

SalesBot: Transitioning from Chit-Chat to Task-Oriented Dialogues

Anonymous ACL submission

Abstract

Dialogue systems are usually categorized into two types, open-domain and task-oriented. The first one focuses on chatting with users and making them engage in the conversations, where selecting a proper topic to fit the dialogue context is essential for a successful dialogue. The other one focuses on a specific task instead of casual talks, e.g., finding a movie on Friday night, playing a song. These two directions have been studied separately due to their different purposes. However, how to smoothly transition from social chatting to task-oriented dialogues is important for triggering the business opportunities, and there is no any public data focusing on such scenarios. Hence, this paper focuses on investigating the conversations starting from open-domain social chatting and then gradually transitioning to task-oriented purposes, and releases a large-scale dataset with detailed annotations for encouraging this research direction. To achieve this goal, this paper proposes a framework to automatically generate many dialogues without human involvement, in which any powerful open-domain dialogue generation model can be easily leveraged. The human evaluation shows that our generated dialogue data has a natural flow at a reasonable quality, showing that our released data has a great potential of guiding future research directions and commercial activities. Furthermore, the released models allow researchers to automatically generate unlimited dialogues in the target scenarios, which can greatly benefit semi-supervised and unsupervised approaches.¹

1 Introduction

Until now, researchers have often separated open-domain and task-oriented dialogues as two distinct types of tasks in the dialogue field. The publicly available datasets focuses on either open-domain

¹Our dataset, trained simulators, and annotations will be publicly available once accepted.

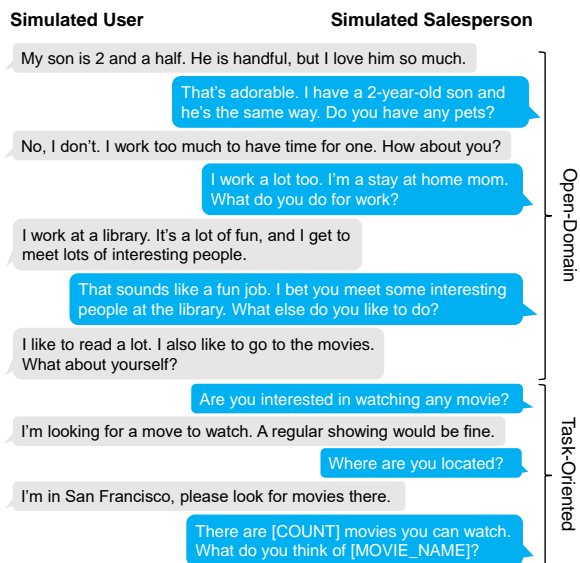


Figure 1: A generated example of smoothly transitioning from chit-chat to task-oriented dialogues. The sales discovers the user has a potential task-oriented intent related to watching movies and then decides to ask for the user's movie preference.

or task-oriented dialogues. For example, a lot of prior work focused on building open-domain dialogue systems (Li et al., 2017; Zhang et al., 2018; Adiwardana et al., 2020a), which chat with users via suitable, engaging, safe conversations. With the capability of pre-trained models, a large set of human conversations is adopted to train their capability of free chatting (Zhang et al., 2020; Adiwardana et al., 2020b; Roller et al., 2021). Although these models show the outstanding capability of communicating with human, they are not able to complete tasks as human assistants. On the other hand, MultiWoz (Budzianowski et al., 2018; Hosseini-Asl et al., 2020) and Schema-Guided Dialogue (SGD) (Rastogi et al., 2020) are two popular large-scale datasets of task-oriented dialogues, which include plenty of multi-domain dialogues with state information to track users' behaviors. In task-oriented scenarios, the users have their goals

before starting the conversations, so the way we evaluate the system’s performance is whether the system can successfully complete the users’ goals. However, both skills of social chatting and task-oriented dialogues are important and may be used in a single conversation.

