

Towards Topic-Guided Conversational Recommender System

Kun Zhou^{1†}, Yuanhang Zhou^{1†}, Wayne Xin Zhao^{2,3*}, Xiaohe Wang¹ and Ji-Rong Wen^{2,3}

¹School of Information, Renmin University of China

²Gaoling School of Artificial Intelligence, Renmin University of China

³Beijing Key Laboratory of Big Data Management and Analysis Methods

francis_kun_zhou@163.com, {sdzyh002, batmanfly}@gmail.com
{yarkonawang, jrwen}@ruc.edu.cn,

Abstract

Conversational recommender systems (CRS) aim to recommend high-quality items to users through interactive conversations. To develop an effective CRS, the support of high-quality datasets is essential. Existing CRS datasets mainly focus on immediate requests from users, while lack proactive guidance to the recommendation scenario. In this paper, we contribute a new CRS dataset named **TG-ReDial (Recommendation through Topic-Guided Dialog)**. Our dataset has two major features. First, it incorporates topic threads to enforce natural semantic transitions towards the recommendation scenario. Second, it is created in a semi-automatic way, hence human annotation is more reasonable and controllable. Based on TG-ReDial, we present the task of topic-guided conversational recommendation, and propose an effective approach to this task. Extensive experiments have demonstrated the effectiveness of our approach on three sub-tasks, namely topic prediction, item recommendation and response generation. TG-ReDial is available at <https://github.com/RUCAIBox/TG-ReDial>.

1 Introduction

Recently, conversational recommender system (CRS) (Chen et al., 2019; Sun and Zhang, 2018; Li et al., 2018; Zhang et al., 2018b; Liao et al., 2019) has become an emerging research topic, which aims to provide high-quality recommendations to users through natural language conversations. Generally, a CRS is composed of a recommender component and a dialog component, which make suitable recommendation and generate proper response, respectively. To develop an effective CRS, high-quality datasets are crucial to learn the model parameters. Existing CRS datasets roughly fall into two main categories, namely attribute-based user simulation (Sun and Zhang, 2018; Lei et al., 2020; Zhang et al., 2018b) and chit-chat based goal completion (Li et al., 2018; Chen et al., 2019; Liu et al., 2020).

These datasets usually assume that a user has clear, immediate requests when interacting with the system. They lack the proactive guidance (or transitions) from non-recommendation scenarios to the desired recommendation scenario. Indeed, it has become increasingly important that recommendations can be naturally triggered according to conversation context (Tang et al., 2019; Kang et al., 2019). This issue has been explored to some extent by DuRecDial dataset (Liu et al., 2020). DuRecDial has characterized the goal-planning process by constructing a goal sequence. However, it mainly focuses on type switch or coverage for dialog sub-tasks (*e.g.*, non-recommendation, recommendation and question-answering). Explicit semantic transition that leads up to the recommendation has not been well studied or discussed in DuRecDial dataset. Besides, most of existing CRS datasets (Liu et al., 2020; Li et al., 2018) mainly rely on human annotators to create user profiles or generate the conversations. It is difficult to capture rich, complicated cases from real-world applications with a limited number of human annotators, since the generated conversations mainly reflect the characteristics (*e.g.*, interest) of annotators or predefined identities.

To tackle the above problems, we construct a new CRS dataset named **Recommendation through Topic-Guided Dialog (TG-ReDial)**. It consists of 10,000 two-party dialogues between a seeker and a

*Corresponding author. †Equal contribution

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

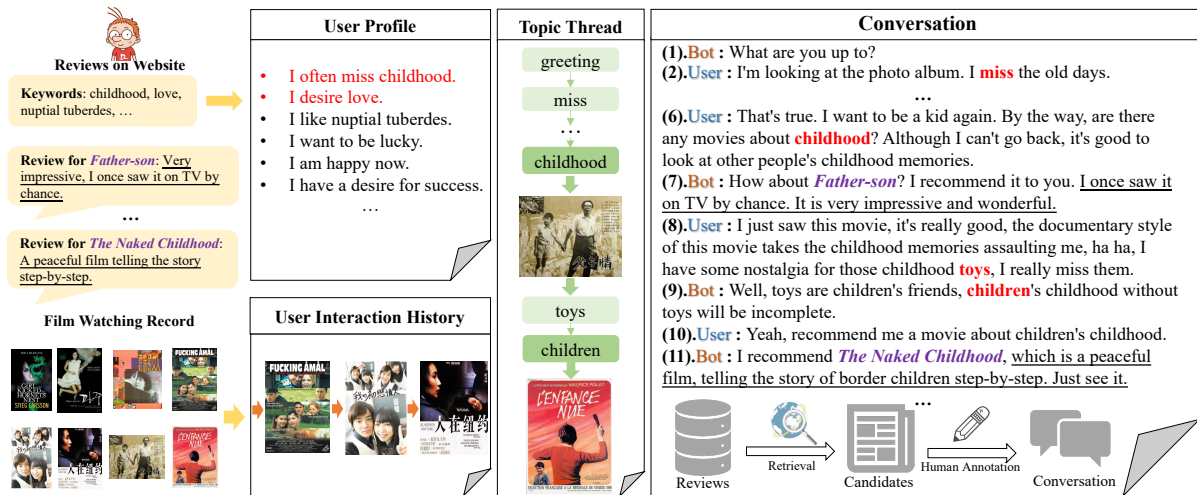


Figure 1: An illustrative example for TG-ReDial dataset. We utilize real data to construct the recommended movies, topic threads, user profiles and utterances. Other user-related information (e.g., historical interaction records) is also available in our dataset.

