

Semantic Role Labeling Guided Multi-turn Dialogue ReWriter

Kun Xu¹, Haochen Tan², Linfeng Song¹, Han Wu², Haisong Zhang¹, Linqi Song², Dong Yu¹

¹Tencent AI Lab

²City University of Hong Kong

{kxkunxu, lfsong, hansongzhang, dyu}@tencent.com

{haochetan2-c, hanwu32-c}@my.cityu.edu.hk

{linqi.song}@cityu.edu.hk

Abstract

For multi-turn dialogue rewriting, the capacity of effectively modeling the linguistic knowledge in dialog context and getting rid of the noises is essential to improve its performance. Existing attentive models attend to all words without prior focus, which results in inaccurate concentration on some dispensable words. In this paper, we propose to use semantic role labeling (SRL), which highlights the core semantic information of *who* did *what* to *whom*, to provide additional guidance for the rewriter model. Experiments show that this information significantly improves a RoBERTa-based model that already outperforms previous state-of-the-art systems.

1 Introduction

Recent research (Vinyals and Le, 2015; Li et al., 2016; Serban et al., 2017; Zhao et al., 2017; Shao et al., 2017) on dialogue generation has been achieving impressive progress for making single-turn responses, while producing coherent multi-turn replies still remains extremely challenging. One important factor that contributes to this difficulty is coreference and information omission, where mention is dropped or replaced by a pronoun for simplicity. These phenomena dramatically introduce the requirements for long-distance reasoning, as they frequently occurred in our daily conversations, especially in pro-drop languages like Chinese and Japanese.

To tackle these problems, sentence rewriting was introduced to ease the burden of dialogue models by simplifying the multi-turn dialogue modeling into a single-turn problem. Several approaches (Su et al., 2019; Zhang et al., 2019; Elgohary et al., 2019) have been proposed to address the rewriting task. Conceptually, these models follow the conventional encoder-decoder architecture

Utterance 1	需要粤语 (I may need Cantonese.)
Utterance 2	粤语 _{ARG0} 是普通话 _{ARG1} 吗 (Is Cantonese Mandarin ?)
Utterance 3	不算 _{predicate} 吧 (Maybe Not.)
Utterance 3'	粤语不算普通话吧 (Cantonese may be not Mandarin.)

Table 1: One example of multi-turn dialogue. The goal of dialogue rewriting is to rewrite utterance 3 into 3'.

that first encodes the dialogue context into a distributional representation and then decodes it to the rewritten utterance. Their decoders mainly use global attention methods that attends to all words in the dialogue context without prior focus, which may result in inaccurate concentration on some dispensable words. We also observe that the accuracy of their models significantly decreases when working on long dialogue contexts. This observation is expected since if the text is lengthy, it would be quite difficult for deep learning models to understand as it suffers from noise and pays vague attention to the text components.

Motivated by these observations, we propose to incorporate the information of Semantic role labeling (SRL) (Gildea and Jurafsky, 2002; Palmer et al., 2010) to improve sentence rewriting. SRL is broadly used to identify the predicate-argument structures of a sentence, where these structures could capture the main semantic information of *who* did *what* to *whom*. As a result, we believe that it can pick out the important words, which are semantically most related to the utterance that needs to be rewritten. As shown in Table 1, our SRL system is able to find that the ARG0 and ARG1 of “不算”(is not) are “粤语”(Cantonese) and “普通话”(Mandarin), respectively. Consequently, our rewriting model can correctly generate the correct output (utterance 3'), which covers all dropped information. We can see that SRL can guide our




Figure 1: The input representation of a running example. We should point out that some tuples that do not contain words in the rewritten utterances could also be used as input predicate-argument triples.

rewriting model to focus on the semantically important words in the dialogue history, especially the omitted information that appears in previous turns.

