

Improving Multi-turn Dialogue Modelling with Utterance ReWriter

Hui Su^{1*}, Xiaoyu Shen^{2*}, Rongzhi Zhang³, Fei Sun⁴, Pengwei Hu⁵
Cheng Niu¹ and Jie Zhou¹

¹Pattern Recognition Center, Wechat AI, Tencent Inc, China

²MPI Informatics & Spoken Language Systems (LSV), Saarland Informatics Campus

³Institute of Software, University of Chinese Academy of Science

⁴Alibaba Group ⁵IBM Research, China

aaronsu@tencent.com, xshen@mpi-inf.mpg.de

Abstract

Recent research has made impressive progress in single-turn dialogue modelling. In the multi-turn setting, however, current models are still far from satisfactory. One major challenge is the frequently occurred coreference and information omission in our daily conversation, making it hard for machines to understand the real intention. In this paper, we propose rewriting the human utterance as a pre-process to help multi-turn dialogue modelling. Each utterance is first rewritten to recover all coreferred and omitted information. The next processing steps are then performed based on the rewritten utterance. To properly train the utterance rewriter, we collect a new dataset with human annotations and introduce a Transformer-based utterance rewriting architecture using the pointer network. We show the proposed architecture achieves remarkably good performance on the utterance rewriting task. The trained utterance rewriter can be easily integrated into online chatbots and brings general improvement over different domains.¹

1 Introduction

Dialogue systems have made dramatic progress in recent years, especially in single-turn chat and FAQ matching (Shang et al., 2015; Ghazvininejad et al., 2018; Molino et al., 2018; Chen et al., 2019). Nonetheless, multi-turn dialogue modelling still remains extremely challenging (Vinyals and Le, 2015; Serban et al., 2016, 2017; Shen et al., 2018a,b). The challenge is multi-sided. One most important difficulty is the frequently occurred coreference and information omission in our daily conversations, especially in pro-drop languages like Chinese or Japanese. From our preliminary study of 2,000 Chinese multi-turn con-

*Both authors contributed equally.

¹The code is available on <https://github.com/chin-gyou/dialogue-utterance-rewriter>.

	Context 1
Utterance 1 (Translation)	Human: 梅西有多高? Human: How tall is Messi?
Utterance 2	ChatBot: 官方说法他的身高是5英尺7英寸。 ChatBot: Officially he is 5ft 7 inches.
Utterance 3	Human: 他和C罗谁是最好的球员? Human: Who is the best, he or C.Ronaldo?
Utterance 3'	Human: 梅西和C罗谁是最好的球员? Human: Who is the best, Messi or C.Ronaldo?
	Context 2
Utterance 1	Human: 你最喜欢什么电影? Human: What movie do you like most?
Utterance 2	ChatBot: 泰坦尼克。 ChatBot: Titanic.
Utterance 3	Human: 为什么呢? Human: Why?
Utterance 3'	Human: 为什么最喜欢泰坦尼克? Human: Why do you like Titanic most?

Table 1: An example of multi-turn dialogue. Each utterance 3 is rewritten into **Utterance 3'**. Green means coreference and blue means omission.

versations, different degrees of coreference and omission exist in more than 70% of the utterances. Capturing the hidden intention beneath them requires deeper understanding of the dialogue context, which is difficult for current neural network-based systems. Table 1 shows two typical examples in multi-turn dialogues. “他”(he) from Context 1 is a coreference to “梅西”(Messi) and “为什么”(Why) from Context 2 omits the further question of “为什么最喜欢泰坦尼克”(Why do you like Titanic most)?. Without expanding the coreference or omission to recover the full information, the chatbot has no idea how to continue the talk.

To address this concern, we propose simplifying the multi-turn dialogue modelling into a *single-turn* problem by rewriting the current utterance. The utterance rewriter is expected to perform (1) coreference resolution and (2) information completion to recover all coreferred and omitted mentions. In the two examples from Table 1, each utterance 3 will be rewritten into utterance 3'. Afterwards, the system will generate a reply by only looking into the utterance 3' without considering the previous turns utterance 1 and 2. This simplification shortens the length of dialogue con-

text while still maintaining necessary information needed to provide proper responses, which we believe will help ease the difficulty of multi-turn dialogue modelling. Compared with other methods like memory networks (Sukhbaatar et al., 2015) or explicit belief tracking (Mrkšić et al., 2017), the trained utterance rewriter is model-agnostic and can be easily integrated into other black-box dialogue systems. It is also more memory-efficient because the dialogue history information is reflected in a single rewritten utterance.