recommender in the movie domain. There are two new features in our dataset. First, we explicitly create a topic thread to guide the entire content flow for each conversation. Starting with a non-recommendation topic, the topic thread naturally guides the user to the recommendation scenario through a sequence of evolving topics. Our dataset enforces natural transitions towards recommendation through chit-chat conversations. Second, our dataset has been created in a semi-automatic way by involving reasonable and controllable human annotation efforts. The key idea is to align user identities in conversations with real users from a popular movie review website. In this way, the recommended movies, the created topic threads and the recommendation reasons are mined or generated based on real-world data. The major role of the human annotators is to revise, polish or rewrite the conversation data when necessary. Therefore, we do not rely on human annotators to create personalized user profiles as previous studies (Li et al., 2018; Liu et al., 2020), making our conversation data closely resembles real-world cases. Figure 1 presents an illustrative example for our TG-ReDial dataset.

Based on the TG-ReDial dataset, we study a new task of topic-guided conversational recommendation, which can be decomposed into three sub-tasks, namely *item recommendation*, *topic prediction*, and *response generation*. Topic prediction aims to create the topic thread that leads to the final recommendation; item recommendation provides suitable items that meet the user needs; and response generation produces proper reply in natural language. In our approach, the recommender module utilizes both historical interaction and dialog text for deriving accurate user preference, which are modeled by sequential recommendation model SASRec (Kang and McAuley, 2018) and pre-trained language model BERT (Devlin et al., 2019), respectively. The dialog module consists of a topic prediction model and a response generation model. The topic prediction model integrates three kinds of useful data (i.e., historical utterances, historical topics and user profile) to predict the next topic. The response generation model is implemented based on GPT-2 (Radford et al., 2019) to produce responses for guiding users or giving persuasive recommendation. To validate the effectiveness of our approach, we conduct extensive experiments on TG-ReDial dataset to compare our approach with competitive baseline models.

Our main contributions are summarized as follows:

(1) We release a new dataset TG-ReDial for conversational recommender systems. It emphasizes natural topic transitions that leads to the final recommendation. Our dataset is created in a semi-automatic way, and hence human annotation is more reasonable and controllable.

(2) Based on TG-ReDial, we present the task of topic-guided conversational recommendation, consisting of item recommendation, topic prediction and response generation. We further develop an effective solution to leverage multiple kinds of data signals based on Transformer and its variants BERT and

GPT-2.

2 Related Work

2.1 Conversation System

Conversation systems (Shang et al., 2015; Li et al., 2016; Dhingra et al., 2017) study how to generate proper responses given multi-turn contextual utterances. Existing works can be categorized into task-oriented systems (Dhingra et al., 2017; Young et al., 2007) to accomplish specific goals (*e.g.*, book the ticket) and chit-chat systems (Li et al., 2016; Shang et al., 2015; Zhou et al., 2019) to provide general-purpose dialogue. Related to our work, topical information have attracted much research interests in the research community for conversation systems (Xing et al., 2017; Tang et al., 2019; Xu et al., 2020), since it can enhance the semantics of the generated conversation. Early works (Xing et al., 2017; Lian et al., 2019) focused on guiding the conversation topic for the next response, while recent studies (Tang et al., 2019; Xu et al., 2020) started to emphasize the multi-turn topic-guided process in the whole conversation. For example, keyword transition (Tang et al., 2019) and knowledge graph (Xu et al., 2020) are incorporated to improve the topic-guided conversation systems.

2.2 Conversational Recommender System

Conversational recommender system (CRS) (Chen et al., 2019; Sun and Zhang, 2018; Li et al., 2018) aims to provide high-quality recommendation through conversations with users. Generally, it consists of a dialog component to interact with a user and a recommender component to select items for recommendation considering user preference. Early conversational recommender systems (Christakopoulou et al., 2016; Sun and Zhang, 2018; Zhou et al., 2020c) mainly asked questions about user preference over pre-defined slots to make recommendations. Recently, several studies (Li et al., 2018; Chen et al., 2019; Liu et al., 2020) started to interact with user through natural language conversation, emphasizing fluent response generation and precise recommendation. Furthermore, follow-up studies (Chen et al., 2019; Kang et al., 2019; Zhou et al., 2020b) incorporated knowledge graph or reinforcement learning to improve the performance of CRSs with enhanced user models or interaction mechanism.