For more details, we first take an SRL parser to recognize the predicate-argument (PA) structures from dialog contexts, before encoding that semantic information into our model. Since conventional SRL benchmarks only contain sentence-level annotations, existing pretrained SRL parsers (Khashabi et al., 2018; Gardner et al., 2018) can fail to extract the cross-turn PA structures in dialogues. To address this problem, we extend the traditional SRL to the conversational scenario by additionally annotating a dialogue dataset with standard SRL labels.

Our rewriting model is based on a pre-trained RoBERTa model (Liu et al., 2019) that takes the outputs of SRL parsing and dialogue history as its inputs, before generating rewriting outputs word by word. Experimental results show that even without the SRL information, our model already outperforms previous state-of-the-art models by a large margin. Augmenting the SRL information, the model performance is further improved significantly *without* adding any new parameters.

2 Task Definition

Formally, an input for dialogue rewriting is a dialogue session $c = (u_1, \dots, u_N)$ of N utterances, and u_N is the most recent utterance that needs to be revised. The output is r , the resulting utterance after recovering all coreference and omitted information in u_N . Our goal is to learn a model that can automatically rewrite u_N based on the dialogue context.

3 Model

Given a dialogue context c , we first apply an SRL parser to identify the predicate-argument structures z ; then conditioned on c and z , the rewritten

utterance is generated as $p(r|c, z)$. The backbone of our infrastructure is similar to the transformer blocks in Dong et al. (2019), which supports both bi-directional encoding and uni-directional decoding flexibly via specific self-attention masks. Specifically, we concatenate z , c and r as a sequence, feeding them into our model for training; during decoding, our model takes the z and c before generating the rewritten utterance word by word. Our model uses a pre-trained Chinese RoBERTa (Liu et al., 2019) for rich features.

3.1 Conversational SRL

SRL has long been treated as a sentence-internal task, and its major benchmarks (Carreras and Màrquez, 2005; Pradhan et al., 2013) contains only sentence-level annotations. We extend SRL to fit the conversational scenario by allowing SRL parsers to search for potential arguments over the whole conversation. As there is no publicly available data with paragraph-level SRL annotations, we directly annotate inter- and cross-utterance arguments for predicates on a public dialogue dataset, Duconv (Wu et al., 2019)¹. Specifically, we annotated 3,000 dialogue sessions, including 33,673 predicates in 27,198 utterances. Among them, 21.89% arguments are not in the same turn with their predicates, respectively. Considering existing standard SRL benchmarks may also be helpful, we first pre-train our SRL model (Shi and Lin, 2019) on the training set of CoNLL 2012 (117,089 examples) and fine-tune it on our annotations. In our experiments, we employ this conversational SRL model to recognize the predicate-argument structures for the dialogue context.

3.2 Input Representation for ReWriter

For each token, its input representation is obtained by summing the embeddings for word, semantic role and position. One example is shown in Fig-

¹More annotation details could be found in the Appendix.

	BLEU-1	BLEU-2	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	EM
Trans-Gen	78.18	70.31	51.85	83.1	67.84	81.98	24.12
Trans-Pointer	83.22	78.32	64.08	87.89	77.94	86.88	36.54
Trans-Hybrid	82.92	77.65	62.54	87.59	76.91	86.66	35.03
Su et al. (2019)	85.41	81.67	70.00	89.75	81.84	88.56	46.33
BERT	88.21	85.17	75.64	90.73	84.35	89.47	57.36
BERT + SRL							
w/ Bi-mask	88.89	85.88	76.36	90.92	85.00	89.72	58.36
w/ Triple-mask	89.66	86.78	77.76	91.82	85.87	90.52	60.49
BERT + Partial-SRL	89.46	86.57	77.75	91.60	85.60	90.50	59.15
BERT + Gold-SRL	93.34	91.38	84.97	94.94	90.45	93.86	71.96

Table 2: Evaluation results on the dataset of Su et al. (2019).

ure 1 and details are described in the following:

