

Mining Clues from Incomplete Utterance: A Query-enhanced Network for Incomplete Utterance Rewriting

Shuzheng Si^{1,2*}, Shuang Zeng^{1,2*}, Baobao Chang^{1†}

¹Key Laboratory of Computational Linguistics, Peking University, MOE, China

²School of Software and Microelectronics, Peking University, China

sishuzheng@stu.pku.edu.cn {zengs, chbb}@pku.edu.cn

Abstract

Incomplete utterance rewriting has recently raised wide attention. However, previous works do not consider the semantic structural information between incomplete utterance and rewritten utterance or model the semantic structure implicitly and insufficiently. To address this problem, we propose a **QUE**ry-Enhanced Network (**QUEEN**). Firstly, our proposed query template explicitly brings guided semantic structural knowledge between the incomplete utterance and the rewritten utterance making model perceive where to refer back to or recover omitted tokens. Then, we adopt a fast and effective edit operation scoring network to model the relation between two tokens. Benefiting from extra information and the well-designed network, QUEEN achieves state-of-the-art performance on several public datasets.

1 Introduction

Multi-turn dialogue modeling, a classic research topic in the field of human-machine interaction, serves as an important application area for pragmatics (Leech, 2003) and Turing Test. The major challenge in this task is that interlocutors tend to use incomplete utterances for brevity, such as referring back to (i.e., coreference) or omitting (i.e., ellipsis) entities or concepts that appear in dialogue history. As shown in Table 1, the incomplete utterance u_3 refers to “Smith” (“史密斯”) from u_1 and u_2 using a pronoun “He” (“他”) and omits “the type of cuisine” (“菜肴的类型”) from u_2 . This may cause referential ambiguity and semantic incompleteness problems if we only read this single utterance u_3 , which is a common case of downstream applications like retrieval-based dialogue systems (Boussaha et al., 2019). Moreover, previous studies (Su et al., 2019; Pan et al., 2019)

Turns	Utterances with Translation
u_1	Smith needs to find an expensive restaurant nearby. 史密斯需要在附近找一家昂贵的餐馆。
u_2	Does Smith care the type of cuisine? 史密斯关心菜肴的类型吗?
u_3	No, he does not care. 不, 他不关心。
u'_3	No, Smith does not care about the type of cuisine. 不, 史密斯不关心菜肴的类型。

Table 1: An example in multi-turn dialogue including dialogue utterance history u_1 and u_2 , incomplete utterance u_3 and rewritten utterance u'_3 .

also find that coreference and ellipsis exist in more than 70% of the utterances, especially in pro-drop languages like Chinese. These phenomena make it imperative to effectively model dialogue in incomplete utterance scenarios.

To cope with this problem, previous works (Kumar and Joshi, 2016a; Elgohary et al., 2019; Su et al., 2019) propose the **Incomplete Utterance Rewriting (IUR)** task. It aims to rewrite an incomplete utterance into a semantically equivalent but self-contained utterance by mining semantic clues from the dialogue history. Then the generated utterance can be understood without reading dialogue history. For example, in Table 1, after recovering the referred and omitted information from u_3 into u'_3 , we could better understand this utterance comprehensively than before.

Early works use coreference resolution methods (Clark and Manning, 2016) to identify the entity that a pronoun refers to. However, they ignore the more common cases of ellipsis. So the text generation-based methods (Su et al., 2019; Pan et al., 2019) are introduced to generate the rewritten sequence from the incomplete sequence by jointly considering coreference and ellipsis problems. Though effective, generation models neglect a key trait of the IUR task, where the main semantic structure of a rewritten utterance is usually similar to the original incomplete utterance. So the

*Equal contribution.