2.3 Dataset for Conversational Recommendation

To facilitate the study of conversational recommendation, multiple datasets (Li et al., 2018; Kang et al., 2019; Liu et al., 2020; Lei et al., 2020) have been released in recent years. Among them, Facebook-rec (Dodge et al., 2016) and EAR (Lei et al., 2020) are synthetic dialog datasets built by natural language templates based on classic recommendation datasets. ReDial (Li et al., 2018), GoReDial (Kang et al., 2019) and DuRecDial (Liu et al., 2020) are created by human annotation with pre-defined goals, such as *item recommendation* and *goal planning*. These goal-oriented datasets combine the elements of chit-chat and task-oriented (recommendation task) dialogs. Compared with them, our dataset emphasizes the topic-guided process that naturally leads the conversation to the recommendation scenario in CRS.

It is worth noting that DuRecDial dataset (Liu et al., 2020) is similar to TG-ReDial dataset in that it utilizes a goal sequence to guide the conversation. However, the goal sequence is composed by multiple types of tasks (*e.g.*, recommendation, recommendation and question answering). As a comparison, TG-ReDial utilizes topic threads to characterize the evolution of content flow, which is easier to be integrated into open-domain dialogs. Another significant difference is that DuRecDial relies on human annotators to generate user-related data, *e.g.*, user profiles and utterances. While, TG-ReDial mainly mines suitable information from a movie review website, which closely resembles the real cases.

3 Dataset Construction

Following (Sun and Zhang, 2018; Li et al., 2018), we consider a two-party setting for conversational recommendation, in which the person and the chat-bot play the roles of the seeker and the recommender, respectively. We expect that the dialogue between the seeker and the recommender starts from a non-recommendation scenario. The recommender proactively guides the conversation into the target topic, and then makes suitable recommendation considering the seeker’s interests.

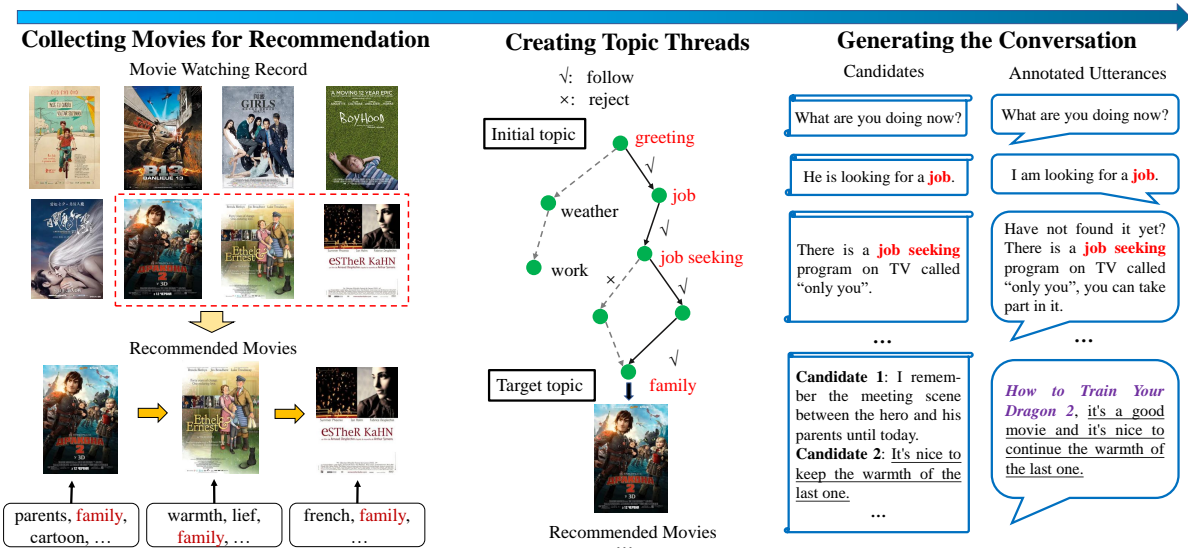


Figure 2: An example of the data collection process of TG-ReDial. We select three movies sharing the same tag of “family” from film watching record for recommendation, then create the topic thread (marked in red font) on ConceptNet, and finally provide high-quality candidate sentences for human annotation.

Different from previous studies, we construct the dataset in a semi-automatic way. We utilize real data records from a popular Chinese movie review website *Douban Movie**. We associate each conversation with a real user from *Douban Movie*, so that the watching records (*likes* and *dislikes*) can be incorporated for recommendation. For a recommended movie, we create an evolving topic thread that leads from previous topic to the target topic of the movie. Finally, the human annotators will generate the proper reply for making recommendation based on user profiles and retrieved high-quality candidates related to the movie. Figure 2 presents an illustrative example for the construction process.

In what follows, we first describe the construction process of TG-ReDial dataset, and then present the detailed statistics about the dataset.

