in-dialogue words (i.e. words in \mathbf{c}). To model edit operations between \mathbf{x} and \mathbf{c} , we define a $M \times N$ matrix Y ,

where the entry $Y_{m,n}$ indicates the edit type between c_m and x_n . There are three kinds of edit types: *Substitute* means replacing the span in \mathbf{x} with the corresponding context span in \mathbf{c} ; *Insert* aims to insert the context span before a certain token in \mathbf{x} , and *None* represents no operation. For example, as shown schematically in Figure 1, we can edit \mathbf{x} by replacing (x_2, x_3) with (c_2, c_3, c_4) and insert (c_6, c_7, c_9, c_{10}) before x_7 . It is notable that we append a special word [E] to \mathbf{x} , to enable *Insert* take place after \mathbf{x} . More concrete examples can be found in Section 5.3.

Then, we propose to emit such a matrix Y in a way analogous to the task of semantic segmentation. Specially, we build a $M \times N$ feature map via capturing the word-to-word relevance between \mathbf{c} and \mathbf{x} . Taking the feature map as an image, the output word-level edit matrix Y is parallel to the pixel-level mask in semantic segmentation, which bridges IUR and semantic segmentation. Such a formulation comes with several key advantages: (i) **Easy**: compared with traditional methods generating the rewritten utterance from scratch, such a formulation introduces edit operations to lower the difficulty of generation; (ii) **Fast**: these edits are predicted concurrently, so our model naturally enjoys the fast inference speed than conventional models which decode word by word; (iii) **Transferable**: taking the formulation as a bridge between IUR and semantic segmentation, one can transfer empirical models from the community of semantic segmentation with ease.

4 Methodology

As shown in Figure 2, our approach firstly obtains the word-level edit matrix through three neural layers. Then based on the word-level edit matrix, it applies a generation algorithm to produce the rewritten utterance. Since the model yields a U-shaped architecture (illustrated later), we name our approach as **Rewritten U-shaped Network (RUN)**.

4.1 Word-level Edit Matrix Construction

To construct a word-level edit matrix, our model passes through three neural layers: a context layer, an encoding layer and a subsequent segmentation layer. The context layer produces a context-aware representation for each word in both \mathbf{c} and \mathbf{x} , based on which the encoding layer forms a feature map matrix \mathbf{F} to capture word-to-word relevance. Finally a segmentation layer is applied to emit the word-level edit matrix.

Context Layer As shown in the left of Figure 2, at first the concatenation of \mathbf{c} and \mathbf{x} passes by the word embedding ϕ to get the representation for each word in both utterances. The embedding is initialized using GloVe (Pennington et al., 2014), and then updated along with other parameters. On top of the joint word embedding sequence $(\phi(c_1), \dots, \phi(c_M), \phi(x_1), \dots, \phi(x_N))$, Bidirectional Long Short-Term Memory Network (BiLSTM) (Hochreiter and Schmidhuber, 1997; Schuster and Paliwal, 1997) is applied to capture contextual information inter and intra utterances. Although \mathbf{c} and \mathbf{x} are jointly encoded by BiLSTM (see the left of Figure 2), below we distinguish their hidden states for clear illustration. For a word $c_m (m = 1, \dots, M)$ in \mathbf{c} , its hidden state is denoted by \mathbf{u}_m obtained through BiLSTM, while the hidden state \mathbf{h}_n is for word $x_n (n = 1, \dots, N)$ in the incomplete utterance.

Encoding Layer On top of the context-aware hidden states, we consider several similarity functions to encode the word-to-word relevance. Concretely, for each word x_n in the incomplete utterance and c_m in the context utterances, their relevance is captured by a D -dimensional feature vector $\mathbf{F}(x_n, c_m)$. It is produced by concatenating element-wise similarity (Ele Sim.), cosine similarity (Cos Sim.) and learned bi-linear similarity (Bi-Linear Sim.) between them as:

$$\mathbf{F}(x_n, c_m) = [\mathbf{h}_n \odot \mathbf{u}_m; \cos(\mathbf{h}_n, \mathbf{u}_m); \mathbf{h}_n \mathbf{W} \mathbf{u}_m], \quad (1)$$

where \mathbf{W} is a learned parameter. These similarity functions are expected to model the word-to-word relevance from different perspectives, important for the follow-up edit type classification. However, they concentrate on local rather than global information (see discussion in Section 5.3). To capture global information, a segmentation layer is proposed.