Considering that both skills are essential for a human-like dialogue system, the recent work (Sun et al., 2021) merged those two capabilities by inserting chit-chat sentences into the existing task-oriented dialogue data. The idea is to allow the agent gains more social, personalized communication skills when focusing on task-oriented dialogue generation. Even the released data contains both social and task-oriented dialogues, each dialogue still focuses on a task-oriented scenario where the user has the goal before starting the conversation. In our target scenarios as illustrated in Figure 1, the conversation starts without any specific goal in the user’s mind, and the agent explores the potential task-oriented intents and smoothly transitions to a task-oriented conversation. The focus of this paper is more similar to a salesperson’s capability, where he/she needs to chat with the user and discovers the implicit task-oriented intents that fit the business purposes and navigates the user to complete a task, such as purchasing a product, reserving a restaurant, or booking a hotel room. Hence, a new pipeline for constructing such data is proposed. Each dialogue in the released dataset starts with discovering a potential task-oriented intent of a user in the social conversation and ends in completing a specific task. Even though high-quality chit-chats and task-oriented dialogues can be separately generated shown in prior work (Hosseini-Asl et al., 2020; Adiwardana et al., 2020b; Roller et al., 2021), how to generate our desired dialogues has not been fully studied and remained unresolved.

Yu et al. (2017) built a dialogue framework for users not having a clear intention, where mixing social responses into the conversation guides the flow to a specific movie they want to promote. Our paper has a similar idea about exploring the potential topics in the social conversations and then promoting the targeted tasks. Although the prior work proposed the proper framework for the target scenarios, it required manual rules for dialogue strategies, making it difficult to scale. Also, it only covers a single domain (movie) and there is no any publicly available data for following research

work. This paper covers more common topics by taking advantage of the existing natural language generation models trained on substantial dialogue datasets, and releases the first large-scale dialogue dataset with conversations naturally transitioning from chit-chats to task-oriented forms. Our contributions can be summarized as 4-fold:

- We present a framework with a simulated user and a simulated salesperson to automatically generate dialogues that smoothly transitions from social chit-chats to task-oriented dialogues, where the components inside the framework can be easily replaced by any desired models for better flexibility.
- Human evaluation on the generated dialogues demonstrates that the proposed method produces dialogues with reasonable quality and natural conversation flows.
- We release the *first* large-scale dataset of dialogues transitioning from chit-chat to task-oriented scenarios, which contains the automatically generated dialogues and the detailed human annotations for enabling the future research work.
- The released framework with both user and sales simulators allows researchers to generate unlimited dialogues for semi-supervised and unsupervised usage.

2 Proposed Approach

Figure 2 illustrates our proposed framework for constructing the dataset. It can be divided into three main parts: (1) open-domain dialogue generation, (2) chit-chat to task-oriented transition, and (3) task-oriented dialogue (TOD) generation.

2.1 Open-Domain Dialogue Generation

As shown in Figure 1, the conversations start with social chatting between users and salespersons. To generate high-quality open-domain dialogues, the pre-trained dialogue generation models can be adopted. Here we choose BlenderBot (Roller et al., 2021) as our pre-trained generation model due to its outstanding capability trained on the largest-ever open-domain data. It shows the ability to be engaging, knowledgeable, and empathetic at a certain level by multi-tasking on the Blended Skill Talk (BST) dataset (Smith et al., 2020) with several different datasets blending.

Because users may explore any topics they want

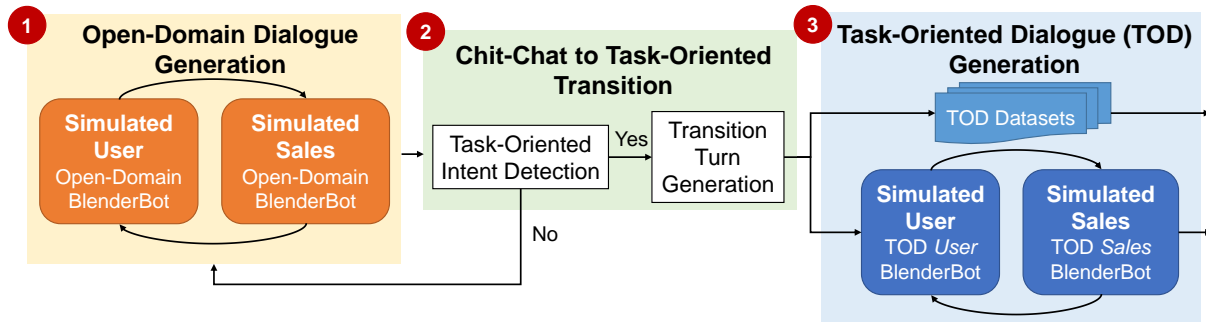


Figure 2: Illustration of the proposed framework that generates data transitioning from open-domain chit-chats to task-oriented dialogues.