3.1 Collecting Movies for Recommendation

To simulate the real recommendation scenario, we first collect the watching records of real users on *Douban* for recommendation. In order to make recommendation topic-related, we attach each movie with several meaningful tags (*e.g.* genre, director and starring). We keep the original tags of a movie from the *Douban Movie*, and further mine its reviews for extracting high-frequency keywords, and then manually select suitable tags. The number of tags for a movie is set between 1 and 38. The entire watching sequence is split into several coherent subsequence, in which the movies are ensured to share at least a common tag. We remove incoherent subsequences. Each kept subsequence corresponds to a unique conversation, and each user is involved in four to five conversations on average. Given a user, we mark the accept/reject status about a movie according to her/his rating on it (*accept*: ≥ 4 and *reject*: ≤ 2 in a five-star scale). Compared with previous studies (Li et al., 2018; Liu et al., 2020), a major difference is that we reuse existing watching records with rated preferences (covering liked and disliked movies), making the generated conversations closely resemble to the real cases.

3.2 Creating Topic Threads

Given the movies for a conversation, we incorporate topic tags to connect them in an ordered way. The initial topic of each conversation is set to *greeting*, and the target topic is a selected tag of the next movie to be recommended. For creating topic threads, we start from the initial topic and traverse over the commonsense knowledge graph *ConceptNet* (Speer et al., 2017). The shortest topic path identified with

*<https://movie.douban.com/>

| | Total | | Average |
|--------------|---------|------------------------------|---------|
| # Users | 1,482 | # Words per Utterance | 19.0 |
| # Dialogues | 10,000 | # Topics per Dialogue | 7.9 |
| # Utterances | 129,392 | # Movies per Dialogue | 3 |
| # Movies | 33,834 | # Profile Sentences per User | 10 |
| # Topics | 2,571 | # Watched Movies per User | 202.7 |

Table 1: Data statistics of our TG-ReDial dataset.

the *Depth First Search* (DFS) algorithm is considered as a *topic thread*. We repeat the above process in multiple times until all the recommended movies can be connected via topic threads.

In order to enhance the personality of the seeker, following (Zhang et al., 2018a), we generate user profiles to better control the quality of the conversation. We first collect the keywords from the profiles, self-description, and her/his review text from the original website, and then utilize 47 handwritten templates to produce sentences to describe user profiles. With user profiles, we can capture two options in generating topic threads, namely *following* and *rejection*. At each step, the following option will incorporate the current topic in the topic thread, while the rejection option will consider another topic to extend. The choice is made according to whether the topic keyword appears in the extracted profile keywords. For a topic that is not among the keywords, we simply reject to extend this topic with a probability of 0.5. Such a sampling way increases the flexibility and variability of our conversation data.

3.3 Generating the Conversation

After obtaining topic threads and recommended movies, we ask the crowd-sourced workers to complete the conversations. Each conversation starts from chit-chat utterances, evolves according to the topic thread, and provides the recommendation on the target topic. Although the above information (*i.e.*, movie sequence, topic thread and user profile) has highly summarized the sketch of the conversation, it is still difficult to conduct the data annotation with a limited number of human annotators. Inspired by MultiWOZ (Budzianowski et al., 2018), we propose a candidate-driven annotation approach based on an open-domain dialogue corpus Douban (Wu et al., 2017) and the crawled movie reviews, to help generate the topic- and recommendation-related utterances, respectively.

Given a topic thread, we need generate an utterance for each topic in it. Given a topic, we first randomly retrieve 20 utterances containing the topic from Douban corpus. Then we utilize a RNN-based matching model (Lowe et al., 2015) to compute their relevance with the last utterance, and select the most relevant sentence as the candidate utterance. In this step, the role of human annotators is to modify the retrieved candidate to ensure semantic consistency of the entire dialogue.

For the target topic, we need to generate a persuasive reason for making the recommendation. Recall that the user has actually watched the movie and published a review about it. We utilize the target topic as a query to retrieve top three relevant review sentences with extreme embedding similarity (Liu et al., 2016a). The annotators will select among the three candidate sentences, and revise or rewrite it according to conversation context when necessary.

Our crowd-sourced workers are from a specialized data annotation company. Each utterance was assigned an annotator (*labeling*) and an inspector (*checking*). Each annotator is required to carefully read the user profile and browse the detailed information on the original website. To guarantee the quality of human-generated data, we further utilize two automatic metrics to identify low-quality cases for re-labeling. In specific, we compute the Distinct metrics (Li et al., 2016) to filter low-informativeness dialogues with small Distinct values; and we compute the BLEU score (Papineni et al., 2002) between the given candidate with the human-annotated utterance, and then filter the dialogues with very little modification. These bad cases would be relabeled until they pass the automatic evaluation.