Segmentation Layer Taking the feature map matrix $\mathbf{F} \in \mathbb{R}^{M \times N \times D}$ as a D -channel image, the segmentation layer is to predict the word-level edit matrix $Y \in \mathbb{R}^{M \times N}$, analogous to a pixel-level mask. Inspired by UNet (Ronneberger et al., 2015), the layer is formed as a U-shaped structure: two down-sampling blocks and two up-sampling blocks with skip connection. A down-sampling block contains two separate ‘‘Conv’’ modules and a subsequent max pooling. Each down-sampling block doubles the number of channels. Intuitively, the down-sampling block expands the receptive fields of each cell, hence providing rich global information for the final decision. An up-sampling block contains two separate ‘‘Conv’’ modules, and a subsequent deconvolution neural network. Each up-sampling block halves the number of channels and concatenates the correspondingly cropped feature map in down-sampling as the output (skip connect in Figure 2). Finally a feedforward neural network is employed to map each feature vector to one of three edit types, obtaining the word-level edit matrix Y . By incorporating an encoding layer and a segmentation layer, our model is able to capture both local and global information.

BERT Enhanced Embedding Since pretrained language models have been proven to be effective on several tasks, we also experiment with employing BERT (Devlin et al., 2019) to augment our model via BERT enhanced embedding.

4.2 Rewritten Utterance Generation

Once a word-level edit matrix is emitted, a subsequent generation algorithm is applied for producing the rewritten utterance. As indicated in Figure 1, to apply edit operations without ambiguity, we assume each edit region in Y is a rectangle. However, the predicted Y is not guaranteed to meet this requirement, indicating the need for a standardization step. Therefore, the overall procedure of generation is divided into two stages: first the algorithm delimits standard edit regions via searching minimal covering rectangles

for each connected region; then it manipulates the incomplete utterance based on these standard edit regions to produce the rewritten utterance. Since the second step has been illustrated in Section 3, in the following we concentrate on the first standardization step.

In the standardization step, we employ the two-pass algorithm (also known as Hoshen–Kopelman algorithm) to find connected regions (Hoshen and Kopelman, 1976). In a nutshell, the algorithm makes two passes over the word-level edit matrix. The first pass is to assign temporary cluster labels and record equivalences between clusters in an order of left to right and top to down. Concretely, for each cell, if its neighbors (i.e. left or top cells with the same edit type) have been assigned temporary cluster labels, it is labeled as the smallest neighboring label. Meanwhile, its neighboring clusters are recorded as equivalent. Otherwise, a new temporary cluster label is created for the cell. The second pass is to merge temporary cluster labels which are recorded as equivalent. Finally, cells with the same label form a connected region. For each connected region, we use its minimal covering rectangle to serve as the output of our model.

4.3 Distant Supervision

As mentioned in Section 3, the expected supervision for our model is the word-level edit matrix, but existing datasets only contain rewritten utterances. Therefore, we use a procedure to automatically derive (noisy) word-level edit matrices (i.e. distant supervision), and use these examples to train our model. We use the following process to build our training set. First, we find a Longest Common Subsequence (LCS) between \mathbf{x} and \mathbf{x}^* . Then, for each word in \mathbf{x}^* , if it is not in LCS, it is marked as ADD. Conversely, for each word in \mathbf{x} but not in LCS, it is marked as DEL. Contiguous words with the same mark are merged into one span. By a span-level comparison, any ADD span in \mathbf{x}^* with a DEL counterpart (i.e. under the same context) relates it to `Substitute`. Otherwise, the span is inserted into \mathbf{x} , corresponding to `Insert`.

Taking the example from Table 1, given \mathbf{x} as “为什么总是这样”(Why is always this) and \mathbf{x}^* as “北京为什么总是阴天”(Why is Beijing always cloudy), their longest common subsequence is “为什么总是”(Why is always). Therefore, with “这样”(this) in \mathbf{x} being marked as DEL and “阴天”(cloudy) in \mathbf{x}^* being marked as ADD, they cor-

respond to the edit type `Substitute`. In comparison, since “北京”(Beijing) cannot find a counterpart, it is related to the edit type `Insert`.