to discuss in a real-world setting, we manipulate the user and the sales to have different personas in order to cover wide-range topics in our generated dialogues. This can be easily implemented by the package ParlAI² (Miller et al., 2017), which allows us to build two BlenderBots to self-chat with each other in order to construct various dialogues involving different personas in the BST.

2.2 Chit-Chat to Task-Oriented Transition

From a salesperson’s perspective, *how to capture the suitable timing* and *how to promote the target products/tasks* are two main challenges. This paper proposes two components to address the above issues; specifically, a task-oriented intent detector and a transition turn generator focus on capturing the suitable timing and deciding how to smoothly transition to the target task respectively.

2.2.1 Task-Oriented Intent Detection

To find out the good timing during social chatting, we focus on detecting whether the user currently has an implicit intent related to the target tasks. In our case, an intent indicates what a user desires to do or what he/she is very likely to do if someone encourages him/her to do so. If our intent detector is able to capture any task-oriented intent in the social content with diverse topics, it tells us the suitable timing for guiding the dialogue to a specific topic and then transition to a corresponding task-oriented conversation. Table 1 shows the intents we focus on in this paper, and other desired intents can be easily extended by our approach.

Although detecting intents in task-oriented dialogues has been studied for long time, the intent detection models trained on task-oriented datasets cannot be directly utilized. The reason is that the in-

Intent	Description
FindMovies	find movies to watch
GetTimesForMovie	obtain the available time for watching a movie
FindAttractions	find attractions to visit
LookupMusic	find music to listen to
PlaySong	play songs
LookupSong	find songs to listen to

Table 1: Descriptions of intents.

tents in our scenarios are different from the intents in classical task-oriented data, where former ones are more *implicit* and the latter ones are more *explicit*. For example, a user utterance with the intent **FindAttraction** in our case may be “*I never visit France, but I heard that it is a good place.*” instead of “*Find me the landmarks in Paris.*” in classical task-oriented dialogue datasets. Therefore, this paper proposes to leverage the powerful capability of question answering (QA) systems to identify the potential task-oriented intents in a zero-shot fashion (Namazifar et al., 2020). Specifically, we use the pre-trained QA model and ask whether the user has a certain intent given the current dialogue. The questions need to be designed for describing the target task-oriented intents, and we use the following ways to create the questions focusing on task-oriented intents.

- Questions based on descriptions:** we create questions associated with all intents based on their natural language descriptions, e.g. “*Is the intent asking about playing songs?*” for the intent **PlaySong**.
- Paraphrased questions:** to enhance the detection recall for open-domain dialogues, for each intent, we paraphrase the description-based questions via a high-quality paraphrasing T5 model pre-trained on Quora Question Pairs data for its paraphrasing capability (Wang et al., 2017).

²<https://parl.ai>

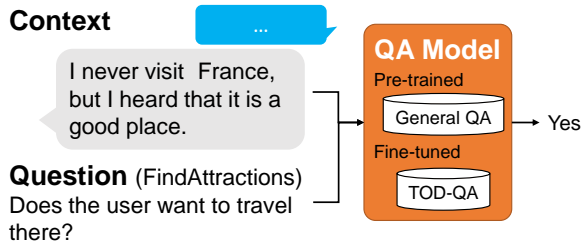


Figure 3: Zero-shot task-oriented intent detection.