To get supervised training data for the utterance rewriting, we construct a Chinese dialogue dataset containing 20k multi-turn dialogues. Each utterance is paired with corresponding manually annotated rewritings. We model this problem as an extractive generation problem using the Pointer Network (Vinyals et al., 2015). The rewritten utterance is generated by copying words from either the dialogue history or the current utterance based on the attention mechanism (Bahdanau et al., 2014). Inspired by the recently proposed Transformer architecture (Vaswani et al., 2017) in machine translation which can capture better intra-sentence word dependencies, we modify the Transformer architecture to include the pointer network mechanism. The resulting model outperforms the recurrent neural network (RNN) and original Transformer models, achieving an F1 score of over 0.85 for both the coreference resolution and information completion. Furthermore, we integrate our trained utterance rewriter into two online chatbot platforms and find it leads to more accurate intention detection and improves the user engagement. In summary, our contributions are:

1. We collect a high-quality annotated dataset for coreference resolution and information completion in multi-turn dialogues, which might benefit future related research.
2. We propose a highly effective Transformer-based utterance rewriter outperforming several strong baselines.
3. The trained utterance rewriter, when integrated into two real-life online chatbots, is shown to bring significant improvement over the original system.

In the next section, we will first go over some related work. Afterwards, in Section 3 and 4, our collected dataset and proposed model are introduced. The experiment results and analysis are

presented in Section 5. Finally, some conclusions are drawn in Section 6.

2 Related Work

2.1 Sentence Rewriting

Sentence rewriting has been widely adopted in various NLP tasks. In machine translation, people have used it to refine the output generations from seq2seq models (Niehues et al., 2016; Junczys-Dowmunt and Grundkiewicz, 2017; Grangier and Auli, 2017; Gu et al., 2017). In text summarization, reediting the retrieved candidates can provide more accurate and abstractive summaries (See et al., 2017; Chen and Bansal, 2018; Cao et al., 2018). In dialogue modelling, Weston et al. (2018) applied it to rewrite outputs from a retrieval model, but they pay no attention to recovering the hidden information under the coreference and omission. Concurrent with our work, Rastogi et al. (2019) adopts a similar idea on English conversations to simplify the downstream SLU task by reformulating the original utterance. Rewriting the source input into some easy-to-process standard format has also gained significant improvements in information retrieval (Riezler and Liu, 2010), semantic parsing (Chen et al., 2016) or question answering (Abujabal et al., 2018), but most of them adopt a simple dictionary or template based rewriting strategy. For multi-turn dialogues, due to the complexity of human languages, designing suitable template-based rewriting rules would be time-consuming.

2.2 Coreference Resolution

Coreference resolution aims to link an antecedent for each possible mention. Traditional approaches often adopt a pipeline structure which first identify all pronouns and entities then run clustering algorithms (Haghighi and Klein, 2009; Lee et al., 2011; Durrett and Klein, 2013; Björkelund and Kuhn, 2014). At both stages, they rely heavily on complicated, fine-grained features. Recently, several neural coreference resolution systems (Clark and Manning, 2016a,b) utilize distributed representations to reduce human labors. Lee et al. (2017) reported state-of-the-art results with an end-to-end neural coreference resolution system. However, it requires computing the scores for all possible spans, which is computationally inefficient on online dialogue systems. The recently proposed Transformer adopted the self-

attention mechanism which could implicitly capture inter-word dependencies in an unsupervised way (Vaswani et al., 2017). However, when multiple coreferences occur, it has problems properly distinguishing them. Our proposed architecture is built upon the Transformer architecture, but perform coreference resolution in a supervised setting to help deal with ambiguous mentions.

3 Dataset

To get parallel training data for the sentence rewriting, we crawled 200k candidate multi-turn conversational data from several popular Chinese social media platforms for human annotators to work on. Sensitive information is filtered beforehand for later processing. Before starting the annotation, we randomly sample 2,000 conversational data and analyze how often coreference and omission occurs in multi-turn dialogues. Table 2 lists the statistics. As can be seen, only less than 30% utterances have neither coreference nor omission and quite a few utterances have both. This further validates the importance of addressing the these situations in multi-turn dialogues.