5 Experiments

In this section, we conduct thorough experiments to demonstrate the superiority of our approach.

5.1 Experimental Setup

Datasets We conduct experiments on four public datasets across different domains: Open-Domain Dialogue (MULTI Pan et al., 2019, REWRITE Su et al., 2019), Task-Oriented Dialogue (TASK Quan et al., 2019) and Question Answering in Context (CANARD Elgohary et al., 2019). We use the same data split for these datasets as their original paper, and some statistics are shown in Table 3.

Baselines We consider a bunch of baselines, including LSTM-based models, Transformer-based models and state-of-the-art models on each dataset. (i) **LSTM-based models** consist of the vanilla sequence to sequence model with attention (L-Gen) (Bahdanau et al., 2015), the pointer network architecture (L-Ptr) (Vinyals et al., 2015) and the hybrid pointer generator network (L-Ptr-Gen) (See et al., 2017). (ii) **Transformer-based models** consist of the basic transformer model (T-Gen) (Vaswani et al., 2017), the transformer-based pointer network (T-Ptr), and the transformer-based pointer generator (T-Ptr-Gen). (iii) **State-of-the-art models** consist of Syntactic (Kumar and Joshi, 2016), PAC (Pan et al., 2019), GECOR (Quan et al., 2019), L-Ptr- λ and T-Ptr- λ (Su et al., 2019). We refer readers to their papers for more details. It is remarkable that above methods all generate rewritten utterances from scratch.

Evaluation We employ both automatic metrics and human evaluations to evaluate our approach. As in literature (Pan et al., 2019), we examine RUN using the widely used automatic metrics BLEU, ROUGE, EM and Rewriting F-score. (i) **BLEU_n** (\mathbf{B}_n) evaluates how similar the rewritten utterances are to the golden ones via the cumulative n -gram BLEU score (Papineni et al., 2002). (ii) **ROUGE_n** (\mathbf{R}_n) measures the n -gram overlapping between the rewritten utterances and the golden ones, while **ROUGE_L** (\mathbf{R}_L) measures the longest matching sequence between them (Lin, 2004). (iii) **EM** stands for the exact match ac-

Model	\mathcal{P}_1	\mathcal{R}_1	\mathcal{F}_1	\mathcal{P}_2	\mathcal{R}_2	\mathcal{F}_2	\mathcal{P}_3	\mathcal{R}_3	\mathcal{F}_3	\mathbf{B}_1	\mathbf{B}_2	\mathbf{R}_1	\mathbf{R}_2
Syntactic †	67.4	37.2	47.9	53.9	30.3	38.8	45.3	25.3	32.5	84.1	81.2	89.3	80.6
L-Gen †	65.5	40.8	50.3	52.2	32.6	40.1	43.6	27.0	33.4	84.9	81.7	88.8	80.3
L-Ptr-Gen †	66.6	40.4	50.3	<u>54.0</u>	33.1	41.1	<u>45.9</u>	28.1	34.9	84.7	81.7	89.0	80.9
RUN (Ours)	66.9	54.9	60.3	<u>53.0</u>	43.4	47.7	43.8	35.7	39.3	91.1	88.0	91.0	83.3
PAC †	70.5	58.1	63.7	55.4	45.1	49.7	45.2	36.6	40.4	89.9	86.3	91.6	82.8
RUN + BERT (Ours)	73.2	64.6	68.6	59.5	53.0	56.0	50.7	45.1	47.7	92.3	89.6	92.4	85.1

Table 2: The experimental results of (Top) general and (Bottom) BERT-based results on MULTI. †: Results from Pan et al. (2019). A bolded number in a column indicates a statistically significant improvement against all the baselines ($p < 0.05$), whereas underline numbers show comparable performances. Both are same for Table 4&5.