- The input is the concatenation of PA structures, dialog context, and rewritten utterance. Note that a PA structure is essentially in a tree format, where the root is a predicate and its children are corresponding semantic arguments. For the linearization, we decomposing each PA structure into several triples of the form $\langle predicate, role, argument \rangle$ and concatenate them in a random order. A special end-of-utterance token (i.e., [EOS]) is appended to the end of each utterance for separation. Another begin-of-utterance token (i.e., [BOS]) is also added at the beginning of the rewritten utterance. The final hidden state of the last token in the final layer is used to predict the next token during generation.
- We expand the segment-type embeddings of BERT to distinguish different types of tokens. In particular, the type embedding E_A is added for the rewritten utterance, as well as dialogue utterances generated by the same speaker in the context; the type embedding E_B is used for the other speaker; E_{SRL} is used as the type embedding of the tokens in predicate-argument triples. Position embeddings are added according to the token position in each utterance. The input embedding is the summation of word embedding, segment embedding, and position embedding.

3.3 Attention Mask

Similar to TransferTransfo (Wolf et al., 2019), we apply a future mask on the rewritten sequence, that is, the tokens in the rewritten utterance only attend on previous tokens in self-attention layers. Recall that, we linearize a PA structure into a concatenation sequence of triples. Since these triples are randomly ordered, it may inevitably introduce noisy information when using a sequence encoder. To better reflect its structural information, we elaborate the attention mask on PA sequence: the to-

kens in the same PA triple have bidirectional attentions while tokens in different PA triples can not attend each other. And the position embeddings of tokens in the PA sequence are added according to their positions in each distinct triple rather than the total PA sequence. In experiments, we find using these two designs help our model to more efficiently use the SRL information. We leave a more detailed discussion in Session 4.

3.4 Training

We employ the NLL loss to train our model:

$$\mathcal{L} = - \sum_{t=1}^T \log p(r_t | c, z, r_{<t}; \theta)$$

where θ represents the model parameters, T is the length of the target response r , and $r_{<t}$ denotes previously generated words.

4 Experiments

We evaluate our model on two rewrite datasets, which are built by Su et al. (2019) and Pan et al. (2019). Both of these two datasets are generated by crawling multi-turn conversational data from several popular Chinese social media platforms. Specifically, the dataset of Su et al. (2019) contains 17,890 examples, which are further split as 80%/10%/10% for training/development/testing, respectively. The dataset of Pan et al. (2019) contains 204k examples, where 194k/5k/5k are for training/development/testing.

The hyper-parameters used in our model are listed as follows. The network parameters of our model are initialized using RoBERTa. The batch size is set to 32. We use Adam (Kingma and Ba, 2014) with learning rate $5e-5$ to update parameters.

Results and Discussion. Following previous works, we used BLEU, ROUGE, and the exact

	B1	B2	R1	R2
Trans-Pointer	84.70	81.70	89.00	80.90
BERT	85.21	82.51	89.53	83.18
BERT + SRL	85.77	82.85	89.59	83.08

Table 3: Evaluation results on the dataset of Pan et al. (2019). B_n represents n-gram BLEU score and R_n represents n-gram ROUGE score.

match score (EM) (the percentage of decoded sequences that exactly match the human references). We implemented three baselines that use the same transformer-based encoder but differ in the choice of the decoder. Specifically, *Trans-Gen* uses a pure generation decoder which generates words from a fixed vocabulary; *Trans-Pointer* applies a pure pointer-based decoder (Vinyals et al., 2015) which can only copy the word from the input; *Trans-Hybrid* uses a hybrid pointer+generation decoder as in See et al. (2017), which can either copy the words from the input or generate words from a fixed vocabulary. Table 2 and Table 3 summarizes the results of our model and these baselines.