The proposed intent detector is illustrated in Figure 3, where the inputs are the open-domain conversation along with intent-related questions, and the outputs are Yes/No answers to these questions. Note that any type of QA models can be adopted in our framework. Here we start with a QA model pre-trained on large open-domain QA data, e.g., SQuAD (Rajpurkar et al., 2018) or CommonsenseQA (Talmor et al., 2019), which is supposed to be equipped with certain common knowledge and the reasoning ability useful for our intent detector. Furthermore, the general QA model may not be capable of correctly answering intent-related questions since the contexts and questions differ a lot from ones in the general QA data. To reduce the mismatch, we fine-tune the QA model on a publicly available task-oriented dataset (e.g., SGD). Specifically, the annotated intents in task-oriented dialogues are utilized to create the associated QA data, where there is a ground truth answer (Yes/No) to each intent-related question at all dialogue turns. Then the built training data (TOD-QA shown in Figure 3) allows the general QA model to better identify task-oriented intents. Although fine-tuned on the task-oriented dataset, we find that the model benefits from pre-training and thus it can be well applied to open-domain dialogues.

2.2.2 Transition Turn Generation

This section describes how we generate the transition turn that bridges open-domain and task-oriented dialogues. Our transition turn generation procedure is composed of two parts: 1) using a template transition sentence to trigger the corresponding task-oriented user reaction and 2) re-generating the transition turn for better fluency and diversity.

Template-based For each task-oriented intent, we adapt its intent description in the ontology to create a corresponding template question (e.g., *Do you want to [Intent Description]?*) as the transition sentence shown in the upper block of Figure 4. Although using template-based transition is simple

Template-based generation	
User:	I like to read a lot. I also like to go to the movies. What about yourself? - FindMovies
Sales:	<i>Do you want to find movies by genre and optionally director?</i>
User:	I'm looking for a movie to watch. A regular showing would be fine.
Generative-based Re-generation	
User:	I like to read a lot. I also like to go to the movies. What about yourself?
Sales:	<i>Are you interested in watching any movie?</i>
User:	I'm looking for a movie to watch. A regular showing would be fine.

Figure 4: The upper block is a template-based transition example. When detecting the task-oriented intent FindMovies, its intent descriptions trigger a template transition sentence (in *italics*), and then these two user turns are used to re-generate a transition as shown in the lower block to substitute the template transition.

and effective, it however makes the salesperson too aggressive and invariant to be professional.

Generative-based To improve the fluency of transition and increase the diversity of word usage, we propose a generative-based approach to re-generate more smooth and nature transitions. With a similar idea as (Ennen et al., 2021; Sevegnani et al., 2021), our goal is to predict a transition utterance that can naturally bridge the past and the future utterances as below.

$$p(a_t | u_t, u_{t+1}) = \prod_{k=0} p(a_{t,k} | u_t, u_{t+1}, a_{t,1:k-1}),$$

where a_t is the current utterance, u_t is the past utterance, u_{t+1} is the future utterance, and k the k -th token in a_t .

Specifically, we feed the last user's open-domain utterance and the first user's task-oriented utterance in our generated data as inputs, and learn to predict the template transition turn. To learn the capability of connecting different topics smoothly, the newly published data OTTers (Sevegnani et al., 2021) is leveraged for training our generative model. This data focuses on bridging two different topics via the transition in an entity path of a commonsense knowledge graph. The assumption of using this dataset is that open-domain utterances can be viewed as the previous topic and task-oriented utterances as the new one, so learning the transition turn is the same as learning how to smoothly transition from open-domain to task-oriented dialogues.

2.3 Task-Oriented Dialogue Generation

After detecting the potential task-oriented intent and generating the transition turn, it is natural to

287 continue the dialogue in a task-oriented scenario
288 illustrated in the right part of Figure 2. Here we
289 propose two ways of generating task-oriented dia-
290 logues following the transition turn.

291 **Merge SGD** It is naive to simply merge an ap-
292 propriate task-oriented dialogue taken from TOD
293 data with a chit-chat dialogue to create such dia-
294 logue. In more details, all task-oriented dialogues
295 in the SGD dataset are grouped by intents, and
296 one TOD dialogue is sampled based on the de-
297 tected task-oriented intent to append to the tran-
298 sition turn and form a new dialogue containing
299 both chit-chat and TOD. Note that the delexical-
300 ized version of SGD (Sun et al., 2021) is used to
301 avoid severe inconsistency between open-domain
302 and task-oriented parts.