	% Rate
Coreference	33.5
Omission	52.4
Neither	29.7

Table 2: Proportion of utterances containing coreference and omission in multi-turn conversation

In the annotation process, human annotators need to identify these two situations then rewrite the utterance to cover all hidden information. An example is shown in Table 1. Annotators are required to provide the rewritten utterance $3'$ given the original conversation [utterance 1,2 and 3]. To ensure the annotation quality, 10% of the annotations from each annotator are daily examined by a project manager and feedbacks are provided. The annotation is considered valid only when the accuracy of examined results surpasses 95%. Apart from the accuracy examination, the project manager is also required to (1) select topics that are more likely to be talked about in daily conversations, (2) try to cover broader domains and (3) balance the proportion of different coreference and omission patterns. The whole annotation takes 4 months to finish. In the end, we get 40k high-

quality parallel samples. Half of them are *negative* samples which do not need any rewriting. The other half are *positive* samples where rewriting is needed. Table 3 lists the statistics. The rewritten utterance contains 10.5 tokens in average, reducing the context length by 80%.

Dataset size:	40,000
Avg. length of original conversation:	48.8
Avg. length of rewritten utterance:	10.5

Table 3: Statistics of dataset. Length is counted in the unit of Chinese characters.

4 Model

4.1 Problem Formalization

We denote each training sample as $(H, U_n \rightarrow R)$. $H = \{U_1, U_2, \dots, U_{n-1}\}$ represents the dialogue history containing the first $n - 1$ turn of utterances. U_n is the n th turn of utterance, the one that needs to be rewritten. R is the rewritten utterance after recovering all coreferenced and omitted information in U_n . R could be identical to U_n if no coreference or omission is detected (negative sample). Our goal is to learn a mapping function $p(R|(H, U_n))$ that can automatically rewrite U_n based on the history information H . The process is to first encode (H, U_n) into a sequence of vectors, then decode R using the pointer network. The next section will explain the steps in order.

4.2 Encoder

We unfold all tokens in (H, U_n) into (w_1, w_2, \dots, w_m) . m is the number of tokens in the whole dialogue. An end-of-turn delimiter is inserted between each two turns. The unfolded sequence of tokens are then encoded with Transformer. We concatenate all tokens in (H, U_n) as the input, in hope that the Transformer can learn rudimentary coreference information within them by means of the self-attention mechanism. For each token w_i , the input embedding is the sum of its word embedding, position embedding and turn embedding:

$$I(w_i) = WE(w_i) + PE(w_i) + TE(w_i)$$

The word embedding $WE(w_i)$ and position embedding $PE(w_i)$ are the same as in normal Transformer architectures (Vaswani et al., 2017). We




Figure 1: Architecture of our proposed model. Green box is the Transformer encoder and pink box is the decoder. The decoder computes the probability λ at each step to decide whether to copy from the context or utterance.

add an additional turn embedding $TE(w_i)$ to indicate which turn each token belongs to. Tokens from the same turn will share the same turn embedding. The input embeddings are then forwarded into L stacked encoders to get the final encoding representations. Each encoder contains a self-attention layer followed by a feedforward neural network.:

$$\mathbf{E}^{(0)} = [I(w_1), I(w_2), \dots, I(w_m)]$$

$$\mathbf{E}^{(l)} = \text{FNN}(\text{MultiHead}(\mathbf{E}^{(l-1)}, \mathbf{E}^{(l-1)}, \mathbf{E}^{(l-1)}))$$

FNN is the feedforward neural network and $\text{MultiHead}(Q, K, V)$ is a multi-head attention function taking a query matrix Q , a key matrix K , and a value matrix V as inputs. Each self-attention and feedforward component comes with a residual connection and layer-normalization step, which we refer to Vaswani et al. (2017) for more details. The final encodings are the output from the L th encoder $\mathbf{E}^{(L)}$.

4.3 Decoder

The decoder also contains L layers, each layer is composed of three sub-layers. The first sub-layer is a multi-head self-attention:

$$\mathbf{M}^l = \text{MultiHead}(\mathbf{D}^{(l-1)}, \mathbf{D}^{(l-1)}, \mathbf{D}^{(l-1)})$$

$\mathbf{D}^{(0)} = R$. The second sub-layer is encoder-decoder attention that integrates $\mathbf{E}^{(L)}$ into the decoder. In our task, as H and U_n serve different purposes, we use separate key-value matrix for tokens

coming from the dialogue history H and those coming from U_n . The encoded sequence $\mathbf{E}^{(L)}$ obtained from the last section is split into $\mathbf{E}_H^{(L)}$ (encodings of tokens from H) and $\mathbf{E}_{U_n}^{(L)}$ (encodings of tokens from U_n) then processed separately. The encoder-decoder vectors are computed as follows:

$$\mathbf{C}(H)^l = \text{MultiHead}(\mathbf{M}^{(l)}, \mathbf{E}_H^{(L)}, \mathbf{E}_H^{(L)})$$

$$\mathbf{C}(U_n)^l = \text{MultiHead}(\mathbf{M}^{(l)}, \mathbf{E}_{U_n}^{(L)}, \mathbf{E}_{U_n}^{(L)})$$