We can see that even without the SRL information, our model still significantly outperforms these baselines on two datasets, indicating that adapting a pre-trained language model could greatly improve the performance of such a generation task. We can also see that the model with the pointer-based decoder achieves better performance than the generation-based and the hybrid one, which is similar to the observation as in Su et al. (2019). This result is expected since there is a high chance the coreference or omission could be perfectly resolved by only using previous dialogue turns. In addition, we find that incorporating the SRL information can further improve the performance by at 1.45 BLEU-1 and 1.6 BLEU-2 points, achieving the state-of-the-art performances on the dataset of Su et al. (2019).

Let us first look at the impact of attention mask design on our model. To incorporate the SRL information into our model, we view the linearized predicate-argument structures as a regular utterance (say u_{pa}) and append it in the front of the input. We experimented with two choices of attention masks. Specifically, the first one is a bidirectional mask (referred as Bi-mask), that is, words in u_{pa} could attend each other; the second one (referred as Triple-mask) only allows words to attend its neighbors in the same triple, i.e., words

in different triples are not visible to each other. From Table 2, we can see that the latter one is significantly better than the first one. We think the main reason is that the second design independently encode each predicate-argument triple, which prevents the unnecessary triple-internal attentions, better mimicking the SRL structures.

Since our framework works in a pipeline fashion, one bottleneck of our system can lie in the performance of the SRL parser. One natural question is how accurate our SRL parser can be and how much performance improvement for the rewriter model we could have by introducing the SRL information. To investigate this, we employ a conventional SRL parser² to analyze the gold rewritten utterance. These extracted PA structures are *considered* as gold SRL annotations to measure the accuracy of our conversational SRL parser. In particular, we evaluate our SRL parser on the micro-averaged F1 over the (*predicate*, *argument*, *label*) tuples. We find our SRL parser achieves 75.66 precision, 74.47 recall, and 75.06 F₁. On the other hand, we use the gold SRL results instead of our SRL parsing results to train and test the model (referred as BERT+Gold-SRL). From Table 2, we can see that all evaluation scores are significantly improved. This result indicates that the performance of our rewriter model is highly relevant to the SRL parser, and the performance of our current SRL parser is still far from satisfactory, which we leave for future work.

We also investigate which type of dialogues our model could benefit from incorporating SRL information? By analyzing the dialogues and our predicted rewritten utterances, we find that the SRL information mainly improves the performance on the dialogues that require information completion. One omitted information is considered as properly completed if the rewritten utterance recovers the omitted words. We find the SRL parser naturally offers important guidance into the selection of omitted words. Examples of rewritten utterances are shown in the Appendix.

Recall that, there is one additional scope option to apply the SRL parser to extract PA structures, i.e., only working on the last utterance that needs to be rewritten. We evaluate this option on our dataset (referred as BERT+Partial-SRL) and results are shown in Table 2. We can see that reducing the SRL scope may slightly hurt the perfor-

²This SRL parser is trained on the CoNLL-2012 dataset.

mance, which we think is due to that larger SRL scope could provide additional guidance for the rewriter model.

5 Conclusions

In this paper, we introduce a novel SRL-guided framework for enhancing dialogue rewriting. For this purpose, we adapted traditional SRL to the conversational scenario by annotating cross-turn annotations on 3,000 dialogues. Experimental results showed that introducing SRL could significantly improve the rewriting performance without adding extra model parameters.

Acknowledgement

This work was supported in part by the Tencent AI Lab Rhino-Bird Gift Fund, the Hong Kong RGC grant ECS 21212419, and the Guangdong Basic and Applied Basic Research Foundation under Key Project 2019B1515120032.