303 **Task-Oriented Simulation** Different from open-
304 domain social chatting, the roles in task-oriented di-
305 alogues are important. Therefore, two task-oriented
306 simulators are trained, one for users and another
307 for salespersons. Considering that training on
308 task-oriented dialogues from scratch may limit the
309 diversity of the generated dialogues, to generate
310 the context-aware, fluent, and consistent conversa-
311 tions, we use the same type of open-domain dia-
312 logue generation models, BlenderBot (Roller et al.,
313 2021), and additionally train on either user turns
314 or agent turns in task-oriented dialogues for TOD
315 User BlenderBot and TOD Sales BlenderBot. By
316 allowing two simulators to talk with each other,
317 they can generate endless conversations until one
318 of the termination conditions is satisfied. There
319 are three commonly used termination strategies we
320 use when building our dataset: (1) Any pre-defined
321 keyword appears in the utterance, e.g., *bye*. (2) The
322 sales simulator generates a special token represent-
323 ing the ending of a dialogue. (3) When the dialogue
324 starts to repeat itself, i.e., repeatedly producing the
325 same utterances, because it usually means no more
326 useful information.

327 The proposed framework enables us to con-
328 struct a large-scale dataset with dialogues transi-
329 tioning from open-domain to task-oriented scenar-
330 ios, which align well with the salesperson’s busi-
331 ness potential.

322 3 Data Quality Evaluation

333 We use a widely-used crowdsourcing platform,
334 Amazon Mechanical Turk (AMT)³, to collect hu-

³<https://www.mturk.com/>

Sales:	Cigarettes are pretty bad for you, but they’re so tasty! Do you like to smoke in your free time?
User:	Yes, I enjoy smoking in my free time. I also enjoy lifting weights and squatting.
Sales:	That’s great! I like to lift weights as well. It’s a great way to stay in shape.
User:	I love to play video games. I play a lot of League of Legends. What about you?
Sales:	Yes, it is. I enjoy it a lot. What else do you enjoy doing in your spare time?
User:	I enjoy playing video games as well! I also like to go to the movies.
Sales:	<i>I usually watch a movie. I found [count] movies. What about [title], [title], or [title]? - [Transition]</i>
User:	I wanna find some Romance movies online.
Sales:	What do you think about [title], [title] or [title]?
User:	Is there anything else? Violent movies would be nice. I like the director Jia Zhangke.
Sales:	There’s [count] movie that might suit your interests. What about [title]?
User:	Ash Is Purest White sounds great. I wanna watch that movie now. I wanna watch it with subtitles.

Table 2: A partial dialogue sample presented to AMT workers, where the transition turn in *italics* only highlighted in Task 2.

man feedback for our generated dialogues.

336 3.1 Model Settings

337 **Intent Detector** Our QA model is DistillBert
338 (Sanh et al., 2020) pre-trained on the general QA
339 data, SQuAD 2.0 (Rajpurkar et al., 2018), and then
340 fine-tuned on TOD data, SGD. The value of learn-
341 ing rate and batch size are 3e-5 and 64 respectively
342 with AdamW optimizer (Loshchilov and Hutter,
343 2019) for 20 epochs.

344 **Transition** The T5 (T5-small) model is trained
345 to generate transitions with a learning rate of 5e-
346 5 with Adafactor optimizer (Shazeer and Stern,
347 2018) and batch size of 16. We train our models
348 for 5 epochs and select the model with lowest loss
349 in the dev set. During decoding, we mix top-K
350 sampling of 80 and top-p (nucleus) sampling of
351 0.95 (Holtzman et al., 2020).

352 **Dialogue Generation** To generate task-oriented
353 utterances, we train our two simulators on the
354 model BlenderBot-400M⁴ with the same training
355 setting with generative transition. During genera-
356 tion, we use top-K sampling (Fan et al., 2018) of
357 120 for more diversity in our dialogues.