†Corresponding author.

inherent structure-unawareness and uncontrollable feature of generation-based models impede their performances. For semantic structure-aware methods, Liu et al. (2020) utilize an edit operation matrix (e.g., substitution, insertion operations) to convert an incomplete utterance into a complete one. They formulate this task as a semantic segmentation problem with a CNN-based model (Ronneberger et al., 2015) on the matrix to capture the semantic structural relations between words implicitly. Xu et al. (2020) attempt to add additional semantic information to language models (Devlin et al., 2019) by annotating semantic role information but it is time-consuming and costly. Huang et al. (2021) propose a semi-autoregressive generator using a tagger to model the some considerable overlapping regions between the incomplete utterance and rewritten utterance, yet only implicitly learn the difference between them. Although these methods maintain some similarities between the incomplete utterance and the rewritten utterance (i.e., the overlap between them), it is difficult for these methods to explicitly model the semantic structure, especially the difference between the two utterances, ignoring the information in the incomplete utterance, such as which tokens are more likely to be replaced and which positions are more likely to require the insertion of new tokens. Therefore, there are still limitations of existing methods for IUR task, especially in jointly considering coreference and ellipsis cases and better utilizing semantic structural information.

This paper proposes a simple yet effective **QUE**ry-Enhanced Network (**QUEEN**) to solve the IUR task. QUEEN jointly considers coreference and ellipsis problems that frequently happen in multi-turn utterances. Specifically, we propose a straightforward query template featuring two linguistic properties and concatenate this query with utterances as input text. This query explicitly brings semantic structural guided information shared between the incomplete and the rewritten utterances, i.e., making model perceive where to refer back to or recover omitted tokens. We regard the rewritten utterance as the output from a series of edit operations on the incomplete utterance by constructing a token-pair edit operation matrix, which attempts to model the the overlap between the incomplete utterance between the rewritten utterance. Different from Liu et al. (2020), we adopt a well-designed edit operation scoring network on

the matrix to perform incomplete utterance rewriting, which is faster and more effective. QUEEN brings semantic structural information from linguistics into the model more explicitly and avoids unnecessary overheads of labeled data from other tasks. Experiments on several IUR benchmarks show that QUEEN outperforms previous state-of-the-art methods. Extensive ablation studies also confirm that the proposed query template makes key contributions to the improvements of QUEEN.

2 Methodology

Overview Our QUEEN mainly consists of two modules: query template construction module (Sec. 2.1) and edit operation scoring network module (Sec. 2.2). From two linguistic perspectives, the former module aims to generate a query template for each incomplete utterance, i.e., coreference-ellipsis-oriented query template, to cope with coreference and ellipsis problems. This query template explicitly hints the model where to refer back and recover omitted tokens. The latter module tries to capture the semantic structural relations between tokens by constructing an edit operation matrix. As shown in Figure 1, our goal is to learn a model to generate correct edit operations on this matrix and compute edit operation scores between token pairs so as to convert the incomplete utterance into the complete one.

2.1 Query Template Construction Module

By observing incomplete and rewritten utterance pairs in existing datasets, we find that pronouns and referential noun phrases in the incomplete utterance often need to be substituted by text spans in dialogue history. And ellipsis often occurs in some specific positions of incomplete utterance, conforming to a certain syntactic structure. In this module, we expect to encode these linguistic prior knowledge into the input of QUEEN. The query template is constructed as follows:

Coreference-oriented Query Template In order to make QUEEN perceive the positions of coreference that need to be substituted by text spans from dialogue history, we use a special token [COREF] to replace pronouns and referential noun phrases in the incomplete utterance so as to get our coreference-oriented query template. For example, the coreference-oriented query template of the incomplete utterance “No, he does not care” (“不, 他不关心”) is “No, [COREF] does not