The third sub-layer is a position-wise fully connected feed-forward neural network:

$$\mathbf{D}^{(l)} = \text{FNN}([\mathbf{C}(H)^l \circ \mathbf{C}(U_n)^l])$$

where \circ denotes vector concatenation.

4.4 Output Distribution

In the decoding process, we hope our model could learn whether to copy words from H or U_n at different steps. Therefore, we impose a soft gating weight λ to make the decision. The decoding probability is computed by combining the atten-

tion distribution from the last decoding layer:

$$\begin{aligned}
p(R_t=w|H, U_n, R_{<t}) &= \lambda \sum_{i:(w_i=w) \wedge (w_i \in H)} a_{t,i} \\
&+ (1-\lambda) \sum_{j:(w_j=w) \wedge (w_j \in U_n)} a'_{t,j} \\
a &= \text{Attention}(\mathbf{M}^{(L)}, \mathbf{E}_{U_n}^{(L)}) \\
a' &= \text{Attention}(\mathbf{M}^{(L)}, \mathbf{E}_H^{(L)}) \\
\lambda &= \sigma(\mathbf{w}_d^\top \mathbf{D}_t^L + \mathbf{w}_H^\top \mathbf{C}(H)_t^L + \mathbf{w}_U^\top \mathbf{C}(U_n)_t^L)
\end{aligned}$$

a and a' are the attention distribution over tokens in H and U_n respectively. $\mathbf{w}_d, \mathbf{w}_H$, and \mathbf{w}_U are parameters to be learned, σ is the sigmoid function to output a value between 0 and 1. The gating weight λ works like a sentinel to inform the decoder whether to extract information from the dialogue history H or directly copy from U_n . If U_n contains neither coreference nor information omission. λ would be always 1 to copy the original U_n as the output. Otherwise λ becomes 0 when a coreference or omission is detected. The attention mechanism is then responsible of finding the proper coreferred or omitted information from the dialogue history. The whole model is trained end-to-end by maximizing $p(R|H, U_n)$.

5 Experiments

We train our model to perform the utterance rewriting task on our collected dataset. In this section, we focus on answering the following two questions: (1) How accurately our proposed model can perform coreference resolution and information completion respectively and (2) How good the trained utterance rewriter is at helping off-the-shelf dialogue systems provide more appropriate responses. To answer the first question, we compare our models with several strong baselines and test them by both automatic evaluation and human judgement. For the second question, we integrate our rewriting model to two online dialogue systems and analyze how it affects the human-computer interactions. The following section will first introduce the compared models and basic settings, then report our evaluation results.

5.1 Compared Models

When choosing compared models, we are mainly curious to see (1) whether the self-attention based Transformer architecture is superior to other networks like LSTMs, (2) whether the pointer-based

generator is better than pure generation-based models and (3) whether it is preferred to split the attention by a coefficient λ as in our model. With these intentions, we implement the following four types of models for comparison:

1. **(L/T)-Gen**: Pure generation-based model. Words are generated from a fixed vocabulary.
2. **(L/T)-Ptr-Net**: Pure pointer-based model as in [Vinyals et al. \(2015\)](#). Words can only be copied from the input.
3. **(L/T)-Ptr-Gen**: Hybrid pointer+generation model as in [See et al. \(2017\)](#). Words can be either copied from the input or generated from a fixed vocabulary.
4. **(L/T)-Ptr- λ** : Our proposed model which split the attention by a coefficient λ .

(L/T) denotes the encoder-decoder structure is the LSTM or Transformer. For the first three types of models, we unfold all tokens from the dialogue as the input. No difference is made between the dialogue history and the utterance to be rewritten.

5.2 Experiment Settings

Transformer-based models We set the hidden size as 512. The attention has 8 individual heads and the encoder/decoder have 6 individual stacked layers. Models are optimized with the Adam optimizer. The initial learning rate is 0.0001 and batch size is 64. All hyperparameters are tuned base on the performance on the validation data.