	MULTI	REWRITE	TASK	CANARD
Language	Chinese	Chinese	English	English
# Ques. (Train)	194 K	18 K	2.2 K	32 K
# Ques. (Dev)	5 K	2 K	0.5 K	4 K
# Ques. (Test)	5 K	NA	NA	6 K
Avg. Con len	25.8	17.7	52.6	85.4
Avg. Cur len	8.6	6.5	9.4	7.5
Avg. Rew len	12.4	10.5	11.3	11.6

Table 3: Statistics of different datasets. NA means the development set is also the test set. “Ques” is short for questions, “Avg” for average, “len” for length, “Con” for context utterance, “Cur” for current utterance, and “Rew” for rewritten utterance.

curacy, which is the strictest evaluation metric. (iv) **Rewriting Precision $_n$** , **Recall $_n$** and **F-score $_n$** ($\mathcal{P}_n, \mathcal{R}_n, \mathcal{F}_n$) emphasize more on words from \mathbf{c} which are argued to be harder to copy (Pan et al., 2019). Therefore, they are calculated on the collection of n -grams that contain at least one word from \mathbf{c} . As validated by Pan et al. (2019), above automatic metrics are credible indicators to reflect the rewrite quality. However, none of automatic metrics reflects the utterance fluency or the improvement on downstream tasks. Therefore, human evaluations are included to evaluate the fluency of rewritten utterances and their boost on the downstream task.

Implementation Details Our implementation was based on PyTorch (Paszke et al., 2019), AllenNLP (Gardner et al., 2018) and HuggingFace’s transformers library (Wolf et al., 2019). Since the distribution of edit types is severely unbalanced (e.g. None accounts for nearly 90%), we employed weighted cross-entropy loss and tuned the weight on development sets. We used Adam (Kingma and Ba, 2015) to optimize our model and set the learning rate as 1e-3, except for BERT as 1e-5. The embedding size and hidden size in BiLSTM are 100 and 200 respectively. Specifically, BERT

Model	EM	\mathbf{B}_2	\mathbf{B}_4	\mathbf{R}_2	\mathbf{R}_L
L-Gen	47.3	81.2	73.6	80.9	86.3
L-Ptr-Gen	50.5	82.9	75.4	83.8	87.8
L-Ptr-Net	51.5	82.7	75.5	84.0	88.2
L-Ptr- λ †	42.3	82.9	73.8	81.1	84.1
T-Gen	35.4	72.7	62.5	74.5	82.9
T-Ptr-Gen	53.1	84.4	77.6	<u>85.0</u>	89.1
T-Ptr-Net	53.0	83.9	77.1	<u>85.1</u>	88.7
T-Ptr- λ	52.6	85.6	78.1	<u>85.0</u>	89.0
RUN (Ours)	53.8	86.1	79.4	<u>85.1</u>	89.5
T-Ptr- λ + BERT	57.5	86.5	79.9	86.9	90.5
RUN + BERT (Ours)	66.4	91.4	86.2	90.4	93.5

Table 4: The experimental results on REWRITE. †: Reproduced from the code released by Su et al. (2019).

mentioned above refer to BERT_{base}. All results of baselines without specific marks were reproduced by ours using OpenNMT with beam size as 4 (Klein et al., 2017).

Connection Words Similar to pointer network (Vinyals et al., 2015), RUN is restricted to predict words which have appeared in the dialogue. Although most examples work well under the restriction, there still exist a few cases which rely on certain words to generate fluent utterances. For example, when rewriting possessive pronouns such as “their”, we usually need an extra word “of” to enhance the fluency. Such common words, named after *connection words*, improve fluency of the rewritten utterances. In practice, we append a small list of connection words to the tail of \mathbf{c} , enabling our model to pick connection words as well. For each dataset, their connection word list is automatically derived from the training data.

5.2 Model Comparison

Table 2 and Table 4 show experimental results of our approach and baselines on MULTI and REWRITE. As shown, our approach outperforms all baselines significantly. Taking MULTI as an

TASK				CANARD						
Model	EM	B ₄	F ₁	Model	B ₁	B ₂	B ₄	R ₁	R ₂	R _L
Ellipsis Recovery [†]	50.4	74.1	44.1	Copy	52.4	46.7	37.8	72.7	54.9	68.5
GECOR 1 [†]	<u>68.5</u>	83.9	66.1	Pronoun Sub	60.4	55.3	47.4	73.1	63.7	73.9
GECOR 2 [†]	66.2	83.0	66.2	L-Ptr-Gen	67.2	<u>60.3</u>	<u>50.2</u>	<u>78.9</u>	62.9	<u>74.9</u>
RUN (Ours)	<u>69.2</u>	85.6	70.6	RUN (Ours)	70.5	<u>61.2</u>	<u>49.1</u>	<u>79.1</u>	61.2	<u>74.7</u>

Table 5: The experimental results on (Left) TASK and (Right) CANARD. †: Results from Quan et al. (2019).

