

# Towards Conversational Recommendation over Multi-Type Dialogs

Zeming Liu<sup>1\*</sup>, Haifeng Wang<sup>2</sup>, Zheng-Yu Niu<sup>2</sup>, Hua Wu<sup>2</sup>, Wanxiang Che<sup>1†</sup>, Ting Liu<sup>1</sup>

<sup>1</sup>Research Center for Social Computing and Information Retrieval,  
Harbin Institute of Technology, Harbin, China

<sup>2</sup>Baidu Inc., Beijing, China

{zmliu, car, tliu}@ir.hit.edu.cn

{wanghaifeng, niuzhengyu, wu.hua}@baidu.com

## Abstract

We focus on the study of conversational recommendation in the context of multi-type dialogs, where the bots can proactively and naturally lead a conversation from a non-recommendation dialog (e.g., QA) to a recommendation dialog, taking into account user’s interests and feedback. To facilitate the study of this task, we create a human-to-human Chinese dialog dataset *DuRecDial* (about 10k dialogs, 156k utterances), which contains multiple sequential dialogs for every pair of a recommendation seeker (user) and a recommender (bot). In each dialog, the recommender proactively leads a multi-type dialog to approach recommendation targets and then makes multiple recommendations with rich interaction behavior. This dataset allows us to systematically investigate different parts of the overall problem, e.g., how to naturally lead a dialog, how to interact with users for recommendation. Finally we establish baseline results on *DuRecDial* for future studies.<sup>1</sup>

## 1 Introduction

In recent years, there has been a significant increase in the work of conversational recommendation due to the rise of voice-based bots (Christakopoulou et al., 2016; Li et al., 2018; Reschke et al., 2013; Warnestal, 2005). They focus on how to provide high-quality recommendations through dialog-based interactions with users. These work fall into two categories: (1) task-oriented dialog-modeling approaches (Christakopoulou et al., 2016; Sun and Zhang, 2018; Warnestal, 2005); (2) non-task dialog-modeling approaches with more free-form interactions (Kang et al., 2019; Li et al., 2018).

Almost all these work focus on a single type of dialogs, either task oriented dialogs for recommendation, or recommendation oriented open-domain conversation. Moreover, they assume that both sides in the dialog (especially the user) are aware of the conversational goal from the beginning.

In many real-world applications, there are multiple dialog types in human-bot conversations (called *multi-type dialogs*), such as chit-chat, task oriented dialogs, recommendation dialogs, and even question answering (Ram et al., 2018; Wang et al., 2014; Zhou et al., 2018b). Therefore it is crucial to study how to proactively and naturally make conversational recommendation by the bots in the context of multi-type human-bot communication. For example, the bots could proactively make recommendations after question answering or a task dialog to improve user experience, or it could lead a dialog from chit-chat to approach a given product as commercial advertisement. However, to our knowledge, there is less previous work on this problem.

To address this challenge, we present a novel task, conversational recommendation over multi-type dialogs, where we want the bot to proactively and naturally lead a conversation from a non-recommendation dialog to a recommendation dialog. For example, in Figure 1, given a starting dialog such as question answering, the bot can take into account user’s interests to determine a recommendation target (the movie <The message>) as a long-term goal, and then drives the conversation in a natural way by following short-term goals, and completes each goal in the end. Here each goal specifies a dialog type and a dialog topic. Our task setting is different from previous work (Christakopoulou et al., 2016; Li et al., 2018). First, the overall dialog in our task contains multiple dialog types, instead of a single dialog type as done in previous work. Second, we emphasize the initiative of the recommender, i.e. the bot proactively plans

\* This work was done at Baidu.

† Corresponding author: Wanxiang Che.

<sup>1</sup>Dataset and codes are publicly available at <https://github.com/PaddlePaddle/models/tree/develop/PaddleNLP/Research/ACL2020-DuRecDial>.

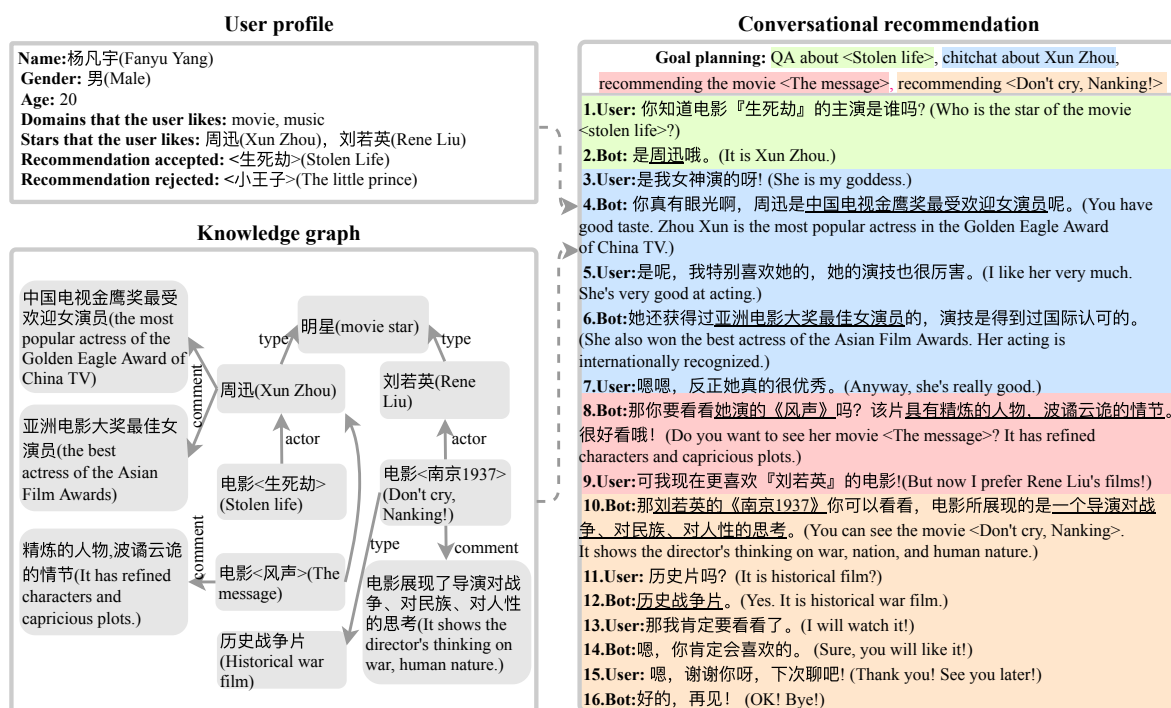


Figure 1: A sample of conversational recommendation over multi-type dialogs. The whole dialog is grounded on knowledge graph and a goal sequence, while the goal sequence is planned by the bot with consideration of user’s interests and topic transition naturalness. Each goal specifies a dialog type and a dialog topic (an entity). We use different colors to indicate different goals and use underline to indicate knowledge texts.

a goal sequence to lead the dialog, and the goals are unknown to the users. When we address this task, we will encounter two difficulties: (1) how to proactively and naturally lead a conversation to approach the recommendation target, (2) how to iterate upon initial recommendation with the user.