LSTM-based Models We encode words with a single-layer bidirectional LSTM and decode with a uni-directional LSTM. We use 128-dimensional word embeddings and 256-dimensional hidden states for both the encoder and decoder.² The batch size is set as 128. Models are trained using Adagrad with learning rate 0.15 and initial accumulator value 0.1, same as in [See et al. \(2017\)](#).

General Setup We built our vocabulary based on character-based segmentation for Chinese scripts. For non-Chinese characters, like frequently mentioned entity names “Kobe” and “NBA”, we split them by space and keep all unique tokens which appear more than twice. The resulting vocabulary size is 5629 (4813 Chinese

²We tried increasing the dimension but find it degrades the performance.

	BLEU-1	BLEU-2	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	EM
L-Gen	65.49	55.38	38.69	65.57	48.57	66.38	47.14 80.18
L-Ptr-Gen	69.78	59.25	43.07	68.24	54.13	70.36	47.35 84.09
L-Ptr-Net	71.70	60.29	44.72	70.81	56.35	72.33	48.24 91.94
L-Ptr- λ	72.26	62.15	47.11	73.47	57.51	74.55	51.66 93.01
T-Gen	68.74	59.09	42.57	69.12	50.92	69.70	48.59 87.61
T-Ptr-Gen	70.67	62.80	45.17	73.96	53.14	72.07	49.86 89.62
T-Ptr-Net	75.10	66.89	48.11	76.10	58.51	75.54	53.30 94.71
T-Ptr- λ	77.85	68.21	52.47	78.49	60.53	77.70	55.84 98.14

Table 4: BLEU, ROUGE (F_1), and EM scores on the test set. EM score is split into the results on the positive (left) and negative (right) test samples. The first half is LSTM-based models and the second half is Transformer-based. Bold denotes best results.

characters and 816 other tokens), including the end-of-turn delimiter and a special UNK token for all unknown words. In the testing stage, all models decode words by beam search with beam size set to 4.

5.3 Quality of Sentence ReWriting

	Precision	Recall	F1
Lee et al. (2017)	0.82	0.78	0.80
L-Gen	0.76	0.66	0.71
L-Ptr-Gen	0.81	0.76	0.78
L-Ptr-Net	0.83	0.78	0.81
L-Ptr- λ	0.85	0.82	0.83
T-Gen	0.80	0.75	0.77
T-Ptr-Gen	0.85	0.81	0.83
T-Ptr-Net	0.88	0.87	0.88
T-Ptr- λ	0.93	0.90	0.92

Table 5: Precision, recall and F1 score of coreference resolution. First row is the current state-of-the-art coreference resolution model

Accuracy of Generation We first evaluate the accuracy of generation leveraging three metrics: BLEU, ROUGE, and the exact match score (EM) (the percentage of decoded sequences that exactly match the human references). For the EM score, we report separately on the positive and negative samples to see the difference. We report BLEU-1, 2, 4 scores and the F1 scores of ROUGE-1, 2, L. The results are listed in Table 4. We can have several observations in response to the three questions proposed in the beginning of Section 5.1:

1. Transformer-based models lead to signif-

icant improvement compare with LSTM-based counterparts. This implies the self-attention mechanism is helpful in identifying coreferred and omitted information. More analysis on how it helps coreference resolution can be seen in the next section.

2. The generation mode does not work well in our setting since all words can be retrieved from either H or U_n . Pointer-based models outperform the more complex generation-based and hybrid ones.
3. Separately processing H and U_n then combine their attention with a learned λ performs better than treating the whole dialogue tokens as a single input, though the improvement is less significant compared with previous two mentions.

Overall our proposed model achieves remarkably good performance, with 55.84% of its generations exactly matches the human reference on the positive samples. For negative samples, our model properly copied the the original utterances in 98.14% of the cases. It suggests our model is already able to identify the utterances that do not need rewriting. Future work should work on improving the rewriting ability on positive samples.