3.4 The TG-ReDial Dataset

The detailed statistics of TG-ReDial are shown in Table 1. TG-ReDial consists of 129,392 utterances from 1,482 users. Our dataset is constructed in a topic-guided way, containing more informative sen-

tences. On average, a dialogue has 7.9 topics and an utterance contains 19.0 words, which are larger than the corresponding numbers of existing CRS datasets (Li et al., 2018; Liu et al., 2020; Kang et al., 2019). Furthermore, a user has 10 profile sentences and 202.7 watching records on average.

A major feature of our dataset is that we organize the conversation by topic threads, so that the transitions from chit-chat to recommendation are more natural. Such a dataset is particularly useful to help integrate the recommender component into general-purpose chat-bots, since it is easy to align our topics with open-domain conversations. Moreover, we associate a conversation with a unique user identity, so that it can closely resemble real-world cases. Especially, we can obtain profiles and watching history for the users in a conversation. To our knowledge, most of existing datasets (Li et al., 2018; Liu et al., 2020) mainly focus on cold-start scenario for CRS, while it is also important that CRS can leverage historical interaction data for existing users. Our dataset provides the possibility of training conversational recommendation algorithms with historical interaction data. It is also feasible to study other personalized tasks, since a user is involved in multiple conversations in our dataset.

Note that, in order to protect user privacy, we only sample users with a large number of watching records. For derived user data (e.g., profile or watching records), we perform the anonymized operation and add randomized modification (e.g., removal, replacement or deletion). We also require that the retrieved review sentences have to be written via paraphrasing. Finally, we ask human annotators to manually trace the user identities with corresponding user data in our dataset. We do not include the data from the users that can be identified in the final dataset.

4 Our Approach

In this section, we first formulate the topic-guided conversational recommendation task. Then we introduce our solution to this task.

4.1 Problem Formulation

Given a user u , we assume that she/he is associated with a profile P_u (a set of descriptive sentences related to the topics that u is interested in) and a historical interaction sequence I_u (a chronologically-ordered sequence of items that u has interacted with). Each dialogue is composed by a list of utterances, denoted by $d = \{s_k\}_{k=1}^n$, in which s_k is the utterance at the k -th turn. We consider the CRS in a topic-guided manner, and each utterance s_k is associated with a topic t_k . When t_k is a target topic, the system will trigger the recommendation of item i_k with the persuasive reason.

Based on these notations, the task of topic-guided conversational recommendation is defined as: given the user profile P_u , user interaction sequence I_u , historical utterances $\{s_1, \dots, s_{k-1}\}$ and corresponding topic sequence $\{t_1, \dots, t_{k-1}\}$, we aim to (1) predict the next topic t_k to reach the target topic, or (2) recommend the movie i_k , and finally (3) produce a proper response s_k about the topic or with persuasive reason. The three sub-tasks are referred to topic prediction, item recommendation and response generation.

4.2 Recommendation Module

The recommendation module aims to predict the item that a user likes given the conversation context. The key point is how to derive an effective user presentation for recommendation. We consider two kinds of data signals for this task. Specially, we utilize the pre-trained language model BERT (Devlin et al., 2019) to encode the historical utterances $\{s_1, \dots, s_{k-1}\}$, and a self-attentive sequential recommendation model SASRec (Kang and McAuley, 2018) to encode the user interaction sequence I_u .

The representation v_u of user u is obtained as follows:

$$v_u = \text{MLP}([v_u^{(1)}; v_u^{(2)}]), \quad (1)$$

where $v_u^{(1)}$ (obtained from BERT) and $v_u^{(2)}$ (obtained from SASRec) are the embeddings to represent the historical utterances and interaction sequence, respectively. Given the user representation, we can compute the probability that recommends an item i from the item set to a user u :

$$P_{rec}(i) = \text{softmax}(e_i^\top \cdot v_u) \quad (2)$$

where e_i is the learned item embedding for item i . We utilize Equation 2 to rank all the items and select the item with the largest probability for recommendation.

4.3 Dialog Module

The dialog module aims to generate proper responses to the user (seeker) for topic guidance or item recommendation. We achieve the two purposes with specific models.

Topic Prediction Model. It predicts the next topic that guides user u towards the target topic. We mainly utilize text data for topic prediction, and implement three different BERT-based encoders, namely conversation-BERT, topic-BERT and profile-BERT, for encoding the historical utterances, historical topic sequence and user profile, respectively. For each BERT variant, we simply concatenate all the available text data and the target topic (with [SEP] tokens for separating). The incorporation of target topic is to enhance the topical semantics. Based on the obtained representations, we compute the probability of a topic t as the next topic by:

$$P_{topic}(t) = \text{softmax}(e_t^\top \cdot \text{MLP}([\mathbf{r}^{(1)}; \mathbf{r}^{(2)}; \mathbf{r}^{(3)}])), \quad (3)$$

where e_t is the learned embedding for topic t , $\mathbf{r}^{(1)}$, $\mathbf{r}^{(2)}$ and $\mathbf{r}^{(3)}$ are the embeddings of historical utterances, topics and user profile obtained from conversation-BERT, topic-BERT and profile-BERT, respectively. We utilize Eq. 3 to rank all the topics and select the topic with the largest probability.