To facilitate the study of this task, we create a human-to-human recommendation oriented multi-type Chinese dialog dataset at Baidu (**DuRecDial**). In *DuRecDial*, every dialog contains multi-type dialogs with natural topic transitions, which corresponds to the first difficulty. Moreover, there are rich interaction variability for recommendation, corresponding to the second difficulty. In addition, each seeker has an explicit profile for the modeling of personalized recommendation, and multiple dialogs with the recommender to mimic real-world application scenarios.

To address this task, inspired by the work of Xu et al. (2020), we present a multi-goal driven conversation generation framework (**MGCG**) to handle multi-type dialogs simultaneously, such as QA/chitchat/recommendation/task etc.. It consists of a goal planning module and a goal-guided responding module. The goal-planning module can conduct dialog management to control the dialog

flow, which determines a recommendation target as the final goal with consideration of user’s interests and online feedback, and plans appropriate short-term goals for natural topic transitions. To our knowledge, this goal-driven dialog policy mechanism for multi-type dialog modeling is not studied in previous work. The responding module produces responses for completion of each goal, e.g., chatting about a topic or making a recommendation to the user. We conduct an empirical study of this framework on *DuRecDial*.

This work makes the following contributions:

- We identify the task of conversational recommendation over multi-type dialogs.
- To facilitate the study of this task, we create a novel dialog dataset *DuRecDial*, with rich variability of dialog types and domains as shown in Table 1.
- We propose a conversation generation framework with a novel mixed-goal driven dialog policy mechanism.

Datasets↓ Metrics→	#Dialogs	#Utterances	Dialog types	Domains	User profile
Facebook_Rec (Dodge et al., 2016)	1M	6M	Rec.	Movie	No
REDIAL (Li et al., 2018)	10k	163k	Rec., chitchat	Movie	No
GoRecDial (Kang et al., 2019)	9k	170k	Rec.	Movie	Yes
OpenDialKG (Moon et al., 2019)	12k	143k	Rec.	Movie, book	No
CMU DoG (Zhou et al., 2018a)	4k	130k	Chitchat	Movie	No
IIT DoG (Moghe et al., 2018)	9k	90k	Chitchat	Movie	No
Wizard-of-wiki (Dinan et al., 2019)	22k	202k	Chitchat	1365 Wikipedia topics	No
OpenDialKG (Moon et al., 2019)	3k	38k	Chitchat	Sports, music	No
DuConv (Wu et al., 2019)	29k	270k	Chitchat	Movie	No
KdConv (Zhou et al., 2020)	4.5k	86k	Chitchat	Movie, music, travel	No
DuRecDial	10.2k	156k	Rec., chitchat, QA, task	Movie, music, movie star, food, restaurant, news, weather	Yes

Table 1: Comparison of our dataset *DuRecDial* to recommendation dialog datasets and knowledge grounded dialog datasets. “Rec.” stands for recommendation.

## 2 Related Work

### Datasets for Conversational Recommendation

To facilitate the study of conversational recommendation, multiple datasets have been created in previous work, as shown in Table 1. The first recommendation dialog dataset is released by Dodge et al. (2016), which is a synthetic dialog dataset built with the use of the classic MovieLens ratings dataset and natural language templates. Li et al. (2018) creates a human-to-human multi-turn recommendation dialog dataset, which combines the elements of social chitchat and recommendation dialogs. Kang et al. (2019) provides a recommendation dialogue dataset with clear goals, and Moon et al. (2019) collects a parallel Dialog $\leftrightarrow$ KG corpus for recommendation. Compared with them, our dataset contains multiple dialog types, multi-domain use cases, and rich interaction variability.

**Datasets for Knowledge Grounded Conversation** As shown in Table 1, CMU DoG (Zhou et al., 2018a) explores two scenarios for Wikipedia-article grounded dialogs: only one participant has access to the document, or both have. IIT DoG (Moghe et al., 2018) is another dialog dataset for movie chats, wherein only one participant has access to background knowledge, such as IMDB’s facts/plots, or Reddit’s comments. Dinan et al. (2019) creates a multi-domain multi-turn conversations grounded on Wikipedia articles. OpenDialKG (Moon et al., 2019) provides a chit-chat dataset between two agents, aimed at the modeling of dialog logic by walking over knowledge graph-Freebase. Wu et al. (2019) provides a Chinese dialog dataset-DuConv, where one participant can proactively lead the conversation with an explicit goal. KdConv (Zhou et al., 2020) is a Chinese dialog dataset,

where each dialog contains in-depth discussions on multiple topics. In comparison with them, our dataset contains multiple dialog types, clear goals to achieve during each conversation, and user profiles for personalized conversation.

### Models for Conversational Recommendation

Previous work on conversational recommender systems fall into two categories: (1) task-oriented dialog-modeling approaches in which the systems ask questions about user preference over predefined slots to select items for recommendation (Christakopoulou et al., 2018, 2016; Lee et al., 2018; Reschke et al., 2013; Sun and Zhang, 2018; Warnestal, 2005; Zhang et al., 2018b); (2) non-task dialog-modeling approaches in which the models learn dialog strategies from the dataset without predefined task slots and then make recommendations without slot filling (Chen et al., 2019; Kang et al., 2019; Li et al., 2018; Moon et al., 2019; Zhou et al., 2018a). Our work is more close to the second category, and differs from them in that we conduct multi-goal planning to make proactive conversational recommendation over multi-type dialogs.

### Goal Driven Open-domain Conversation Generation

Recently, imposing goals on open-domain conversation generation models having attracted lots of research interests (Moon et al., 2019; Li et al., 2018; Tang et al., 2019b; Wu et al., 2019) since it provides more controllability to conversation generation, and enables many practical applications, e.g., recommendation of engaging entities. However, these models can just produce a dialog towards a single goal, instead of a goal sequence as done in this work. We notice that the model by Xu et al. (2020) can conduct multi-goal planning for conversation generation. But their goals are limited

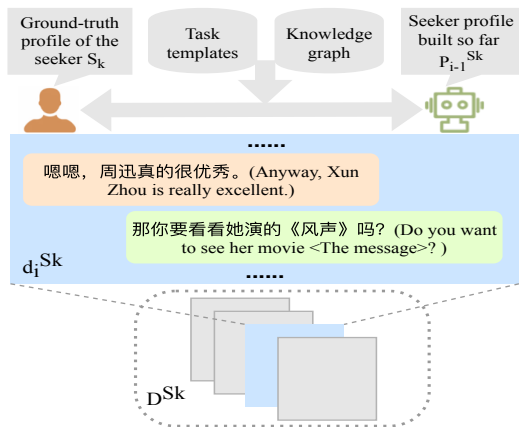


Figure 2: We collect multiple sequential dialogs  $\{d_i^{S_k}\}$  for each seeker  $s_k$ . For annotation of every dialog, the recommender makes personalized recommendations according to task templates, knowledge graph and the seeker profile built so far. The seeker must accept/reject the recommendations.

to in-depth chitchat about related topics, while our goals are not limited to in-depth chitchat.