Coreference Resolution Apart from the standard metrics for text generation, we specifically test the precision, recall and F1 score of coreference resolution on our task. A pronoun or a noun is considered as properly coreferred if the rewritten utterance contains the correct mention in the corresponding referent. The result is shown in Table 5. To compare with current state-of-the-

History	U1: 你看莎士比亚吗 U2: 特别喜欢罗密欧与朱丽叶	U1: 你玩英雄联盟吗 U2: 是的
(Translation)	U1: Do you read Shakespeare U2: I especially like Romeo and Juliet	U1: Do you play League of Legends U2: Yes.
Utterance	U3: 喜欢哪个角色 U3: Which character do you like	U3: 什么时候开始的 U3: When did it start
Ground Truth	你喜欢罗密欧与朱丽叶哪个角色 Which character do you like in Romeo and Juliet	什么时候开始玩英雄联盟的 When did you start to play League of Legends
L-Gen	你喜欢莎士比亚吗 // Do you like Shakespeare	什么时候开始开始开始 // When start start start
L-Ptr-Gen	你喜欢罗密欧角色角色 // You like Romeo character character	什么时候开始的 // When did it start
L-Ptr-Net	你喜欢罗密欧与朱丽叶 // You like Romeo and Juliet	什么时候英雄联盟开始的 // When did League of Legends start
L-Ptr- λ	你喜欢罗密欧与朱丽叶角色 // You like Romeo and Juliet character	什么时候开始玩英雄联盟的 // When did you start to play League of Legends
T-Gen	你喜欢罗密欧与朱丽叶 // You like Romeo and Juliet	是什么时候开始玩的 // Yes When start to play
T-Ptr-Gen	你喜欢罗密欧与朱丽叶哪个 // Which do you like in Romeo and Juliet	什么时候开始的 // When did it start
T-Ptr-Net	你喜欢罗密欧与朱丽叶角色 // Character you like Romeo and Juliet	英雄联盟什么时候开始玩的 // League of Legends When did you start to play
T-Ptr- λ	你喜欢罗密欧与朱丽叶哪个角色 // Which character do you like Romeo and Juliet	什么时候开始玩英雄联盟的 // When did you start to play League of Legends

Table 6: Examples of rewritten utterances. Highlighted utterances are exactly the same as the ground truth.

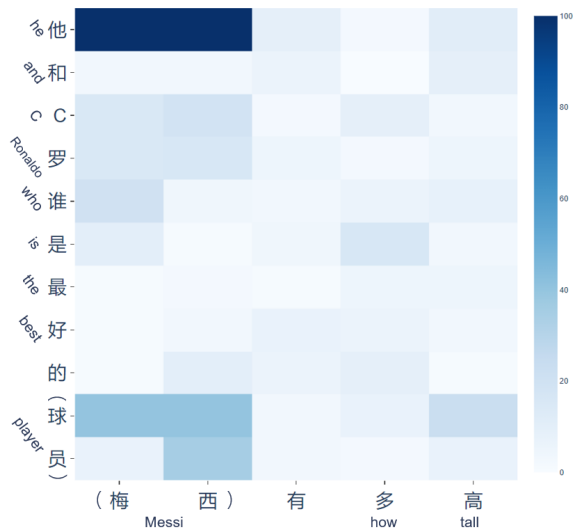


Figure 2: Visualization of the self-attention weights in Transformer. “他”(he) is properly aligned to “梅西”(Messi).

art models. We train the model from Lee et al. (2017) on our task and report the results on the first row. The result is quite consistent with the findings from the last section. Our final model outperforms the others by a large margin, reaching a precision score of 93% and recall score of 90%. It implies our model is already quite good at finding the proper coreference. Future challenges would be more about information completion. Figure 2 further provides an examples of how the Transformer can help implicitly learn the coreference resolution through the self-attention mechanism. The same example is also shown in Table 1. The pronoun “他”(he) in the utterance is properly aligned to the mention “梅西”(Messi) in the dialogue history, also partially to “球员”(player) which is the occupation of him. The implicitly learned coreference relation should be part of the reason that Transformers outperform LSTM models on the coreference resolution task.

Model	Recall	Precision	F1	Fluency
L-Gen	0.65	0.70	0.67	4.31
L-Ptr-Gen	0.70	0.74	0.72	4.52
L-Ptr-Net	0.78	0.81	0.79	4.74
L-Ptr- λ	0.80	0.82	0.81	4.82
T-Gen	0.71	0.74	0.73	4.74
T-Ptr-Gen	0.77	0.81	0.79	4.85
T-Ptr-Net	0.82	0.84	0.83	4.87
T-Ptr- λ	0.85	0.87	0.86	4.90
Human	-	-	-	4.97

Table 7: Recall, Precision, F1 score on information completion and Human evaluation results on fluency.

Information Completion Similar as coreference resolution, we evaluate the quality of information completeness separately. One omitted information is considered as properly completed if the rewritten utterance recovers the omitted words. Since it inserts new words to the original utterance, we further conduct a human evaluation to measure the fluency of rewritten utterances. We randomly sample 600 samples from our positive test set. Three participants were asked to judge whether the rewritten utterance is a fluent sentence with the score 1(not fluent)-5(fluent). The fluency score for each model is averaged over all human evaluated scores.