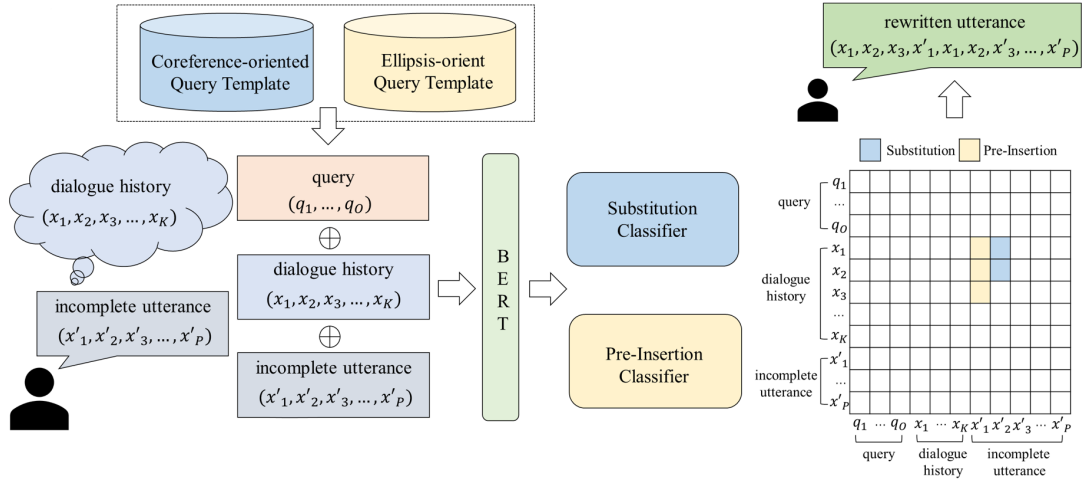


Figure 1: General architecture of QUEEN.

care” (“不, [COREF] 不关心”). To get the target complete utterance, this query explicitly tells the model we should replace the “He” (“他”) with text spans (such as ‘Smith’ (“史密斯”)) from dialogue history, rather than replacing other words. Here, we find all pronouns that required to be replaced using a predefined pronoun collection.

Ellipsis-oriented Query Template To make QUEEN perceive the positions of ellipsis that need to be inserted by text spans from dialogue history, we define a special token [ELLIP] and put it in a linguistically right place of the incomplete utterance. Since a self-contained utterance usually contains a complete S-V-O (Subject-Verb-Object) structure, if an incomplete utterance lack any of these key elements, we could assume there is a case of ellipsis in its corresponding text position. So we perform dependency parsing on the incomplete utterance to get the structure of the incomplete utterance. For example, the parsing result of the incomplete utterance “No, he does not care” (“不, 他不关心”) is an S-V structure and lack object element, thus we put [ELLIP] at the end of the sentence to get the ellipsis-oriented query template as “No, he does not care [ELLIP]” (“不, 他不关心 [ELLIP]”). Other circumstances are detailed in the Appendix.

Then we fuse these two query templates into the final coreference-ellipsis-oriented query template. For incomplete utterance “No, he does not care” (“不, 他不关心”), we get “No, [COREF] does not care [ELLIP]” (“不, [COREF] 不关心 [ELLIP]”) as our final query template. Under supervised setting, the models will perceive the positions to refer back and recover omitted

tokens for this utterance. For a multi-turn dialogue $d = (u_1, \dots, u_{N-1}, u_N)$ containing N utterances where $u_1 \sim u_{N-1}$ are dialogue history and the last utterance u_N needs to be rewritten, we could get the dialogue history text $s = (w_1^1, \dots, w_i^n, \dots, w_{L_n}^N)$ where w_i^n is the i -th token in the n -th utterance and L_n is the length of n -th utterance. We then concatenate our coreference-ellipsis-oriented query template with the dialogue history text to get our final input text $s' = (w_1^q, \dots, w_k^q, \dots, w_M^q, w_1^1, \dots, w_i^n, \dots, w_{L_n}^N)$ where w_k^q is the k -th token of the query template and M is the length of query template.

2.2 Edit Operation Scoring Network Module

Since pre-trained language models have been proven to be strongly effective on several natural language processing tasks, we employ BERT (Devlin et al., 2019) to encode our input text to get the contextualized hidden representation $H = (h_1^q, \dots, h_k^q, \dots, h_M^q, h_1^1, \dots, h_i^n, \dots, h_{L_n}^N)$. Our model attempts to predict whether there is an edit operation between each token pair. To this end, we define an operation scoring function as follows. Since the order of utterance is also important for dialogue, we further use RoPE (Su et al., 2021) to provide relative position information :

$$q_i^\alpha = W^\alpha h_i + b^\alpha \quad (1)$$

$$k_j^\alpha = W^\alpha h_j + b^\alpha \quad (2)$$

$$s_{ij}^\alpha = (R_i q_i^\alpha)^T (R_j k_j^\alpha) \quad (3)$$

where α is edit operation type including *Substitution* and *Pre-Insertion*. For different operations, we use different trainable parameters W^α and b^α .