References

- Xavier Carreras and Lluís Màrquez. 2005. [Introduction to the CoNLL-2005 shared task: Semantic role labeling](#). In *CoNLL-2005*, Ann Arbor, Michigan.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. In *Advances in Neural Information Processing Systems*, pages 13042–13054.
- 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 5917–5923.
- 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.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational linguistics*, 28(3):245–288.
- Daniel Khashabi, Mark Sammons, Ben Zhou, Tom Redman, Christos Christodoulopoulos, Vivek Srikrumar, Nick Rizzolo, Lev Ratinov, Guanheng Luo, Quang Do, et al. 2018. Cogcompnlp: Your swiss army knife for nlp. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119.
- Zuchao Li, Shexia He, Hai Zhao, Yiqing Zhang, Zhuosheng Zhang, Xi Zhou, and Xiang Zhou. 2019. Dependency or span, end-to-end uniform semantic role labeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6730–6737.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Hiroki Ouchi, Hiroyuki Shindo, and Yuji Matsumoto. 2018. A span selection model for semantic role labeling. *arXiv preprint arXiv:1810.02245*.
- Martha Palmer, Daniel Gildea, and Nianwen Xue. 2010. Semantic role labeling. *Synthesis Lectures on Human Language Technologies*, 3(1):1–103.
- 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.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. Towards robust linguistic analysis using ontonotes. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152.
- 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 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*.

- Yuanlong Shao, Stephan Gouws, Denny Britz, Anna Goldie, Brian Strope, and Ray Kurzweil. 2017. Generating high-quality and informative conversation responses with sequence-to-sequence models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2210–2219.
- Peng Shi and Jimmy Lin. 2019. Simple bert models for relation extraction and semantic role labeling. *arXiv preprint arXiv:1904.05255*.
- 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.
- Zhixing Tan, Mingxuan Wang, Jun Xie, Yidong Chen, and Xiaodong Shi. 2018. Deep semantic role labeling with self-attention. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- 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. In *ICML Deep Learning Workshop*.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.
- Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. Proactive human-machine conversation with explicit conversation goal. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3794–3804.
- Xiyuan Zhang, Chengxi Li, Dian Yu, Samuel Davidson, and Zhou Yu. 2019. Filling conversation ellipsis for better social dialog understanding. *arXiv preprint arXiv:1911.10776*.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664.

A Conversational SRL Dataset

In this section, we first introduce the dialog set that we annotate on and then discuss more details about the annotation.

A.1 Dialogue Dataset: DuConv

DuConv is a publicly available knowledge-driven dialogue dataset, focusing on the domain of movies and stars. It consists of 30k dialogues with 270k dialogue turns and provides a corresponding knowledge graph (KG).

A.2 Semantic Roles

We follow PropBank (Carreras and Màrquez, 2005), the most widely used standard for annotating predicate-argument structures. It has 32 standard semantic roles. By analyzing the conversation dataset, we adopt 9 core semantic roles in our dialogue SRL:

- Numbered arguments (ARG0-ARG4): Arguments defining verb-specific roles. Their semantics depends on the verb and the verb usage in a sentence, or verb sense. In general, ARG0 stands for the *agent* and ARG1 corresponds to the *patient* or *theme* of the proposition, and these two are the most frequent roles. Numbered arguments reflect either the arguments that are required for the valency of a predicate, or if not required, those that occur with high-frequency in actual usage.
- Adjuncts: General arguments that any verb may take optionally. In PropBank, there are 13 types of adjuncts, while in our dataset we only consider the most frequent four types of adjuncts, i.e., AM-LOC, AM-TMP, AM-PRP and AM-NEG. Specifically, the locative modifiers (AM-LOC) indicate where the action takes place. The temporal arguments (AM-TMP) show when an action takes place, such as 很快 (soon) or 马上 (immediately). Note that, the adverbs of frequency (e.g., 偶尔 (sometimes), 总是 (always)), adverbs of duration (e.g., 过两天 (in two days)) and repetition (e.g., 又 (again)) are also labeled as AM-TMP. Purpose clauses (AM-PRP) are used to show the motivation for an action. Clauses beginning with 为了 (in order to) and 因为 (because) are canonical purpose clauses. AM-NEG is used for elements such as ‘没有’ (not) and ‘绝不’ (no longer).