The results are shown in Table 7. Basically the condition is similar as in Table 5. T-Ptr- λ achieves the best performance, with the F1 score of 0.86. The performance is slightly worse than coreference resolution since information omission is more implicit. Retrieving all hidden information is sometimes difficult even for humans. Moreover, the fluency of our model’s generations is very good, only slightly worse than the human reference (4.90 vs 4.97). Information completeness does not have much effects on the fluency. Exam-

Task-Oriented Chatbot	
Context	U1: 北京天气怎么样 U2: 天气晴朗, 温度适宜
(Translation)	U1: How is the weather in Beijing U2: The weather is fine and the temperature is suitable
Utterance	U3: 那穿什么衣服合适 U3: 北京天气穿什么合适
Intention	U3: Then what clothes are suitable to wear U3: What clothes are suitable for weather in Beijing
	生活购物 城市天气
	Life Shopping City Weather
Chatbot Answer	您想要购买什么类型的衣服 根据天气推荐穿一件外套
	What type of clothes do you want to buy You'd better wear a coat according to the weather
Chit-Chat Chatbot	
Context	U1: 库里的三分真准啊 U2: 勇士今年又是冠军
	U1: Curry's 3-pointer is really good U2: The Warriors are the champion again this year
Utterance	U3: 我也觉得 U3: 我也觉得 勇士今年又是冠军
	U3: I agree U3: I agree that the Warriors are the champion again this year
Chatbot Answer	觉得什么 勇士真的厉害啊
	agree what The Warriors are so strong

Table 8: Examples of integrated test. Left column is the original system and right is the one with utterance rewriter. Blue words denote completed information by the utterance rewriter.

Model	Intention Precision	CPS
Original	80.77	6.3
With Rewrite	89.91	7.7

Table 9: Results of integrated testing. Intention precision for task-oriented and conversation-turns-per-session (CPS) for chitchat.

ples of rewritten utterances are shown in Table 6.

5.4 Integration Testing

In this section, we study how the proposed utterance rewriter can be integrated into off-the-shelf online chatbots to improve the quality of generated responses. We use our best model T-Ptr- λ to rewrite each utterance based on the dialogue context. The rewritten utterance is then forwarded to the system for response generation. We apply on both a task-oriented and chitchat setting. The results are compared with the original system having no utterance rewriter.

Task-oriented Our task-oriented dialogue system contains an intention classifier built on FastText(Bojanowski et al., 2017) and a set of templates that perform policy decision and slot-value filling sequentially. Intention detection is a most important component in task-oriented dialogues and its accuracy will affect all the following steps. We define 30 intention classes like weather, hotel booking and shopping. The training data contains 35,447 human annotations. With the combination of our rewriter, the intention classifier is able

to achieve a precision of 89.91%, outperforming the original system by over 9%. The improved intention classification further lead to better conversations. An example is shown in Table 8, a multi-turn conversation about the weather. The user first asks “How is the weather in Beijing”, then follows with a further question about “Then what clothes are suitable to wear”. The original system wrongly classified the user intention as shopping since this is a common conversational pattern in shopping. In contrast, our utterance rewriter is able to recover the omitted information “under the weather in Beijing”. Based on the rewritten utterance, the classifier is able to correctly detect the intention and provide proper responses.

Chitchat Our social chatbot contains two separate engines for multi-turn and single-turn dialogues. Each engine is a hybrid retrieval and generation model. In real-life applications, a user query would be simultaneously distributed to these two engines. The returned candidate responses are then reranked to provide the final response. Generally the model is already able to provide rather high-quality responses under the single-turn condition, but under multi-turn conversations, the complex context dependency makes the generation difficult. We integrate our utterance rewriter into the single-turn engine and compare with the original model by conducting the online A/B test. Specifically, we randomly split the users into two groups. One talks with the original system and the other talks with the system integrated with the utterance rewriter. All users are unconscious of the

details about our system. The whole test lasted one month. Table 9 shows the Conversation-turns Per Session (CPS), which is the average number of conversation-turns between the chatbot and the user in a session. The utterance rewriter increases the average CPS from 6.3 to 7.7, indicating the user is more engaged with the integrated model. Table 8 shows an example of how the utterance rewriter helps with the generation. After the rewriting, the model can better understand the dialogue is about the NBA team Warriors, but the original model feels confused and only provides a generic response.