Model	EM	B_2	B_4	R_2	R_L
T-Gen†	35.4	72.7	62.5	74.5	82.9
L-Ptr-λ†	42.3	82.9	73.8	81.1	84.1
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
T-Ptr-Gen†	53.1	84.4	77.6	85.0	89.1
T-Ptr-λ†	52.6	85.6	78.1	85.0	89.0
T-Ptr-λ + BERT†	57.5	86.5	79.9	86.9	90.5
CSRL	60.5	86.8	77.8	85.9	90.5
RUN†	66.4	91.4	86.2	90.4	93.5
QUEEN	70.1	92.1	86.9	90.9	94.6

Table 2: Experimental results on REWRITE. †: Results from Liu et al. (2020).

Model	EM	B_1	B_2	R_1	R_2
Syntactic†	-	84.1	81.2	89.3	80.6
L-Gen†	-	84.9	81.7	88.8	80.3
L-Ptr-Gen†	-	84.7	81.7	89.0	80.9
BERT‡	-	85.2	82.5	89.5	80.9
PAC†	-	89.9	86.3	91.6	82.8
CSRL‡	-	85.8	82.9	89.6	83.1
SARG	-	92.2	89.6	92.1	86.0
RUN†	49.3	92.3	89.6	92.4	85.1
QUEEN	53.5	92.4	89.8	92.5	86.3

Table 3: Experimental results on Restoration-200K. Additionally, we reproduce from the released code to get EM of RUN. †: Results from Liu et al. (2020). ‡: Results from Xu et al. (2020).

R is a transformation matrix from RoPE to inject position information and s_{ij}^α is the score for α -th edit operation from i -th token in dialogue history to j -th token in incomplete utterance.

During decoding for α -th operation, edit operation label \mathcal{Y}_{ij}^α satisfies:

$$\mathcal{Y}_{ij}^\alpha = \begin{cases} 1 & s_{ij}^\alpha \geq \theta \\ 0 & s_{ij}^\alpha < \theta \end{cases} \quad (4)$$

where θ is a hyperparameter. Once \mathcal{Y}_{ij}^α equals to 1, the edit operation α should be performed between token i and token j .

Since the label distribution of edit operation is very unbalanced (most elements are zeros), we employ Circle Loss (Sun et al., 2020) as our objective function to mitigate this problem :

$$\log(1 + \sum_{(i,j) \in \Omega_{pos}} e^{-s_{i,j}^\alpha}) + \log(1 + \sum_{(i,j) \in \Omega_{neg}} e^{s_{i,j}^\alpha}) \quad (5)$$

3 Experiments

3.1 Experimental Setup

Datasets We evaluate our model on four IUR benchmarks from different domains and lan-

Model	EM	B_4
Ellipsis Recovery†	50.4	74.1
GECOR 1†	68.5	83.9
GECOR 2†	66.2	83.0
RUN†	69.2	85.6
QUEEN	71.6	86.3

Table 4: Experimental results on TASK. †: Results and evaluation metrics from Liu et al. (2020).

Model	B_1	B_2	B_4	R_1	R_2	R_L
Copy†	52.4	46.7	37.8	72.7	54.9	68.5
Pronoun Sub†	60.4	55.3	47.4	73.1	63.7	73.9
L-Ptr-Gen†	67.2	60.3	50.2	78.9	62.9	74.9
RUN†	70.5	61.2	49.1	79.1	61.2	74.7
QUEEN	72.4	65.2	54.4	82.5	68.1	81.8

Table 5: Experimental results on CANARD. †: Results and evaluation metrics from Liu et al. (2020).

guages: REWRITE (Chinese, Su et al., 2019), Restoration-200K (Chinese, Pan et al., 2019), TASK (English, Quan et al., 2019), CANARD (English, Elgohary et al., 2019). More statistical details for datasets are shown in the Appendix.

Evaluation We use BLEU (Papineni et al., 2002), ROUGE (Lin, 2004) and the exact match (EM score) as our evaluation metrics.