	Win	Tie	Loss
RUN v.s. L-Ptr-λ	41.6 %	42.4 %	16.0 %
RUN v.s. T-Ptr-Gen	23.6 %	56.4 %	20.0 %
RUN v.s. T-Ptr-λ	22.6 %	57.0 %	20.4 %

Table 6: Pairwise human evaluation results about the rewritten utterance fluency on randomly sampled 500 dialogues from REWRITE. Our approach achieves similar or better fluency compared with top baselines.

	Origin	L-Gen	L-Ptr-Gen	RUN	Gold
Avg. Score	0.92	0.93	0.91	1.09	1.10
NR	100%	74%	68%	51%	46%

Table 7: Human rating evaluations about the response quality on sampled 300 dialogues from the development set of MULTI. The score ranges from 0 to 2. “NR” represents the proportion of rewritten utterances which are equal to current utterances.

example, our approach exceeds the best baseline L-Ptr-Gen by a large margin, reaching a new state-of-the-art performance on almost all automatic metrics. To illustrate, our approach improves the previous best model by 6.4 points and 10.0 points on B₁ and F₁ respectively. Furthermore, our approach leaves a striking impression when augmented with BERT. It not only fully surpasses the best sequence generation baseline with BERT (i.e. T-Ptr-λ+BERT on REWRITE), but also obtains a considerable boost over a cascade model designed for stimulating potential of BERT (i.e. PAC on MULTI). Even for the most challenging metric EM on REWRITE, RUN with BERT improves 8.9 points, demonstrating the superiority of our model. Besides, our approach also achieves comparable or better results against all baselines on TASK and CANARD, as shown in Table 5.

Besides automatic results, we perform two groups of human evaluation to answer (i) how fluent the rewritten utterances are and (ii) how much IUR can contribute to downstream tasks. For the evaluation of fluency, we randomly sampled 500

Beam	Model	B ₄	ΔB ₄	Latency	Speedup
4	L-Gen	73.6	0.0	82 ms	1.00 ×
	L-Ptr-Net	75.5	+1.9	116 ms	0.71 ×
	L-Ptr-Gen	75.4	+1.8	110 ms	0.75 ×
	T-Gen	62.5	-11.1	322 ms	0.25 ×
	T-Ptr-Net	77.1	+3.5	576 ms	0.14 ×
	T-Ptr-Gen	77.6	+4.0	415 ms	0.20 ×
1	L-Gen	73.5	-0.1	55 ms	1.49 ×
	L-Ptr-Net	76.2	+3.0	95 ms	0.86 ×
	L-Ptr-Gen	73.3	-0.3	59 ms	1.39 ×
	T-Gen	60.9	-12.7	240 ms	0.38 ×
	T-Ptr-Net	77.9	+4.3	401 ms	0.20 ×
	T-Ptr-Gen	77.1	+3.5	374 ms	0.22 ×
-	RUN (Ours)	79.4	+5.8	21 ms	3.90 ×

Table 8: The inference speed comparison between RUN and baselines. Beam stands for the beam size in beam search, not applicable for RUN. Latency is computed as the time to produce a single sentence without data batching, averaged over the development set of REWRITE. All models are implemented in PyTorch on a single NVIDIA V100.

dialogues in the development set of REWRITE. Then we fed them to representative IUR models and presented generated rewritten utterances to 10 judges, who are asked to decide which of the rewritten utterances is of higher fluency in pairwise comparisons. Ties are acceptable. Table 6 shows the evaluation results. In comparison to the best baseline T-Ptr-λ, our model only loses in 20.4% cases, which is extremely competitive.