6 Conclusion

In this paper, we propose improving multi-turn dialogue modelling by imposing a separate utterance rewriter. The rewriter is trained to recover the coreferred and omitted information of user utterances. We collect a high-quality manually annotated dataset and designed a Transformer-pointer based architecture to train the utterance rewriter. The trained utterance rewriter performs remarkably well and, when integrated into two online chatbot applications, significantly improves the intention detection and user engagement. We hope the collected dataset and proposed model can benefit future related research.

Acknowledgments

We thank all anonymous reviewers and the dialogue system team of Wechat AI for valuable comments. Xiaoyu Shen is supported by IMPRS-CS fellowship.

References

Abdalghani Abujabal, Rishiraj Saha Roy, Mohamed Yahya, and Gerhard Weikum. 2018. Never-ending learning for open-domain question answering over knowledge bases. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, pages 1053–1062. International World Wide Web Conferences Steering Committee.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Anders Björkelund and Jonas Kuhn. 2014. Learning structured perceptrons for coreference resolution with latent antecedents and non-local features. In

Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), volume 1, pages 47–57.

- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Ziqiang Cao, Wenjie Li, Sujian Li, and Furu Wei. 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)*, volume 1, pages 152–161.
- Bo Chen, Le Sun, Xianpei Han, and Bo An. 2016. Sentence rewriting for semantic parsing. *CoRR*, abs/1901.02998.
- Shiqian Chen, Chenliang Li, Feng Ji, Wei Zhou, and Haiqing Chen. 2019. Driven answer generation for product-related questions in e-commerce. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 411–419. ACM.
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. *arXiv preprint arXiv:1805.11080*.
- Kevin Clark and Christopher D Manning. 2016a. Deep reinforcement learning for mention-ranking coreference models. *arXiv preprint arXiv:1609.08667*.
- Kevin Clark and Christopher D Manning. 2016b. Improving coreference resolution by learning entity-level distributed representations. *arXiv preprint arXiv:1606.01323*.
- Greg Durrett and Dan Klein. 2013. Easy victories and uphill battles in coreference resolution. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1971–1982.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- David Grangier and Michael Auli. 2017. Quikedit: Editing text & translations via simple delete actions. *arXiv preprint arXiv:1711.04805*.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor OK Li. 2017. Search engine guided non-parametric neural machine translation. *arXiv preprint arXiv:1705.07267*.
- Aria Haghighi and Dan Klein. 2009. Simple coreference resolution with rich syntactic and semantic features. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 3-Volume 3*, pages 1152–1161. Association for Computational Linguistics.

- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2017. An exploration of neural sequence-to-sequence architectures for automatic post-editing. *arXiv preprint arXiv:1706.04138*.
- Heeyoung Lee, Yves Peirsman, Angel Chang, Nathanael Chambers, Mihai Surdeanu, and Dan Jurafsky. 2011. Stanford’s multi-pass sieve coreference resolution system at the conll-2011 shared task. In *Proceedings of the fifteenth conference on computational natural language learning: Shared task*, pages 28–34. Association for Computational Linguistics.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. *arXiv preprint arXiv:1707.07045*.
- Piero Molino, Huaixiu Zheng, and Yi-Chia Wang. 2018. Cota: Improving the speed and accuracy of customer support through ranking and deep networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 586–595. ACM.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1777–1788.
- Jan Niehues, Eunah Cho, Thanh-Le Ha, and Alex Waibel. 2016. Pre-translation for neural machine translation. *arXiv preprint arXiv:1610.05243*.
- Pushpendre Rastogi, Arpit Gupta, Tongfei Chen, and Lambert Mathias. 2019. Scaling multi-domain dialogue state tracking via query reformulation. *NAACL*.
- Stefan Riezler and Yi Liu. 2010. Query rewriting using monolingual statistical machine translation. *Computational Linguistics*, 36(3):569–582.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. *AAAI*.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. *AAAI*.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. *arXiv preprint arXiv:1503.02364*.
- Xiaoyu Shen, Hui Su, Wenjie Li, and Dietrich Klakow. 2018a. Nexus network: Connecting the preceding and the following in dialogue generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4316–4327.
- Xiaoyu Shen, Hui Su, Shuzi Niu, and Vera Demberg. 2018b. Improving variational encoder-decoders in dialogue generation. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in neural information processing systems*, pages 2440–2448.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Jason Weston, Emily Dinan, and Alexander Miller. 2018. Retrieve and refine: Improved sequence generation models for dialogue. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 87–92.