Baseline Models We compare our model with a large number of baselines and state-of-the-art models. (i) **Baselines and Generation models** include sequence to sequence model (L-Gen) (Bahdanau et al., 2014), the hybrid pointer generator network (L-Ptr-Gen) (See et al., 2017a), the basic transformer model (T-Gen) (Vaswani et al., 2017) and the transformer-based pointer generator (T-Ptr-Gen) (See et al., 2017a), Syntactic (Kumar and Joshi, 2016b), PAC (Pan et al., 2019), L-Ptr-λ and T-Ptr-λ (Su et al., 2019), GECOR (Quan et al., 2019). Above methods need to generate rewritten utterances from scratch, neglecting the semantic structure between a rewritten utterance and the original incomplete utterance. (ii) **Structure-aware models** include CSRL (Xu et al., 2020), RUN (Liu et al., 2020), SARG (Huang et al., 2021).

3.2 Results and Analysis

Main Results We report the experiment results in Table 2, Table 3, Table 4 and Table 5. On all datasets with different languages and evaluation metrics, our approach outperforms all previous state-of-the-art methods. The improvement in EM shows that our model has a stronger ability to

find the correct span, due to our model making full use of the prior information of semantic structure from our coreference-ellipsis-oriented query template. On the Chinese datasets Table 2 and Table 3, QUEEN outperforms previous methods. Since Chinese is a pro-drop language where coreference and ellipsis often happen, the improvement confirms that QUEEN is superior in finding the correct ellipsis and referring back positions. The results on data sets of different domains and languages also show that our model is robust and effective.

Ablation Study To verify the effectiveness of the query in our model, we present an ablation study in Table 8 in the Appendix. It is clear that query is important to improve performance on all evaluation metrics. Meanwhile, only using coreference-oriented or ellipsis-oriented template still improves the performance, as it can also bring semantic structure information.

Inference Speed Meanwhile, we compare inference speed between our Edit Operation Scoring Network and the current fastest model RUN in A.5. As we remove the heavy model U-net of RUN and apply the simple network, our Edit Operation Scoring Network is faster than previous SOTA RUN.

Case Study We also do a case study in the Appendix. Our model avoids the uncontrolled situations that the generation-based model is prone to, and our model can more easily capture the correct semantic span.

4 Conclusion

We propose a simple yet effective query-enhanced network for IUR task. Our well-designed query template explicitly brings semantic structural information to improve the ability to predict correct edit operations between incomplete utterance and complete one. Experiments have shown that our model with this well-designed query achieves promising results than previous methods.

5 Acknowledgements

This paper is supported by the National Science Foundation of China under Grant No.61876004, 61936012, the National Key R&D Program of China under Grand No.2020AAA0106700.

6 Ethics Consideration

We collect our data from public datasets that permit academic use. The open-source tools we use

are freely accessible online without copyright conflicts.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Basma El Amel Boussaha, Nicolas Hernandez, Christine Jacquin, and Emmanuel Morin. 2019. [Deep retrieval-based dialogue systems: A short review](#). *CoRR*, abs/1907.12878.
- Wanxiang Che, Yunlong Feng, Libo Qin, and Ting Liu. 2020. N-ltp: A open-source neural chinese language technology platform with pretrained models. *arXiv preprint arXiv:2009.11616*.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. 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.
- Jie Hao, Linfeng Song, Liwei Wang, Kun Xu, Zhaopeng Tu, and Dong Yu. 2021. [RAST: Domain-robust dialogue rewriting as sequence tagging](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4913–4924, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mengzuo Huang, Feng Li, Wuhe Zou, and Weidong Zhang. 2021. [SARG: A novel semi autoregressive generator for multi-turn incomplete utterance restoration](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI*