### 3 Dataset Collection<sup>2</sup>

#### 3.1 Task Design

We define one person in the dialog as the recommendation seeker (the role of users) and the other as the recommender (the role of bots). We ask the recommender to proactively lead the dialog and then make recommendations with consideration of the seeker’s interests, instead of the seeker to ask for recommendation from the recommender. Figure 2 shows our task design. The data collection consists of three steps: (1) collection of seeker profiles and knowledge graph; (2) collection of task templates; (3) annotation of dialog data.<sup>3</sup> Next we will provide details of each step.

**Explicit seeker profiles** Each seeker is equipped with an explicit unique profile (a ground-truth profile), which contains the information of name, gender, age, residence city, occupation, and his/her preference on domains and entities. We automatically generate the ground-truth profile for each seeker, which is known to the seeker, and unknown to the recommender. We ask that the utterances of each seeker should be consistent with his/her profile. We expect that this setting could encourage the seeker to clearly and self-consistently explain what he/she likes/dislikes. In addition, the

<sup>2</sup>Please see Appendix 1. for more details.

<sup>3</sup>We have a strict data quality control process, please see Appendix 1.2 for more details.

recommender can acquire seeker profile information only through dialogs with the seekers.

**Knowledge graph** Inspired by the work of document grounded conversation (Ghazvininejad et al., 2018; Moghe et al., 2018), we provide a knowledge graph to support the annotation of more informative dialogs. We build them by crawling data from Baidu Wiki and Douban websites. Table 3 presents the statistics of this knowledge graph.

**Multi-type dialogs for multiple domains** We expect that the dialog between the two task-workers starts from a non-recommendation scenario, e.g., question answering or social chitchat, and the recommender should proactively and naturally guide the dialog to a recommendation target (an entity). The targets usually fall into the seeker’s interests, e.g., the movies of the star that the seeker likes.

Moreover, to be close to the setting in practical applications, we ask each seeker to conduct multiple sequential dialogs with the recommender. In the first dialog, the recommender asks questions about seeker profile. Then in each of the remaining dialogs, the recommender makes recommendations based on the seeker’s preferences collected so far, and then the seeker profile is automatically updated at the end of each dialog. We ask that the change of seeker profile should be reflected in later dialogs. The difference between these dialogs lies in sub-dialog types and recommended entities.

**Rich variability of interaction** How to iterate upon initial recommendation plays a key role in the interaction procedure for recommendation. To provide better supervision for this capability, we expect that the task workers can introduce diverse interaction behaviors in dialogs to better mimic the decision-making process of the seeker. For example, the seeker may reject the initial recommendation, or mention a new topic, or ask a question about an entity, or simply accept the recommendation. The recommender is required to respond appropriately and follow the seeker’s new topic.

**Task templates as annotation guidance** Due to the complexity of our task design, it is very hard to conduct data annotation with only high-level instructions mentioned above. Inspired by the work of MultiWOZ (Budzianowski et al., 2018), we provide a task template for each dialog to be annotated, which guides the workers to annotate in the way we expect them to be. As shown in Table 2, each template contains the following information: (1) a goal sequence, where each goal consists of two

Goals	Goal description
Goal1: QA (dialog type) about the movie <Stolen life> (dialog topic)	The seeker takes the initiative, and asks for the information about the movie <Stolen life>; the recommender replies according to the given knowledge graph; finally the seeker provides feedback.
Goal2: chitchat about the movie star Xun Zhou	The recommender proactively changes the topic to movie star Xun Zhou as a short-term goal, and conducts an in-depth conversation;
Goal3: Recommendation of the movie <The message>	The recommender proactively changes the topic from movie star to related movie <The message>, and recommend it with movie comments, and the seeker changes the topic to Rene Liu’s movies;
Goal4: Recommendation of the movie <Don’t cry, Nanking!>	The recommender proactively recommends Rene Liu’s movie <Don’t cry, Nanking!> with movie comments. The seeker tries to ask questions about this movie, and the recommender should reply with related knowledge. Finally the user accepts the recommended movie.

Table 2: One of our task templates that is used to guide the workers to annotate the dialog in Figure 1. We require that the recommendation target (the long-term goal) is consistent with the user’s interests and the topics mentioned by the user, and short-term goals provide natural topic transitions to approach the long-term goal.

elements, a dialog type and a dialog topic, corresponding to a sub-dialog. (2) a detailed description about each goal. We create these templates by (1) first automatically enumerating appropriate goal sequences that are consistent with the seeker’s interests and have natural topic transitions and (2) then generating goal descriptions with the use of some rules and human annotation.

### 3.2 Data Collection

To obtain this data, we develop an interface and a pairing mechanism. We pair up task workers and give each of them a role of seeker or recommender. Then the two workers conduct data annotation with the help of task templates, seeker profiles and knowledge graph. In addition, we ask that the goals in templates must be tagged in every dialog.

**Data structure** We organize the dataset of *DuRecDial* according to seeker IDs. In *DuRecDial*, there are multiple seekers (each with a different profile) and only one recommender. Each seeker  $s_k$  has multiple dialogs  $\{d_i^{s_k}\}_i$  with the recommender. For each dialog  $d_i^{s_k}$ , we provide a knowledge graph and a goal sequence for data annotation, and a seeker profile updated with this dialog.

**Data statistics** Table 3 provides statistics of knowledge graph and *DuRecDial*, indicating rich

#Domains	7
Knowledge#Entities	21,837
graph #Attributes	454
#Triples	222,198
#Dialogs	10,190
DuRecDial #Sub-dialogs for QA/Rec/task/chitchat	6,722/8,756/3,234/10,190
#Utterances	155,477
#Seekers	1362
#Entities	11,162/8,692/2,470/
recom-mended/accepted/rejected	

Table 3: Statistics of knowledge graph and *DuRecDial*.

variability of dialog types and domains.

**Data quality** We conduct human evaluations for data quality. A dialog will be rated "1" if it follows the instruction in task templates and the utterances are fluent and grammatical, otherwise "0". Then we ask three persons to judge the quality of 200 randomly sampled dialogs. Finally we obtain an average score of 0.89 on this evaluation set.

## 4 Our Approach

### 4.1 Problem Definition and Framework

#### Overview

**Problem definition** Let  $D^{s_k} = \{d_i^{s_k}\}_{i=0}^{N_{D^{s_k}}}$  denote a set of dialogs by the seeker  $s_k$  ( $0 \leq k < N_s$ ), where  $N_{D^{s_k}}$  is the number of dialogs by the seeker  $s_k$ , and  $N_s$  is the number of seekers. Recall that we attach each dialog (say  $d_i^{s_k}$ ) with an updated seeker profile (denoted as  $\mathcal{P}_i^{s_k}$ ), a knowledge graph  $\mathcal{K}$ , and a goal sequence  $\mathcal{G} = \{g_t\}_{t=0}^{T_g-1}$ . Given a context  $X$  with utterances  $\{u_j\}_{j=0}^{m-1}$  from the dialog  $d_i^{s_k}$ , a goal history  $\mathcal{G}' = \{g_0, \dots, g_{t-1}\}$  ( $g_{t-1}$  as the goal for  $u_{m-1}$ ),  $\mathcal{P}_{i-1}^{s_k}$  and  $\mathcal{K}$ , the aim is to provide an appropriate goal  $g_c$  to determine where the dialog goes and then produce a proper response  $Y = \{y_0, y_1, \dots, y_n\}$  for completion of the goal  $g_c$ .