To access the influence of IUR on downstream tasks, we choose multi-turn response selection as a representative, which aims to retrieve suitable responses from a candidate pool considering the context. Concretely, an SMN model trained on the Douban Conversation Corpus is selected as the backbone in multi-turn response selection (Wu et al., 2017). At first we sampled 300 dialogues from the development set of MULTI as the input to IUR models. Then their predicted rewritten utterances and the context utterances were fed into the SMN model, to help it select suitable responses.

Variant	\mathcal{F}_1	\mathcal{F}_2	\mathcal{F}_3	\mathbf{B}_2	\mathbf{R}_2
RUN	60.3	47.8	39.4	87.9	83.2
w/o Edit	0.0	0.0	0.0	77.4	75.6
w/o U-shape Seg.	55.2	41.4	33.1	86.1	82.5
w/o Ele Sim.	60.4	47.1	38.3	86.8	82.6
w/o Cos Sim.	62.3	48.4	39.2	85.3	82.3
w/o Bi-Linear Sim.	61.6	48.0	39.0	85.8	82.6

Table 9: The ablation results on the development set of MULTI. “w/o Edit” means directly using the current utterance as the rewritten utterance. “w/o U-shape seg.” means that our segmentation layer is replaced by a feed-forward neural network with comparable parameters. The remaining variants ablate different similarity functions in the encoding layer.

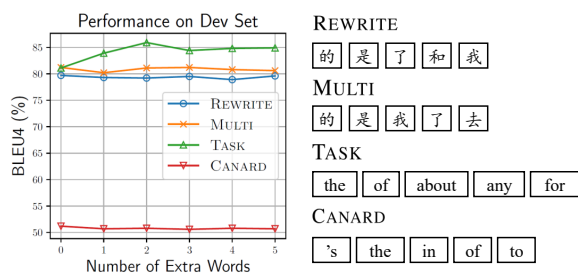


Figure 3: (Left) BLEU4 (\mathbf{B}_4) performance with different number of connection words on the development sets of different datasets. (Right) Connection words in decreasing order of frequency on each dataset.

The response candidate pool was formed by all utterances in MULTI. Finally, 5 workers were asked to evaluate responses following a multi-scale rating from 0 to 2: 0 means the response is not related to the dialogue; 1 means the response is related but not interesting enough; and 2 means the response is satisfying. To illustrate more clearly, we also conduct human rating evaluation on responses under the settings of original dialogue (i.e. without rewriting, relying on the SMN model itself to understand the context) and gold dialogue (i.e. human rewriting). As shown in Table 7, our model achieves the highest response quality score among IUR models, improving the original setting by 19% relatively. Considering that the SMN model is capable of aggregating implicit context information, it is non-trivial for our model to further improve the response quality.

5.3 Closer Analysis

We conduct a series of experiments to analyze our model deeply. First we conduct an inference speed comparison between our model and representative

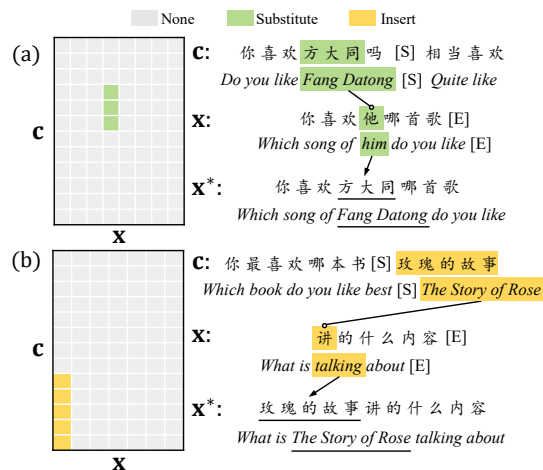


Figure 4: The illustration of (Left) the word-level edit matrix and (Right) the rewritten utterance generation process of two real cases (a) and (b) from REWRITE.

baselines under the same run-time environment. Then we verify the effectiveness of components in our model by a thorough ablation study. Meanwhile, we touch how the amount of connection words affect the performance. Finally, we present two real cases to illustrate our model concretely.

Inference Speed Table 8 compares inference speed between our model and baselines. Since L-Ptr- λ and T-Ptr- λ are not implemented under PyTorch, we do not show their inference time for fair consideration. Noticing the beam size would affect the inference time of baselines, we also show the results with beam size as 1. Using the simplest L-Gen as a standard, one can find that our model is nearly four times faster, with the highest improvement $\Delta \mathbf{B}_4$. Meanwhile, our model is the only one which can improve both performance and inference speed, significantly surpassing all baselines.