- 2021, *Virtual Event, February 2-9, 2021*, pages 13055–13063. AAAI Press.
- 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*.
- Vineet Kumar and Sachindra Joshi. 2016a. [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. 2016b. 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.
- Geoffrey Leech. 2003. [Pragmatics and dialogue](#). In *The Oxford handbook of computational linguistics*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Qian Liu, Bei Chen, Jian-Guang Lou, Bin Zhou, and Dongmei Zhang. 2020. Incomplete utterance rewriting as semantic segmentation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- 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.
- 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*, pages 311–318, Philadelphia, Pennsylvania, USA. 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 4539–4549.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017a. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1073–1083. Association for Computational Linguistics.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017b. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- 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.
- Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2021. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*.
- Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. 2020. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6398–6407.
- 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, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Kun Xu, Haochen Tan, Linfeng Song, Han Wu, Haisong Zhang, Linqi Song, and Dong Yu. 2020. [Semantic Role Labeling Guided Multi-turn Dialogue ReWriter](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6632–6639, Online. Association for Computational Linguistics.
- Ningyu Zhang, Xiang Chen, Xin Xie, Shumin Deng, Chuanqi Tan, Mosha Chen, Fei Huang, Luo Si, and Huajun Chen. 2021. [Document-level relation extraction as semantic segmentation](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 3999–4006. ijcai.org.

A Experimental Details

A.1 Constructing Supervision

The expected supervision for our model is the edit operation matrix, but existing datasets only contain rewritten utterances. So we adopt Longest Common Subsequence (LCS) and 'Distant Supervision' (Liu et al., 2020) to get correct supervision, which contains edit operations *Substitute* and *Pre-Insert*.

A.2 Details in Query Construction

Coreference-oriented Query During the training, we use the ground truth of pronouns and referential noun phrases to construct the coreference-oriented query. During the inference, we use the constructed pronoun collection to construct the coreference-oriented query, which contains pronouns and referential noun phrases from training data and common pronouns (except “我”, “me” etc.).

Ellipsis-oriented Query Construction If the parsing result of the incomplete utterance is an S-V (Subject-Verb) structure and lacks subject element, we insert an [ELLIP] at the end of the incomplete utterance as the query. When there is not the S-V structure after parsing, we insert an [ELLIP] at the beginning of the incomplete utterance as the query. In other cases, we insert [ELLIP] at both the beginning and end of the incomplete utterance as the query. We use spaCy¹ for English and LTP(Che et al., 2020) for Chinese to get the result of parsing.

Extra findings During the experiment, we find two interesting points: (i) As [COREF] and [ELLIP] are sparse respectively, we use a unified token [UNK] to replace [COREF] and [ELLIP] in the query to relieve the sparsity. (ii) In most cases, if there is the referring back in the utterance, there is generally no ellipsis in the utterance. Redundant [ELLIP] tokens can't bring correct guided information in this case. Therefore, once we construct Coreference-oriented Query Template successfully, we will not try to construct the Ellipsis-oriented Query Template. Our experimental results are improved by the above two tricks.

A.3 Datasets details

Some statistics are shown in Table 6. REWRITE and Restoration-200K are constructed from Chi-

¹<https://spacy.io/>

	Restoration-200K REWRITE TASK CANARD			
Language	Chinese	Chinese	English	English
His Avglen	25.8	17.7	52.6	85.4
Inc Avglen	8.6	6.5	9.4	7.5
Rew Avglen	12.4	10.5	11.3	11.6

Table 6: Statistics of different datasets. 'Avglen' for average length, 'His' for historical utterance, 'Inc' for incomplete utterance, and 'Rew' for rewritten utterance.

Model	Inference Speed	Speedup
RUN+BERT	1.69it/s	1.00×
QUEEN	2.13it/s	1.26×

Table 7: The inference speed comparison between RUN and QUEEN on REWRITE dataset.

nese Open-Domain Dialogue. TASK is from English Task-oriented Dialogue. CANARD is constructed from English Context Question Answering. We follow the same data split as their original paper.

A.4 Details in Scoring Network

Hyper-parameters We implement our model on top of a BERT-base model (Devlin et al., 2019). We initialize QUEEN with bert-base-uncased for English and bert-base-chinese for Chinese. We use Adam (Kingma and Ba, 2015) with learning rate $1e-5$. The batch size is set to 16 for REWRITE and TASK, 12 for Restoration-200K, 4 for on CANARD. Meanwhile, θ in Equation 4 is set to 0.1 for REWRITE and TASK, 0.05 for Restoration-200K and CANARD.

Circle Loss We tune Circle Loss the same as Zhang et al. (2021) and Su^2 . We refer readers to their paper for more details.