Response Generation Model. It aims to generate proper responses for topic-guided conversations. We leverage the pre-trained text generation model GPT-2 (Radford et al., 2019) for response generation. GPT-2 utilizes a stacked of masked multi-head self-attention layers trained on massive web-text data by the generic language model (Radford et al., 2019; Devlin et al., 2019). We consider two cases in this model. For non-recommendation case, we generate the response conditioned on the predicted topic, and concatenate t_k with the historical utterances $\{s_1, \dots, s_{k-1}\}$ (separated by [SEP] tokens). For recommendation case, we generate the persuasive reason conditioned on the recommended item, and concatenate the recommended movie i_k with the historical utterances $\{s_1, \dots, s_{k-1}\}$. For both cases, we can unify the input as a long sequence, which will be encoded and fed into GPT-2 for decoding.

5 Experiments

We evaluate the proposed approach on TG-ReDial dataset, which is split into training, validation and test sets using a ratio of 8:1:1. For each conversation, we start from the first utterance, and generate reply utterances or recommendations in turn by our model. We perform the evaluation on the three sub-tasks, namely item recommendation, topic prediction and response generation.

5.1 Evaluation on Item Recommendation

In this subsection, we conduct a series of experiments on the effectiveness of our proposed model for the recommendation task. Following (Kang and McAuley, 2018; Liu et al., 2016b), we adopt $\text{NDCG}@k$ and $\text{MRR}@k$ ($k = 10, 50$) as evaluation metrics for ranking all the possible items.

Baselines. We consider the following baselines for performance comparisons: (1) *Popularity* ranks items according to popularity measured by the number of interactions. (2) *ReDial* (Li et al., 2018) is proposed specially for CRSs by utilizing an auto-encoder for recommendation. (3) *KBRD* (Chen et al., 2019) is the state-of-the-art CRS model using knowledge graphs to enhance the semantics of contextual items or entities for recommendation. (4) *GRU4Rec* (Liu et al., 2016b) applies GRU to model user interaction history without using conversation data. (5) *SASRec* (Kang and McAuley, 2018) adopts the Transformer architecture to encode user interaction history without using conversation data. (6) *TextCNN* (Kim, 2014) adopts a CNN-based model to extract textual features from contextual utterances for recommendation. (7) *BERT* (Devlin et al., 2019) is a pre-training language model that directly encodes the concatenated historical utterances.

| Models | NDCG@10 | NDCG@50 | MRR@10 | MRR@50 |
|-------------|---------------|---------------|---------------|---------------|
| Popularity | 0.0015 | 0.0036 | 0.0011 | 0.0015 |
| ReDial | 0.0009 | 0.0029 | 0.0005 | 0.0009 |
| KBRD | 0.0064 | 0.0111 | 0.0040 | 0.0049 |
| GRU4Rec | 0.0028 | 0.0062 | 0.0014 | 0.0020 |
| SASRec | 0.0092 | 0.0179 | 0.0050 | 0.0068 |
| TextCNN | 0.0144 | 0.0215 | 0.0119 | 0.0133 |
| BERT | 0.0246 | 0.0439 | 0.0182 | 0.0221 |
| Ours | 0.0348 | 0.0527 | 0.0240 | 0.0277 |

Table 2: Results on item recommendation task.

| Models | Hit@1 | Hit@3 | Hit@5 |
|------------------------|---------------|---------------|---------------|
| PMI | 0.0349 | 0.0927 | 0.1290 |
| MGCG | 0.6098 | 0.8128 | 0.8294 |
| Conversation-BERT | 0.6114 | 0.8189 | 0.8341 |
| Topic-BERT | 0.6155 | 0.8275 | 0.8405 |
| Profile-BERT | 0.4986 | 0.8205 | 0.8344 |
| Ours <i>w/o</i> target | 0.4420 | 0.5923 | 0.6374 |
| Ours | 0.6231 | 0.8370 | 0.8497 |

Table 3: Results on topic prediction task.

Result and Analysis. Table 2 presents the performance of different methods on recommendation task. As we can see, Popularity performs better than ReDial but worse than KBRD. ReDial and KBRD utilize the items in historical utterances for recommendation. Besides, KBRD incorporates external knowledge graph to enhance the representations of items. Second, SASRec outperforms GRU4Rec and the two CRS models (*i.e.*, KBRD and ReDial). It indicates that self-attentive architecture is particularly suitable for modeling the interaction history. Furthermore, text-based TextCNN (*i.e.*, TextCNN and BERT) perform better than other baselines, which indicates that it is useful to leverage the historical utterances for recommendation. Among the two text-based models, BERT outperforms TextCNN, since it adopts more powerful architecture trained with large-scale data. Finally, our proposed model outperforms all the baselines significantly. Our model is able to utilize both historical utterances and interaction sequence, combining the merits of BERT and SASRec.