**Framework overview** The overview of our framework MGCG is shown in Figure 3. The goal-planning module outputs goals to proactively and naturally lead the conversation. It first takes as input  $X$ ,  $\mathcal{G}'$ ,  $\mathcal{K}$  and  $\mathcal{P}_{i-1}^{s_k}$ , then outputs  $g_c$ . The responding module is responsible for completion of each goal by producing responses conditioned on  $X$ ,  $g_c$ , and  $\mathcal{K}$ . For implementation of the responding module, we adopt a retrieval model and a generation model proposed by Wu et al. (2019), and modify them to suit our task.

For model training, each [context, response] in  $d_i^{s_k}$  is paired with its ground-truth goal,  $\mathcal{P}_i^{s_k}$  and

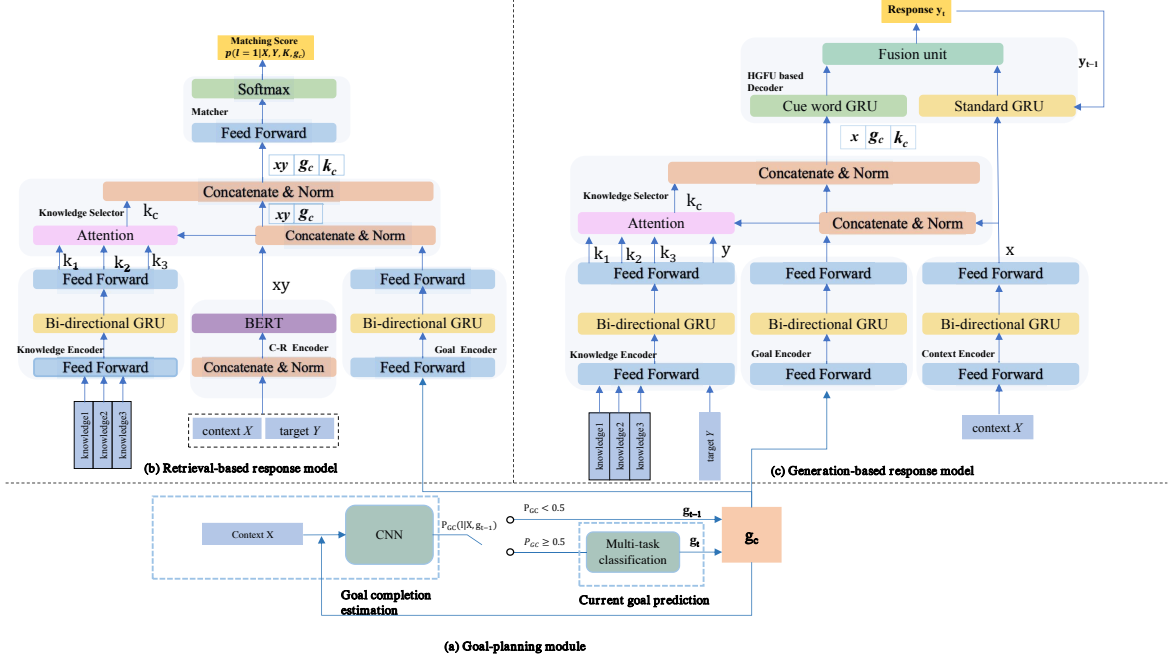


Figure 3: The architecture of our multi-goal driven conversation generation framework (denoted as **MGCG**).

$\mathcal{K}$ . These goals will be used as answers for training of the goal-planning module, while the tuples of [context, a ground-truth goal,  $\mathcal{K}$ , response] will be used for training of the responding module.

## 4.2 Goal-planning Model

As shown in Figure 3(a), we divide the task of goal planning into two sub-tasks, goal completion estimation, and current goal prediction.

**Goal completion estimation** For this subtask, we use Convolutional neural network (CNN)(Kim, 2014) to estimate the probability of goal completion by:

$$P_{GC}(l = 1|X, g_{t-1}). \quad (1)$$

**Current goal prediction** If  $g_{t-1}$  is not completed ( $P_{GC} < 0.5$ ), then  $g_c = g_{t-1}$ , where  $g_c$  is the goal for  $Y$ . Otherwise, we use CNN based multi-task classification to predict the current goal by maximizing the following probability:

$$g_t = \arg \max_{g^{ty}, g^{tp}} P_{GP}(g^{ty}, g^{tp}|X, \mathcal{G}', \mathcal{P}_i^{s_k}, \mathcal{K}), \quad (2)$$

$$g_c = g_t, \quad (3)$$

where  $g^{ty}$  is a candidate dialog type and  $g^{tp}$  is a candidate dialog topic.

## 4.3 Retrieval-based Response Model

In this work, conversational goal is an important guidance signal for response ranking. Therefore,

we modify the original retrieval model to suit our task by emphasizing the use of goals.

As shown in Figure 3(b), our response ranker consists of five components: a context-response representation module (C-R Encoder), a knowledge representation module (Knowledge Encoder), a goal representation module (Goal Encoder), a knowledge selection module (Knowledge Selector), and a matching module (Matcher).

The C-R Encoder has the same architecture as BERT (Devlin et al., 2018), and it takes a context  $X$  and a candidate response  $Y$  as segment.a and segment.b in BERT, and leverages a stacked self-attention to produce the joint representation of  $X$  and  $Y$ , denoted as  $xy$ .

Each related knowledge  $knowledge_i$  is also encoded as a vector by the Knowledge Encoder using a bi-directional GRU (Chung et al., 2014), which can be formulated as  $k_i = [\vec{h}_{T_k}, \overleftarrow{h}_0]$ , where  $T_k$  denotes the length of knowledge,  $\vec{h}_{T_k}$  and  $\overleftarrow{h}_0$  represent the last and initial hidden states of the two directional GRU respectively.

The Goal Encoder uses bi-directional GRUs to encode a dialog type and a dialog topic for goal representation (denoted as  $g_c$ ).

For knowledge selection, we make the context-response representation  $xy$  attended to all knowledge vectors  $k_i$  and get the attention distribution:

$$p(k_i|x, y, g_c) = \frac{\exp(\mathbf{MLP}([xy; g_c]) \cdot k_i)}{\sum_j \exp(\mathbf{MLP}([xy; g_c]) \cdot k_j)} \quad (4)$$

We fuse all related knowledge information into a single vector  $k_c = \sum_i p(k_i|x, y, g_c) * k_i$ .

We view  $k_c, g_c$  and  $xy$  as the information from knowledge source, goal source and dialogue source respectively, and fuse the three information sources into a single vector via concatenation. Finally we calculate a matching probability for each  $Y$  by:

$$p(l = 1|X, Y, K, g_c) = \text{softmax}(\text{MLP}([xy; k_c; g_c])) \quad (5)$$

#### 4.4 Generation-based Response Model

To highlight the importance of conversational goals, we also modify the original generation model by introducing an independent encoder for goal representation. As shown in Figure 3(c), our generator consists of five components: a Context Encoder, a Knowledge Encoder, a Goal Encoder, a Knowledge Selector, and a Decoder.