Ablation Study To verify the effectiveness of different components in our model, we present a thorough ablation study in Table 9. As expected, “w/o Edit” causes a huge drop on all evaluation metrics. Notably, the extreme drop on \mathcal{F}_n indicates that it is more suited for IUR than common metrics. “w/o U-shape Seg.”, which ablates the segmentation layer, also brings a great performance drop. Without our segmentation layer capturing global information, an encoding layer only achieves comparable performance with L-Gen, suggesting there are considerable benefits with bridging IUR and semantic segmentation. We also ablate different feature similarity functions (i.e. “w/o Ele Sim.”, “w/o Cos Sim.” and “w/o Bi-

Linear Sim.”) for an in-depth analysis. As shown in Table 9, ablating each similarity function will hurt most metrics. Meanwhile, our model does not depend on any similarity function severely, showing its robustness. Furthermore, we explore how the amount of connection words affect the performance in Figure 3. As indicated, except TASK, the number of connection words affect slightly. Nevertheless, it shows a positive effect overall, providing a way to generate out-of-dialogue words. We present two real cases in Figure 4 from REWRITE to illustrate the rewritten process of our model concretely. For both (a) coreference and (b) ellipsis, our model deals with them flexibly.

5.4 Discussion

While our approach has made some progress, it still has several limitations. First, our model severely relies on the word order implied by the dialogue. It makes our model vulnerable to some complex cases (i.e. multiple `Insert` corresponds to one position). The second limitation is that we predict edit types of each cell independently, ignoring the relationship between neighboring edit types. It is hopefully resolved by the conditional random field algorithm (Arnab et al., 2018).

The above limitations may raise concerns about the performance upper bound of our approach. In fact, it is not an issue. On three out of four datasets used in the experiments, more than 85% examples could be tackled perfectly by our approach (87.6% in TASK, 91.0% in REWRITE, 95.3% in MULTI). The number in CANARD is relatively low (42.5%) since human annotators introduce many new words in rewriting. Nevertheless, the BLEU upper bound in CANARD could be as high as 72.5% with our approach, which is acceptable.

The last point we focus on is why similarities can be good features for determining edits. We think it can be elaborated from two aspects. For coreference, the similarity function is suitable for identifying whether two spans refer to the same entity. For ellipsis, the similarity function is an effective indicator to find matching anchors, which indicate the possible insertion positions.

6 Conclusion & Future Work

In this paper, we present a novel and extensive approach which formulates the incomplete utterance rewriting as a semantic segmentation task. On top of the formulation, we carefully design a U-

shaped rewritten network, which outperforms existing baselines significantly on several datasets. In the future, we will investigate on extending our approach to more areas.

Acknowledgments

We thank all the anonymous reviewers for their valuable comments. This work was supported in part by National Natural Science Foundation of China (U1736217 and 61932003), and National Key R&D Program of China (2019YFF0302902).