5.2 Evaluation on Topic Prediction

We continue to evaluate the performance of our approach on the topic prediction task. Following (Tang et al., 2019), we adopt Hit@ k ($k = 1, 3, 5$) as evaluation metrics for ranking all the possible topics.

Baselines. We consider the following baselines for performance comparison: (1) *PMI* measures the point-wise mutual information with the last topic for ranking. (2) *MGCG* (Liu et al., 2020) is a recently proposed CRS model based on multi-type GRUs (with a special GRU to encode user profiles). (3) *Conversation/Topic/Profile-BERT* utilizes the conversation/topic/profile-BERT to encode historical utterances/topics/user profiles for predicting the next topic, respectively. (4) *Ours w/o target* is the ablation model of our proposed model by removing the target topic from input.

Result and Analysis. Table 3 presents the performance of different methods on topic prediction task. As we can see, PMI does not perform well, since it cannot consider the target topic. Second, Conversation/Topic-BERT performs better than MGCG. It indicates that the pre-trained language model BERT is particularly suitable to capture topical semantics. Among the BERT-based models, Profile-BERT performs worse, since historical utterances and topics are more important to consider in this task. Furthermore, our model outperforms all the baselines, since it jointly utilizes historical utterances, topics and user profiles encoded by different BERT models. Finally, after removing target topic from the input of BERT models, the performance decreases significantly, indicating that the target topic is important.

| Models | PPL | BLEU-1/2/3 | Distinct-1/2 | Relevance | Fluency | Informativeness |
|-------------|--------------|---------------------------|---------------------|--------------|--------------|-----------------|
| ReDial | 81.614 | 0.177/0.028/0.006 | 0.004/0.025 | 0.660 | 0.400 | 0.560 |
| KBRD | 28.022 | 0.221/0.028/0.009 | 0.004/0.008 | 0.860 | 1.060 | 1.300 |
| Transformer | 32.856 | 0.287/0.071/0.032 | 0.013/0.083 | 1.400 | 1.360 | 1.440 |
| GPT-2 | 13.383 | 0.279/0.066/ 0.032 | 0.017/ 0.094 | 1.460 | 1.440 | 1.500 |
| Ours | 7.223 | 0.280/0.065/0.031 | 0.021/0.094 | 1.520 | 1.360 | 1.640 |

Table 4: Results on response generation task.

5.3 Evaluation on Response Generation

Finally, we evaluate the performance of our approach on the response generation task. Following (Chen et al., 2019; Qiu et al., 2019; Tao et al., 2018a), we adopt perplexity (PPL) and BLEU-1,2,3 for evaluating the relevance between generated response with the ground truth, and adopt Distinct-1,2 for evaluating the informativeness of the generated utterances. Furthermore, we invite human annotators to score the *Relevance*, *Fluency* and *Informativeness* of the generated results with the rating range of $[0, 2]$.

Baselines. We consider the following baselines for performance comparison: (1) *ReDial* (Li et al., 2018) adopts the hierarchical RNN for response generation. (2) *KBRD* (Chen et al., 2019) applies Transformer with enhanced modeling of word weight based on knowledge graphs. (3) *Transformer* (Vaswani et al., 2017) applies a Transformer-based encoder-decoder framework to generate proper responses. (4) *GPT-2* (Radford et al., 2019) is a pre-training text generation model and fine-tuned on TG-ReDial dataset.

Result and Analysis. Table 4 presents the performance of different methods on response generation task. The first observation is that ReDial does not perform well on our dataset. A major reason is that ReDial utilizes a hierarchical RNN for response generation, which is not suitable to encode long utterances (recall that utterances in our dataset are longer and more informative). Second, Transformer is better than KBRD in most of metrics, since KBRD utilizes knowledge graph information to promote the predictive probability of entities and items, which may have an adverse effect on text generation. Furthermore, Transformer, GPT-2 and our model give similar BLEU scores. For PPL, Distinct and human evaluation, Transformer achieves the worst results, while our model achieves very good results. Indeed, BLEU may not be suitable for evaluating CRSs (Tao et al., 2018b), because it is easier to be affected by meaningless words such as stopwords. Finally, our model outperforms all the baselines in most cases, since it can utilize the predicted topic or item to enhance the quality of the generated text.

6 Conclusion

We introduced a high-quality dataset TG-ReDial for conversational recommender systems, which was constructed by human annotation based on real-world user data. Based on TG-ReDial, we presented the task of topic-guided conversational recommendation and a solution to this task. Extensive experiments have demonstrated the effectiveness of the proposed approach on three sub-tasks.

Currently, the potential of TG-ReDial dataset has not been fully explored. It can be useful as a testbed for more tasks, such as personalized chit-chat (Zhang et al., 2018a), target-guided conversation (Tang et al., 2019) and sequential recommendation (Zhou et al., 2020a). As future work, we will investigate the study of these tasks on TG-ReDial dataset. Besides, we will also consider how to construct more effective approaches to topic-guided conversational recommendation.