Given a context  $X$ , conversational goal  $g_c$  and knowledge graph  $K$ , our generator first encodes them as vectors with the use of above encoders (based on bi-directional GRUs).

We assume that using the correct response will be conducive to knowledge selection. Then minimizing KLDivloss will make the effect of knowledge selection in the prediction stage (not use the correct response) close to that of knowledge selection with correct response. For knowledge selection, the model learns knowledge-selection strategy through minimizing the KLDivLoss between two distributions, a *prior distribution*  $p(k_i|x, g_c)$  and a *posterior distribution*  $p(k_i|x, y, g_c)$ . It is formulated as:

$$p(k_i|x, y, g_c) = \frac{\exp(k_i \cdot \text{MLP}([x; y; g_c]))}{\sum_{j=1}^N \exp(k_j \cdot \text{MLP}([x; y; g_c]))} \quad (6)$$

$$p(k_i|x, g_c) = \frac{\exp(k_i \cdot \text{MLP}([x; g_c]))}{\sum_{j=1}^N \exp(k_j \cdot \text{MLP}([x; g_c]))} \quad (7)$$

$$L_{KL}(\theta) = \frac{1}{N} \sum_{i=1}^N p(k_i|x, y, g_c) \log \frac{p(k_i|x, y, g_c)}{p(k_i|x, g_c)} \quad (8)$$

In training procedure, we fuse all related knowledge information into a vector  $k_c = \sum_i p(k_i|x, y, g_c) * k_i$ , same as the retrieval-based method, and feed it to the decoder for response generation. In testing procedure, the fused knowledge is estimated by  $k_c = \sum_i p(k_i|x, g_c) * k_i$  without ground-truth responses. The decoder is implemented with the *Hierarchical Gated Fusion Unit* described in (Yao et al., 2017), which is a standard

GRU based decoder enhanced with external knowledge gates. In addition to the loss  $L_{KL}(\theta)$ , the generator uses the following losses:

**NLL Loss:** It computes the negative log-likelihood of the ground-truth response ( $L_{NLL}(\theta)$ ).

**BOW Loss:** We use the BOW loss proposed by Zhao et al. (2017), to ensure the accuracy of the fused knowledge  $k_c$  by enforcing the relevancy between the knowledge and the true response.<sup>4</sup> Specifically, let  $w = \text{MLP}(k_c) \in \mathcal{R}^{|V|}$ , where  $|V|$  is vocabulary size. We define:

$$p(y_t|k_c) = \frac{\exp(w_{y_t})}{\sum_{v=1}^V \exp(w_v)} \quad (9)$$

Then, the BOW loss is defined to minimize:

$$L_{BOW}(\theta) = -\frac{1}{m} \sum_{t=1}^m \log p(y_t|k_c) \quad (10)$$

Finally, we minimize the following loss function:

$$L(\theta) = \alpha \cdot L_{KL}(\theta) + \alpha \cdot L_{NLL}(\theta) + L_{BOW}(\theta) \quad (11)$$

where  $\alpha$  is a trainable parameter.

## 5 Experiments and Results

### 5.1 Experimental Setting

We split *DuRecDial* into train/dev/test data by randomly sampling 65%/10%/25% data at the level of seekers, instead of individual dialogs. To evaluate the contribution of goals, we conduct an ablation study by replacing input goals with ‘‘UNK’’ for responding model. For knowledge usage, we conduct another ablation study, where we remove input knowledge by replacing them with ‘‘UNK’’.

### 5.2 Methods<sup>5</sup>

**S2S** We implement a vanilla sequence-to-sequence model (Sutskever et al., 2014), which is widely used for open-domain conversation generation.

**MGCG\_R:** Our system with automatic goal planning and a retrieval based responding model.

**MGCG\_G:** Our system with automatic goal planning and a generation based responding model.

<sup>4</sup>The BOW loss is to introduce an auxiliary loss that requires the decoder network to predict the bag-of-words in the response to tackle the vanishing latent variable problem.

<sup>5</sup>Please see Appendix 2. for model parameter settings.

Methods↓ Metrics→	Hits@1/Hits@3	F1/ BLEU2	PPL	DIST-2	Knowledge P/R/F1
S2S- gl.- kg.	6.78% / 24.55%	23.97 / 0.065	27.31	0.011	0.275 / 0.209 / 0.216
S2S+gl.- kg.	8.03% / 27.71%	24.78 / 0.077	24.82	0.012	0.287 / 0.223 / 0.231
S2S+gl.+kg.	8.37% / 27.67%	24.66 / 0.072	23.96	0.011	0.295 / 0.239 / 0.253
MGCG_R- gl.- kg.	19.58% / 42.75%	33.22 / 0.207	-	0.171	0.344 / 0.301 / 0.306
MGCG_R+gl.- kg.	19.77% / 42.99%	33.78 / 0.223	-	0.185	0.351 / 0.322 / 0.309
MGCG_R+gl.+kg.	<b>20.33% / 43.61%</b>	33.93 / <b>0.232</b>	-	<b>0.187</b>	0.349 / 0.331 / 0.316
MGCG_G- gl.- kg.	13.26% / 36.07%	33.11 / 0.189	18.51	0.037	0.386 / 0.349 / 0.358
MGCG_G+gl.- kg.	14.21% / 38.91%	35.21 / 0.213	17.78	0.049	0.393 / 0.352 / 0.351
MGCG_G+gl.+kg.	14.38% / 39.70%	<b>36.81</b> / 0.219	<b>17.69</b>	0.052	<b>0.401 / 0.377 / 0.383</b>

Table 4: Automatic evaluation results. +(-)gl. represents “with(without) conversational goals”. +(-)kg. represents “with(without) knowledge”. For “S2S +gl.+kg.”, we simply concatenate the goal predicted by our model, all the related knowledge and the dialog context as its input.

Methods↓ Metrics→	Turn-level results				Dialog-level results	
	Fluency	Appro.	Infor.	Proactivity	Goal success rate	Coherence
S2S +gl. +kg.	1.08	0.23	0.37	0.94	0.37	0.49
MGCG_R +gl. +kg.	<b>1.98</b>	0.60	1.28	1.22	0.68	0.83
MGCG_G +gl. +kg.	1.94	<b>0.75</b>	<b>1.68</b>	<b>1.34</b>	<b>0.82</b>	<b>0.91</b>

Table 5: Human evaluation results at the level of turns and dialogs.

### 5.3 Automatic Evaluations

**Metrics** For automatic evaluation, we use several common metrics such as BLEU (Papineni et al., 2002), F1, perplexity (PPL), and DISTINCT (DIST-2) (Li et al., 2016) to measure the relevance, fluency, and diversity of generated responses. Following the setting in previous work (Wu et al., 2019; Zhang et al., 2018a), we also measure the performance of all models using Hits@1 and Hits@3.<sup>6</sup> Here we let each model to select the best response from 10 candidates. Those 10 candidate responses consist of the ground-truth response generated by humans and nine randomly sampled ones from the training set. Moreover, we also evaluate the knowledge-selection capability of each model by calculating knowledge precision/recall/F1 scores as done in Wu et al. (2019).<sup>7</sup> In addition, we also report the performance of our goal planning module, including the accuracy of goal completion estimation, dialog type prediction, and dialog topic prediction.