A.3 Annotation Details

There are two main types of semantic roles: span based (Ouchi et al., 2018; Tan et al., 2018) and de-

	Overall Ratio	Cross-turn Ratio
ARG0	42.1%	22.9%
ARG1	40.2%	16.9%
ARG2	10.1%	30.2%
ARG3	3.0%	24.8%
ARG4	0.3%	41.4%
AM-TMP	3.2%	0.3%
AM-LOC	1.0%	2.1%
AM-PRP	0.1%	4.0%

Table 4: Percent of each type of argument and its cross-turn ratio (shown inside parenthesis).

pendency based (Li et al., 2019). The former involves the start and end boundaries for each component, and the latter only considers the head word in a dependency tree for each component. We follow the span-based form, which has been adopted by most previous work.

Preprocessing For each dialogue session, we first convert it to a paragraph by concatenating each utterance in the dialogue history. We then use Stanford CoreNLP (Manning et al., 2014) for sentence segmentation, tokenizing, and POS-tagging. We identify verbs by POS tag with heuristics to filter out auxiliary verbs.

Labeling instructions We ask five annotators who are familiar with PropBank semantic roles to annotate these dialogue sessions. Following the span-based annotation standard, annotators label the index ranges for each predicate and its arguments. In contrast to the standard sentence-level SRL, conversational SRL aims to additionally address the ellipsis and anaphora problems, which frequently occurred in the dialogue scenario. To this end, the annotators are instructed that a *valid* annotation must satisfy the following criteria: (1) the argument should only appear in the current or previous turns; (2) the argument should not be assigned to a pronoun unless its reference could not be found in previous turns; (3) if the argument is the speaker or listener, it should be explicitly assigned to the special token we used to indicate the speaker (i.e., A or B). (4) in cases when there exist multiple choices for labeling an argument, we select the one that is the closest to the predicate.

Statistics We annotated 3,000 dialogue sessions from DuConv (33,673 predicates in 27,198 utterances). Table 4 analyzes our datasets by listing the percent of each argument type and its cross-turn ratio. We can see that, for all the three datasets, ar-

Example #1	
Utterance 1	十一种孤独 _{ARG0} 作者是谁 (Who is the author of Eleven Kinds & of Loneliness ?)
Utterance 2	理查德耶茨, 对吧 (Richard Yates, right ?)
Utterance 3	这本书讲 _{predicate} 的啥 _{ARG1} (What is this book talking about?)
Gold:	十一种孤独讲的啥 (What is Eleven Kinds of Loneliness talking about?)
BERT:	理查德耶茨讲的啥 (What is Richard Yates talking about?)
BERT + SRL:	十一种孤独讲的啥 (What is Eleven Kinds of Loneliness talking about?)
Example #2	
Utterance 1	济南大学 _{ARG0} (University of Jinan.)
Utterance 2	南京一所著名工科强校 (It is a famous school of engineering in Nanjing.)
Utterance 3	不, 它在 _{predicate} 济南 _{ARG1} (No, it is in Jinan.)
Gold:	不, 济南大学在济南 (No, the University of Jinan is in Jinan.)
BERT:	不, 济南大学在济南大学 (No, the University of Jinan is in Jinan University.)
BERT + SRL:	不, 济南大学在济南 (No, the University of Jinan is in Jinan.)

Table 5: Examples of multi-turn dialogue. The outputs of our SRL model are annotated in the utterances.

guments ARG0, ARG1 and ARG2 count for the major proportion of the arguments. For adjunct-type arguments, AM-TMP and AM-LOC appear more than AM-PRP. It is likely because humans tend to avoid mentioning reasons for simplicity. Besides, the adjunct-type arguments have very low cross-turn ratios. This fits our intuition that humans usually mention the time and location when describing an event or a piece of news.

B Examples of Model Prediction

Table 5 gives some running examples of our model predictions. We can see that with accurate SRL guidance, our model could generate better utterances.