Acknowledgement

This work was partially supported by the National Natural Science Foundation of China under Grant No. 61872369 and 61832017, Beijing Academy of Artificial Intelligence (BAAI) under Grant No. BAAI2020ZJ0301, and Beijing Outstanding Young Scientist Program under Grant No. BJJWZYJH012019100020098, the Fundamental Research Funds for the Central Universities, the Research Funds of Renmin University of China under Grant No.18XNLG22 and 19XNQ047. Xin Zhao is the corresponding author.

References

- Pawel Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz - A large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 5016–5026.
- Qibin Chen, Junyang Lin, Yichang Zhang, Ming Ding, Yukuo Cen, Hongxia Yang, and Jie Tang. 2019. Towards knowledge-based recommender dialog system. 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 2019, Hong Kong, China, November 3-7, 2019*, pages 1803–1813.
- Konstantina Christakopoulou, Filip Radlinski, and Katja Hofmann. 2016. Towards conversational recommender systems. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pages 815–824.
- 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.
- Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. 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 484–495.
- Jesse Dodge, Andreea Gane, Xiang Zhang, Antoine Bordes, Sumit Chopra, Alexander H. Miller, Arthur Szlam, and Jason Weston. 2016. Evaluating prerequisite qualities for learning end-to-end dialog systems. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Wang-Cheng Kang and Julian J. McAuley. 2018. Self-attentive sequential recommendation. In *IEEE International Conference on Data Mining, ICDM 2018, Singapore, November 17-20, 2018*, pages 197–206.
- Dongyeop Kang, Anusha Balakrishnan, Pararth Shah, Paul Crook, Y-Lan Boureau, and Jason Weston. 2019. Recommendation as a communication game: Self-supervised bot-play for goal-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 2019, Hong Kong, China, November 3-7, 2019*, pages 1951–1961.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1746–1751.
- Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020. Estimation-action-reflection: Towards deep interaction between conversational and recommender systems. In *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020*, pages 304–312.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 110–119.
- Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards deep conversational recommendations. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 9748–9758.
- Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to select knowledge for response generation in dialog systems. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5081–5087.
- Lizi Liao, Ryuichi Takanobu, Yunshan Ma, Xun Yang, Minlie Huang, and Tat-Seng Chua. 2019. Deep conversational recommender in travel. *CoRR*, abs/1907.00710.

- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016a. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2122–2132.
- Qiang Liu, Shu Wu, Diyi Wang, Zhaokang Li, and Liang Wang. 2016b. Context-aware sequential recommendation. In *IEEE 16th International Conference on Data Mining, ICDM 2016, December 12-15, 2016, Barcelona, Spain*, pages 1053–1058.
- Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Towards conversational recommendation over multi-type dialogs. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1036–1049.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *Proceedings of the SIGDIAL 2015 Conference, The 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 2-4 September 2015, Prague, Czech Republic*, pages 285–294.
- 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, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318.
- Lisong Qiu, Juntao Li, Wei Bi, Dongyan Zhao, and Rui Yan. 2019. Are training samples correlated? learning to generate dialogue responses with multiple references. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3826–3835.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1577–1586.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4444–4451.
- Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, pages 235–244.
- Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric P. Xing, and Zhiting Hu. 2019. Target-guided open-domain conversation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5624–5634.
- Chongyang Tao, Shen Gao, Mingyue Shang, Wei Wu, Dongyan Zhao, and Rui Yan. 2018a. Get the point of my utterance! learning towards effective responses with multi-head attention mechanism. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4418–4424.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018b. RUBER: an unsupervised method for automatic evaluation of open-domain dialog systems. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 722–729.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 496–505.

- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3351–3357.
- Jun Xu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Conversational graph grounded policy learning for open-domain conversation generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1835–1845.
- Steve J. Young, Jost Schatzmann, Karl Weilhammer, and Hui Ye. 2007. The hidden information state approach to dialog management. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2007, Honolulu, Hawaii, USA, April 15-20, 2007*, pages 149–152.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018a. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2204–2213.
- Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018b. Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM 2018, Torino, Italy, October 22-26, 2018*, pages 177–186.
- Kun Zhou, Kai Zhang, Yu Wu, Shujie Liu, and Jingsong Yu. 2019. Unsupervised context rewriting for open domain conversation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 1834–1844.
- Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020a. S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, pages 1893–1902.
- Kun Zhou, Wayne Xin Zhao, Shuqing Bian, Yuanhang Zhou, Ji-Rong Wen, and Jingsong Yu. 2020b. Improving conversational recommender systems via knowledge graph based semantic fusion. In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pages 1006–1014.
- Kun Zhou, Wayne Xin Zhao, Hui Wang, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020c. Leveraging historical interaction data for improving conversational recommender system. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, pages 2349–2352.