**Results** Our goal planning model can achieve accuracy scores of 94.13%, 91.22%, and 42.31% for goal completion estimation, dialog type prediction, and dialog topic prediction. The accuracy of dialog topic prediction is relatively low since the num-

ber of topic candidates is very large (around 1000), leading to the difficulty of topic prediction. As shown in Table 4, for response generation, both MGCG\_R and MGCG\_G outperform S2S by a large margin in terms of all the metrics under the same model setting (without gl.+kg., with gl., or with gl.+kg.). Moreover, MGCG\_R performs better in terms of Hits@k and DIST-2, but worse in terms of knowledge F1 when compared to MGCG\_G.<sup>8</sup> It might be explained by that they are optimized on different metrics. We also found that the methods using goals and knowledge outperform those without goals and knowledge, confirming the benefits of goals and knowledge as guidance information.

### 5.4 Human Evaluations

**Metrics:** The human evaluation is conducted at the level of both turns and dialogs.

For turn-level human evaluation, we ask each model to produce a response conditioned on a given context, the predicted goal and related knowledge.<sup>9</sup> The generated responses are evaluated by three annotators in terms of fluency, appropriateness, informativeness, and proactivity. The appropriateness measures if the response can complete current goal and it is also relevant to the context. The informativeness measures if the model makes full use of knowledge in the response. The proactivity mea-

<sup>6</sup>Candidates (including golden response) are scored by PPL using the generation-based model, then candidates are sorted based on the scores, and Hits@1 and Hits@3 are calculated.

<sup>7</sup>When calculating the knowledge precision/recall/F1, we compare the generated results with the correct knowledge.

<sup>8</sup>We calculate an average of F1 over all the dialogs. It might result in that the value of F1 is not between P and R.

<sup>9</sup>Please see Appendix 3. for more details.



ures if the model can successfully introduce new topics with good fluency and coherence.

For dialogue-level human evaluation, we let each model converse with a human and proactively make recommendations when given the predicted goals and related knowledge.<sup>10</sup> For each model, we collect 100 dialogs. These dialogs are then evaluated by three persons in terms of two metrics: (1) goal success rate that measures how well the conversation goal is achieved, and (2) coherence that measures relevance and fluency of a dialog as a whole.

All the metrics has three grades: good(2), fair(1), bad(0). For proactivity, “2” indicates that the model introduces new topics relevant to the context, “1” means that no new topics are introduced, but knowledge is used, “0” means that the model introduces new but irrelevant topics. For goal success rate, “2” means that the system can complete more than half of the goals from goal planning module, “0” means the system can complete no more than one goal, otherwise “1”. For coherence, “2”/“1”/“0” means that two-thirds/one-third/very few utterance pairs are coherent and fluent.

**Results** All human evaluations are conducted by three persons. As shown in Table 5, our two systems outperform S2S by a large margin, especially in terms of appropriateness, informativeness, goal success rate and coherence. In particular, S2S tends to generate safe and uninformative responses, failing to complete goals in most of dialogs. Our two systems can produce more appropriate and informative responses to achieve higher goal success rate with the full use of goal information and knowledge. Moreover, the retrieval-based model performs better in terms of fluency since its response is selected from the original human utterances, not automatically generated. But it performs worse on all the other metrics when compared to the generation-based model. It might be caused by the limited number of retrieval candidates. Finally, it can be seen that there is still much room for performance improvement in terms of appropriateness and goal success rate, which will be left as the future work.

## 5.5 Result Analysis

In order to further analyze the relationship between knowledge usage and goal completion, we provide the number of failed goals, completed goals, and used knowledge for each method over different dialog types in Table 6. We see that the number of

<sup>10</sup>Please see Appendix 4. for more details.

Methods→		S2S	MGCG_R	MGCG_G
Metrics↓	Types↓	+gl. +kg.	+gl. +kg.	+gl. +kg.
#Failed gl./ #Com- pleted gl.	Rec.	106/7	95/18	93/20
	Chitchat	120/93	96/117	80/133
	QA	66/5	61/10	60/11
	Task	45/4	36/13	39/10
Overall		337/109	288/158	272/174
#Used kg.	Rec.	0	8	7
	Chitchat	9	25	33
	QA	5	10	15
	Task	0	3	2
Overall		14	46	57

Table 6: Analysis of goal completion and knowledge usage across different dialog types.

used knowledge is proportional to goal success rate across different dialog types or different methods, indicating that the knowledge selection capability is crucial to goal completion through dialogs. Moreover, the goal of chitchat dialog is easier to complete in comparison with others, and QA and recommendation dialogs are more challenging to complete. How to strengthen knowledge selection capability in the context of multi-type dialogs, especially for QA and recommendation, is very important, which will be left as the future work.

## 6 Conclusion

We identify the task of conversational recommendation over multi-type dialogs, and create a dataset *DuRecDial* with multiple dialog types and multi-domain use cases. We demonstrate usability of this dataset and provide results of state of the art models for future studies. The complexity in *DuRecDial* makes it a great testbed for more tasks such as knowledge grounded conversation (Ghazvininejad et al., 2018), domain transfer for dialog modeling, target-guided conversation (Tang et al., 2019a) and multi-type dialog modeling (Yu et al., 2017). The study of these tasks will be left as the future work.

## Acknowledgments

We would like to thank Ying Chen for dataset annotation and thank Yuqing Guo and the reviewers for their insightful comments. This work was supported by the National Key Research and Development Project of China (No. 2018AAA0101900) and the Natural Science Foundation of China (No. 61976072).

## References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. MultiWOZ - a large-scale multi-domain wizard-of-Oz dataset for task-oriented dialogue modelling. In *EMNLP*.
- Qibin Chen, Junyang Lin, Yichang Zhang, Ming Ding, Yukuo Cen, Hongxia Yang, and Jie Tang. 2019. Towards knowledge-based recommender dialog system. In *ACL*.
- Konstantina Christakopoulou, Alex Beutel, Rui Li, Sagar Jain, and Ed H. Chi. 2018. Q and r: A two-stage approach toward interactive recommendation. In *KDD*.
- Konstantina Christakopoulou, Katja Hofmann, and Filip Radlinski. 2016. Towards conversational recommender systems. In *KDD*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *arXiv preprint arXiv:1412.3555*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *arXiv preprint arXiv:1810.04805*.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of wikipedia: knowledge-powered conversational agents. In *Proceedings of ICLR*.
- Jesse Dodge, Andreea Gane, Xiang Zhang, Antoine Bordes, Sumit Chopra, Alexander H. Miller, and Arthur Szlam and Jason Weston. 2016. Evaluating prerequisite qualities for learning end-to-end dialog systems. In *ICLR*.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *AAAI*.
- 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 *EMNLP*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar. Association for Computational Linguistics.
- Sunhwan Lee, Robert Moore, Guang-Jie Ren, Raphael Arar, and Shun Jiang. 2018. Making personalized recommendation through conversation: Architecture design and recommendation methods. In *AAAI Workshops*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *NAACL-HLT*, pages 110–119.
- Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards deep conversational recommendations. In *NIPS*.
- Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M. Khapra. 2018. Towards exploiting background knowledge for building conversation systems. In *EMNLP*.
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. Opendialkg: Explainable conversational reasoning with attention-based walks over knowledge graphs. In *ACL*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, Eric King, Kate Bland, Amanda Wartick, Yi Pan, Han Song, Sk Jayadevan, Gene Hwang, and Art Pettigrew. 2018. Conversational AI: the science behind the alexa prize. In *CoRR, abs/1801.03604*.
- Kevin Reschke, Adam Vogel, and Daniel Jurafsky. 2013. Generating recommendation dialogs by extracting information from user reviews. In *ACL*.
- Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *SIGIR*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Jianheng Tang, Tiancheng Zhao, Chengyan Xiong, Xiaodan Liang, Eric P Xing, and Zhiting Hu. 2019a. Target-guided open-domain conversation. In *ACL*.
- Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric P. Xing, and Zhiting Hu. 2019b. Target-guided open-domain conversation. In *Proceedings of ACL*.
- Zhuoran Wang, Hongliang Chen, Guanchun Wang, Hao Tian, Hua Wu, and Haifeng Wang. 2014. Policy learning for domain selection in an extensible multi-domain spoken dialogue system. In *EMNLP*.
- Pontus Warnestal. 2005. Modeling a dialogue strategy for personalized movie recommendations. In *The Beyond Personalization Workshop*.
- 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 *ACL*.