References

- Anurag Arnab, Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Måns Larsson, Alexander Kirillov, Bogdan Savchynskyy, Carsten Rother, Fredrik Kahl, and Philip H. S. Torr. 2018. Conditional random fields meet deep neural networks for semantic segmentation: Combining probabilistic graphical models with deep learning for structured prediction. *IEEE Signal Process. Mag.*, 35(1):37–52.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kevin Clark and Christopher D. Manning. 2015. Entity-centric coreference resolution with model stacking. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1405–1415, Beijing, China. Association for Computational Linguistics.
- Kevin Clark and Christopher D. Manning. 2016. Improving coreference resolution by learning entity-level distributed representations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 643–653, Berlin, Germany. Association for Computational Linguistics.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ahmed Elgohary, Denis Peskov, and Jordan Boyd-Graber. 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), pages 5918–5924, Hong Kong, China. Association for Computational Linguistics.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Hua He and Jimmy Lin. 2016. [Pairwise word interaction modeling with deep neural networks for semantic similarity measurement](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 937–948, San Diego, California. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural computation*, 9(8):1735–1780.
- Joseph Hoshen and Raoul Kopelman. 1976. [Percolation and cluster distribution. i. cluster multiple labeling technique and critical concentration algorithm](#). *Physical Review B*, 14:3438–3445.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. 2019. [BERT for coreference resolution: Baselines and analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5803–5808, Hong Kong, China. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada. Association for Computational Linguistics.
- Vineet Kumar and Sachindra Joshi. 2016. [Non-sentential question resolution using sequence to sequence learning](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2022–2031, Osaka, Japan. The COLING 2016 Organizing Committee.
- Vineet Kumar and Sachindra Joshi. 2017. [Incomplete follow-up question resolution using retrieval based sequence to sequence learning](#). In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17*, page 705–714, New York, NY, USA. Association for Computing Machinery.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. [End-to-end neural coreference resolution](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197, Copenhagen, Denmark. Association for Computational Linguistics.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. [Higher-order coreference resolution with coarse-to-fine inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 687–692, New Orleans, Louisiana. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Qian Liu, Bei Chen, Haoyan Liu, Jian-Guang Lou, Lei Fang, Bin Zhou, and Dongmei Zhang. 2019a. [A split-and-recombine approach for follow-up query analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5316–5326, Hong Kong, China. Association for Computational Linguistics.
- Qian Liu, Bei Chen, Jian-Guang Lou, Ge Jin, and Dongmei Zhang. 2019b. [FANDA: A novel approach to perform follow-up query analysis](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6770–6777. AAAI Press.
- Eric Malmi, Sebastian Krause, Sascha Rothe, Daniil Mirylenka, and Aliaksei Severyn. 2019. [Encode](#),

- tag, realize: [High-precision text editing](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5054–5065, Hong Kong, China. Association for Computational Linguistics.
- Zhufeng Pan, Kun Bai, Yan Wang, Lianqiang Zhou, and Xiaojiang Liu. 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)*, pages 1824–1833, Hong Kong, China. Association for Computational Linguistics.
- Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. [Text matching as image recognition](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 2793–2799. AAAI Press.
- 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, ACL 2002*, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. Curran Associates, Inc.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Jun Quan, Deyi Xiong, Bonnie Webber, and Changjian Hu. 2019. [GECOR: An end-to-end generative ellipsis and co-reference resolution model for task-oriented dialogue](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4547–4557, Hong Kong, China. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. [U-net: Convolutional networks for biomedical image segmentation](#). In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer.
- Mike Schuster and Kuldip K Paliwal. 1997. [Bidirectional recurrent neural networks](#). *Trans. Sig. Proc.*, 45(11):2673–2681.
- 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 (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. [A hierarchical latent variable encoder-decoder model for generating dialogues](#). In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. [Neural responding machine for short-text conversation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1577–1586, Beijing, China. Association for Computational Linguistics.
- Hui Su, Xiaoyu Shen, Rongzhi Zhang, Fei Sun, Pengwei Hu, Cheng Niu, and Jie Zhou. 2019. [Improving multi-turn dialogue modelling with utterance ReWriter](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 22–31, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer networks](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2692–2700. Curran Associates, Inc.

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.
- Wei Wu, Fei Wang, Arianna Yuan, Fei Wu, and Jiwei Li. 2020. [CorefQA: Coreference resolution as query-based span prediction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6953–6963, Online. Association for Computational Linguistics.
- Yu Wu, Furu Wei, Shaohan Huang, Yunli Wang, Zhoujun Li, and Ming Zhou. 2019. [Response generation by context-aware prototype editing](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7281–7288.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. [Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 496–505, Vancouver, Canada. Association for Computational Linguistics.
- Rui Yan, Yiping Song, and Hua Wu. 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, SIGIR ’16*, page 55–64, New York, NY, USA. Association for Computing Machinery.
- Weinan Zhang, Yiming Cui, Yifa Wang, Qingfu Zhu, Lingzhi Li, Lianqiang Zhou, and Ting Liu. 2018. [Context-sensitive generation of open-domain conversational responses](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2437–2447, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Kun Zhou, Kai Zhang, Yu Wu, Shujie Liu, and Jingsong Yu. 2019. [Unsupervised context rewriting for open domain conversation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1834–1844, Hong Kong, China. Association for Computational Linguistics.