A.5 Inference Speed

To compare the inference speed between the current fastest model RUN and our Edit Operation Scoring Network, we conduct experiments using the code released³. Both models are implemented in PyTorch on a single NVIDIA V100. The batch size is set to 16. Meanwhile, In order to fairly compare the speed of the two networks, we performed Distant Supervision and Query Construction before comparing. The results are shown in Table 7.

²A helpful Chinese blog: <https://kexue.fm/archives/7359>

³https://github.com/microsoft/ContextualSP/tree/master/incomplete_utterance_rewriting

VARIANT	EM	B_1	B_2	B_3	B_4	R_1	R_2	R_L
QUEEN	70.1	94.7	92.1	89.5	86.9	96.0	90.9	94.6
-w/o query	67.4	92.9	90.5	88.1	85.7	95.2	90.1	94.0
-w/o CQT	67.9	93.5	91.0	88.5	86.0	95.5	90.4	94.2
-w/o EQT	67.4	93.5	91.1	88.5	85.9	95.4	90.3	94.3

Table 8: The ablation results on the development set of REWRITE. “w/o query” means that we do not append a designed query before encoding that semantic information into our model. “w/o CQT” means that we only perform Ellipsis-oriented Query Template. “w/o EQT” means that we only perform Coreference-oriented Query Template and use the incomplete utterance as query if match fails.

A.6 Case Study

Table 9 gives 3 examples that indicate the representative situations as Hao et al. (2021). The first example illustrates the cases when RUN inserts unexpected characters into the wrong places. T-Ptr-Gen just copies the incomplete utterance. Due to our generated query, the position that needs to be inserted has been explicitly promoted by the query. The second example shows a common situation for generation-based models. T-Ptr-Gen messes up by repeating stupidly. However, this situation doesn’t happen to our model, as it is not a generation-based model. The last example refers to a long and complex entity. For these cases, it is easier for our model to get the correct span. This is because our model learns the span boundaries from the edit operation matrix. Compared to the generation-based model, we don’t generate sentences from scratch and this reduces the difficulty. Meanwhile, our model is not based on CNN as RUN, which suffers from the limitation of receptive-field to find a longer span.

A.7 Limitation

One limitation of current edit-based IUR models, is that only tokens that have appeared in the history dialogue can be selected. Therefore, these models, including ours, cannot generate novel words, e.g., conjunctions and prepositions, to cater to other metrics, like fluency. However, this can be alleviated by incorporating an additional word dictionary as See et al. (2017b) and Liu et al. (2020) deals with the out-of-vocabulary (OOV) words to improve fluency. For fairness, we keep the same words during the experiment as RUN to mitigate it. We will consider this question as a promising direction for future works.

Example # 1

Historical Utterance 1	我意见很大 I have a lot of complaints
Historical Utterance 2	有意见保留 Keep it yourself if there's any
Incomplete Utterance	不想保留 Don't want to keep it myself
Gold	不想保留意见 Don't want to keep the complaints myself
T-Ptr-Gen	不想保留 Don't want to keep it myself
RUN	意见不想意保留 Complaints don't want to keep the complaints myself
QUEEN	不想保留意见 Don't want to keep the complaints myself

Example # 2

Historical Utterance 1	你帮我考雅思 Please help me on IELTS
Historical Utterance 2	雅思第一项考什么 What is tested first for IELTS
Incomplete Utterance	考口语啊 It's oral test
Gold	雅思第一项考口语啊 It's oral test for IELTS
T-Ptr-Gen	考口语考口语啊 It's oral test oral test
RUN	雅思第一项考口语啊 It's oral test for IELTS
QUEEN	雅思第一项考口语啊 It's oral test for IELTS

Example # 3

Historical Utterance 1	帮我找一下西安到商洛的顺风车 Can you help me find a free ride from Xi'an to Shangluo
Historical Utterance 2	哪的 Where is it
Incomplete Utterance	能不能找到 Can you find any
Gold	能不能找到西安到商洛的顺风车 Can you find any free ride from Xi'an to Shangluo
T-Ptr-Gen	能不能找到商洛的顺风车 Can you find any free ride to Shangluo
RUN	能不能找到商洛的顺风车 Can you find any free ride to Shangluo
QUEEN	能不能找到西安到商洛的顺风车 Can you find any free ride from Xi'an to Shangluo