Jun Xu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, and Wanxiang Che. 2020. Knowledge graph grounded goal planning for open-domain conversation generation. In *AAAI*.

Lili Yao, Yaoyuan Zhang, Yansong Feng, Dongyan Zhao, and Rui Yan. 2017. Towards implicit content-introducing for generative short-text conversation systems. In *EMNLP*.

Zhou Yu, Alexander I. Rudnicky, and Alan W. Black. 2017. Learning conversational systems that interleave task and non-task content. In *IJCAI*.

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 *ACL*.

Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018b. Towards conversational search and recommendation: System ask, user respond. In *CIKM*.

Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *ACL*.

Hao Zhou, Chujie Zheng, Kaili Huang, Minlie Huang, and Xiaoyan Zhu. 2020. Kdconv: A chinese multi-domain dialogue dataset towards multi-turn knowledge-driven conversation. In *Proceedings of ACL*.

Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018a. A dataset for document grounded conversations. In *Proceedings of EMNLP*.

Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. 2018b. The design and implementation of xiaoice, an empathetic social chatbot. In *CoRR, abs/1812.08989*.

## Appendix

### 1. Dataset collection process

#### 1.1 Collection of seeker profiles/knowledge graph/task templates

**Collection of seeker profile** The attributes of seeker profiles are shown as follows: name, gender, age range, city of residence, occupation status, and seeker preference. Seeker preference includes : domain preference, seed entity preference, entity list rejected by the seeker, entity list accepted by the seeker.

- **Name:** We generate the first-name Chinese character(or last-name Chinese character) by randomly sampling Chinese characters from a set of candidate characters used as first name (or last name) for the gender of the seeker.

- **Gender:** We randomly select "male" or "female" as the seeker's gender.
- **Age range:** We randomly choose one from the 5 age ranges.
- **Residential city:** We randomly choose one from 55 cities in China as the seeker's residential city.
- **Occupation status:** We randomly choose one from "student", "worker" and "retirement" based on above age range.
- **Domain preference:** We randomly select one or two domains as ones that the seeker likes (e.g., movie, food), and one domain as the one that the seeker dislikes (e.g., news). It will affect the setting of task templates for this seeker.
- **Seed entity preference:** We randomly select one or two entities from KG entities of the domains preferred by the seeker as his/her preference at entity level. It will affect the setting of task templates for this seeker.
- **Rejected entity list and accepted entity list:** Both of them are empty at the beginning, and they will be updated as the conversation progresses; the two lists will affect the recommendation results of subsequent conversations to some extent.

**Collection of Knowledge graph (KG)** The domain of knowledge graph include stars, movies, music, news, food, POI(Point of Interest), weather.

- **Stars:** including the introduction, achievements, awards, comments, birthday, birthplace, height, weight, blood type, constellation, zodiac, nationality, friends, etc
- **Film:** including film rating, comments, region, leading role, director, category, evaluation, award, etc
- **Music:** singer information, comments, etc
- **News:** including the topic, content, etc
- **Food:** including the name, ingredients, category, etc
- **POI:** including restaurant name, average price, score, order quantity, address, city, specialty, etc

- **Weather:** historical weather of 55 cities from July 2017 to August 2019

**Collection of task templates** First we manually annotate a list of around 20 high-level goal sequences as candidates. Most of these goal sequences include 3 to 5 high-level goals. Here each high-level goal contains a dialog type and a domain (not entity or chatting topic). Then for each seeker, we select appropriate high-level goal sequences from the above list, which contains the domains that fall into the seeker’s preferred domain list.

To collect goal sequences at entity level, we first use the seed entities of the seeker to enrich the information of above high-level goal sequences. If the seed entities are not enough, or there is no seeds for some domains in the high-level goal sequences, we select some entities from KG for each goal domain based on embedding based similarity scores of the seed entities (of current seeker) and the candidate entity. Then we obtain goal sequences at entity level. Finally we use some rules to generate a description for each goal (e.g., which side, the seeker or the recommender, to start the dialog, how to complete the goal). Thus we have task templates for guidance of data annotation.

To introduce diverse interaction behavior for recommendation, we design some fine-grained interaction operations, e.g., the seeker may reject the initial recommendation, or mention a new topic, or ask a question about an entity, or simply accept the recommendation. Each interaction operation corresponds to a goal. We randomly sample one of above operations and insert it into the entity-level goal sequences to diversify recommendation dialogs. The entities associated with the above interaction operations are selected from the KG based on their similarity scores with current seeker’s seed entities. If the entity will be accepted by the seeker as described in the task templates (including entity-level goal sequence and its description), then its similarity score with the seeker’s seed entities should be relatively high. If the entity will be rejected by the seeker as described in the task templates, then its similarity score with the seeker’s seed entities should be relatively low.

## 1.2 Dataset annotation process

We first release a small amount of data for training, and then carry out video training for annotation problems. After that, a small amount of data is released again to select the final task workers. To

ensure that at least two workers enter the task at the same time, we arrange multiple workers to log in the annotation platform. During annotation, each conversation is randomly assigned to two workers, one of whom plays the role of BOT and the other plays the role of User. Two workers conduct annotation based on the seeker profile, knowledge graph and task templates.

Due to the complexity of our task design, the quality of data annotation may have a high variance. To address this problem, we provide a strict data annotation standard to guide the workers to annotate in the way we expect them to be. After data annotation, multiple data specialists will review it. As long as one specialist thinks it does not meet the requirements, it will be marked back and re-annotation until all specialists think that it fully meets requirements.

## 2. Model Parameter Settings

All models are implemented using PaddlePaddle.<sup>11</sup> The parameters of all the modules are shown in Table 7.<sup>12</sup>

## 3. Turn-level Human Evaluation Guideline

**Fluency** measures if the produced response itself is fluent:

- score 0 (bad): unfluent and difficult to understand.
- score 1 (fair): there are some errors in the response text but still can be understood.
- score 2 (good): fluent and easy to understand.

**Appropriateness** measures if the response can respond to the context:

- score 0 (bad): Sub-dialogs for Rec and chitchat: not semantically relevant to the context or logically contradictory to the context. Sub-dialogs for task-oriented: No necessary slot value is involved in the conversation. Sub-dialogs for QA: Incorrect answer.
- score 1 (fair): relevant to the context as a whole, but using some irrelevant knowledge, or not answering questions asked by the users.
- score 2 (good): otherwise.

<sup>11</sup>It is an open source deep learning platform(<https://www.paddlepaddle.org.cn/>)

<sup>12</sup>Due to S2S model uses the same parameters as OpenNMT(<https://github.com/OpenNMT/OpenNMT-py>), its parameters are not listed.

module	Parameter	value
Goal-planning model	Embedding Size	256
	Hidden Size	256
	Batch Size	128
	Learning Rate	0.002
	Optimizer	Adam
Retrieval-based model	Dropout	0.1
	Embedding Size	512
	Hidden Size	512
	Batch Size	32
	Learning Rate	0.001
	Optimizer	Adam
	Weight Decay	0.01
Generation-based model	Proportion Warmup	0.1
	Embedding Size	300
	Hidden Size	800
	Batch Size	16
	Learning Rate	0.0005
	Grad Clip	5
	Dropout	0.2
	Beam Size	10
Optimizer	Adam	

Table 7: Model parameter settings.

**Informativeness** measures if the model makes full use of knowledge in the response:

- score 0 (bad): no knowledge is mentioned at all.
- score 1 (fair): only one knowledge triple is mentioned in the response.
- score 2 (good): more than one knowledge triple is mentioned in the response.

**Proactivity** measures if the model can introduce new knowledge/topics in conversation:

- score -1 (bad): some new topics are introduced but irrelevant to the context.
- score 0 (fair): no new topics/knowledge are used.
- score 1 (good): some new topics relevant to the context are introduced.

#### 4. Dialogue-level Human Evaluation Guideline

**Goal Completion** measures how good the given conversation goal is finished:

- score 0 (bad): less than half goals are achieved..
- score 1 (fair): less than half goals are achieved with minor use of knowledge or goal information.

- score 2 (good): more than half goals are achieved with full use of knowledge and goal information.

**Coherence** measures the overall fluency of the whole dialogue:

- score 0 (bad): two-thirds responses irrelevant or logically contradictory to the previous context.
- score 1 (fair): less than one-third responses irrelevant or logically contradictory to the previous context.
- score 2 (good): very few response irrelevant or logically contradictory to the previous context.

#### 5. Case Study

Figure 4 shows the conversations generated by the models via conversing with humans, given the conversation goal and the related knowledge. It can be seen that our knowledge-aware generator can use more correct knowledge for diverse conversation generation. Even though the retrieval-based method can also produce knowledge-grounded responses, the used knowledge is relatively few and inappropriate. The seq2seq model can't successfully complete the given goal, as the knowledge is not fully used as our proposed knowledge-aware generator, making the generated conversation less diverse and sometimes dull.

Goal sequence	问答(张柏芝)---->闲聊(明星)---->电影推荐(左右逢源)---->再见(再见) QA(Cecilia Cheung) ----> Chitchat(Star) ----> Movie recommendation(Left Right Love Destiny) ----> Goodbye(Goodbye)		
Knowledge			
	S2S	MGCG_R	
	<p>User:你知道电影『钟无艳』主演是谁吗？(Who is the main actress of Wu Yen?)</p> <p>Bot: 你有喜欢的明星吗？(Do you have a favorite star?)</p> <p>User:我喜欢张柏芝啊！(I like Cecilia Cheung)</p> <p>Bot: 是钟无艳的歌！(It's Wu Yen's song)</p> <p>User:你说的是什么啊？(What are you talking about?)</p> <p>Bot: 是啊，还有一部电影叫《恋爱大赢家》，你看过吗？(Yes, there is another movie called "The Love Winner". Have you seen it?)</p> <p>User:暂时没看过啊。(Not yet.)</p> <p>Bot: 好的，拜拜(OK, goodbye.)</p> <p>User:再见(Bye)</p> <p>Bot: 好的，拜拜。(OK, goodbye.)</p>	<p>User:你知道电影『钟无艳』主演是谁吗？(Who is the main actress of Wu Yen?)</p> <p>Bot: 是张柏芝哦(Cecilia Cheung)</p> <p>User:谢谢啊，你好聪明。(Thank you. You're smart.)</p> <p>Bot: 是吗？真眼光呢，她可是香港第一美女呢(Is it? You have good taste. She is the first beauty in Hong Kong.)</p> <p>User:除了漂亮，我很喜欢她演戏时的忘我与投入。(Apart from beauty, I like her selflessness and devotion in acting.)</p> <p>Bot: 你看过她演的那部电影《钟无艳》吗？(Have you seen her movie "Wu Yen"?)</p> <p>User:嗯，她很漂亮。(Well, she's very beautiful.)</p> <p>Bot: 你会喜欢上的。(You'll like it)</p> <p>User:暂时没看过啊。(I haven't seen it yet.)</p> <p>Bot: 好的，再见(OK, goodbye.)</p> <p>User:嗯，我比较喜欢。(Well, I like it very much.)</p> <p>Bot: 再见。(Bye.)</p>	<p>User:你知道电影『钟无艳』主演是谁吗？(Who is the main actress of Wu Yen?)</p> <p>Bot: 张柏芝(Cecilia Cheung)</p> <p>User:谢谢啊，你好聪明。(Thank you. You're smart.)</p> <p>Bot: 你是她的粉丝吗？(Are you a fan of her?)</p> <p>User:是啊，我很喜欢张柏芝。(Yes, I like Cecilia Cheung very much.)</p> <p>Bot: 她是香港第一美女呢，你真有眼光呢。(She is the first beauty in Hong Kong. You have good taste.)</p> <p>User:除了漂亮，我很喜欢她演戏时的忘我与投入。(Apart from beauty, I like her selflessness and devotion in acting)</p> <p>Bot: 给你推荐一部张柏芝主演的电影《左右情缘》，很法式的一部短片。(I'd like to recommend you "Left Right Love Destiny" starring Cecilia Cheung, which is a short film with so much french style.)</p> <p>User:这个电影评分多少呢？(What's the rating of this movie?)</p> <p>Bot: 评分是8.3呢。(The rating is 8.3.)</p> <p>User:那还不错，我抽时间要看看。(That's well. I'll see it later.)</p> <p>Bot: 嗯呢，相信你会喜欢的。(Mm-hmm. I'm sure you'll like it.)</p> <p>User:嗯，我有点事，回见。(Well, I have something to do. See you later.)</p> <p>Bot: 好的，再见(OK, goodbye.)</p>

Figure 4: Conversations generated by three different models: texts in red color represent correct knowledge being appropriate in current context, while texts in blue color represent inappropriate knowledge. Texts in purple color indicate that the use of knowledge is correct, but the response is not appropriate.