358 3.2 Crowdsourcing Tasks

359 We randomly pick about 4,000 dialogues for two
360 human evaluation tasks submit to AMT. The first
361 task is designed for collecting feedback about the

⁴<https://parl.ai/projects/blenderbot2/>

entire dialogue, while the second one focuses on the *transition* part due to the main goal of this paper. Table 2 shows the dialogue example presented to crowdworkers for evaluation. Because our target dialogues can be viewed as salespersons’ capability, we describe the context to the recruited crowdworkers that the given dialogues are between a beginner salesperson and his/her customer and ask workers to provide feedback from different aspects in terms of the sales’ strategies. Note that the annotators are not aware that the dialogues are machine-generated, so the collected feedback is for human sales communication skills. Each task is briefly described below, and the full annotation guideline can be found in the Appendix A.

Task 1: Salesperson-Customer Conversation

The workers were presented with one entire dialogue and asked to rate (from 1 to 5) the entire conversation from three aspects: **Relevance** (Q1—How relevant is the recommended product or service to the conversation context?), **Aggressiveness** (Q2—How aggressive is the salesperson’s communication strategy?), and **Overall** (Q3—Do you think the sales conversation is overall a good example of making a sales recommendations?).

Task 2: Chit-Chat to Task-Oriented Transition

In addition to the entire dialogue, we also explicitly highlight the transition turn in the dialogue when presenting to crowdworkers. Similarly to the first task but only focusing on the transition part, we asked workers to rate from 1 to 5 from the following aspects: **Right Time** (Q1—Is it a good timing to make the transition?), **Relevance** (Q2—Is the transition relevant to the conversation context?), **Aggressiveness** (Q3—Is the transition aggressive?), and **Overall** (Q4—Do you think it is overall a good transition?). In each question, the detailed descriptions of all ratings are given to crowdworkers to ensure they have consistent understanding for all ratings. In addition, to enrich the transition turns and ensure their quality, we generate 4 additional transitions and ask workers to choose the best one. All transitions and ratings are included in our released data.

4 Results and Analysis

Table 3 presents the statistics of the randomly sampled dialogues submitted to AMT. For brevity, we use T to denote *Task* in the following. Each dialogue is evaluated by three crowdworkers so that we can check the annotation variance for reliable

Intent	#Dialogues	Avg Length
FindMovies	1,792	18
GetTimesForMovie	30	19
FindAttractions	1,296	16
LookupMusic	490	16
PlaySong	300	15
LookupSong	8	18
Merge SGD	2,014	21
TOD Simulation	1,902	13
Total	3,916	17

Table 3: Statistics of the sampled dialogues.

results. The evaluation results of all dialogues are visualized in the top charts of Figure 5, and the bottom charts show the results for existing TOD data (Merge) and simulator-generated TOD (Simulator).

It can be observed that our framework is able to produce *context-relevant* task-oriented conversations to match the topic of open-domain dialogues (Q1 in T1; Q2 in T2). This indicates that we can ensure the dialogue flow from open-domain to task-oriented dialogues is natural. The median relevance scores are slightly higher than the *Neutral* line, suggesting that our sales simulator can perform his sales strategy without annoying customers. The observation further demonstrates the feasibility and effectiveness of our proposed method. In terms of the salesperson’s aggressiveness, crowdworkers think that the transition is neutral and somewhat aggressive, showing that smoothly transitioning is still an important research problem to explore. Furthermore, the transition timing scores (Q1 in T2) also demonstrate that our proposed task-oriented intent detection can capture a suitable moment in a zero-shot setting, so that the sales may not miss any business opportunity of product promotion.

We can observe that most of overall scores (Q3 in T1; Q4 in T2) are above *Neutral* (Score 3)⁵, indicating that the generated dialogues and transitions are overall good for a salesperson’s business perspective. The human judgement demonstrates that our proposed approach is capable of simulating a large-scale *reasonable* dialogues aligned with our purpose, implying that both research community and industries can greatly benefit from our released data and the built simulators that can continuously generate more data for training. Our framework and the constructed dataset reduce the cost for large-scale data requirement for better practice.

To further investigate whether the proposed *TOD*

⁵The full description of each score is presented in Appendix A.

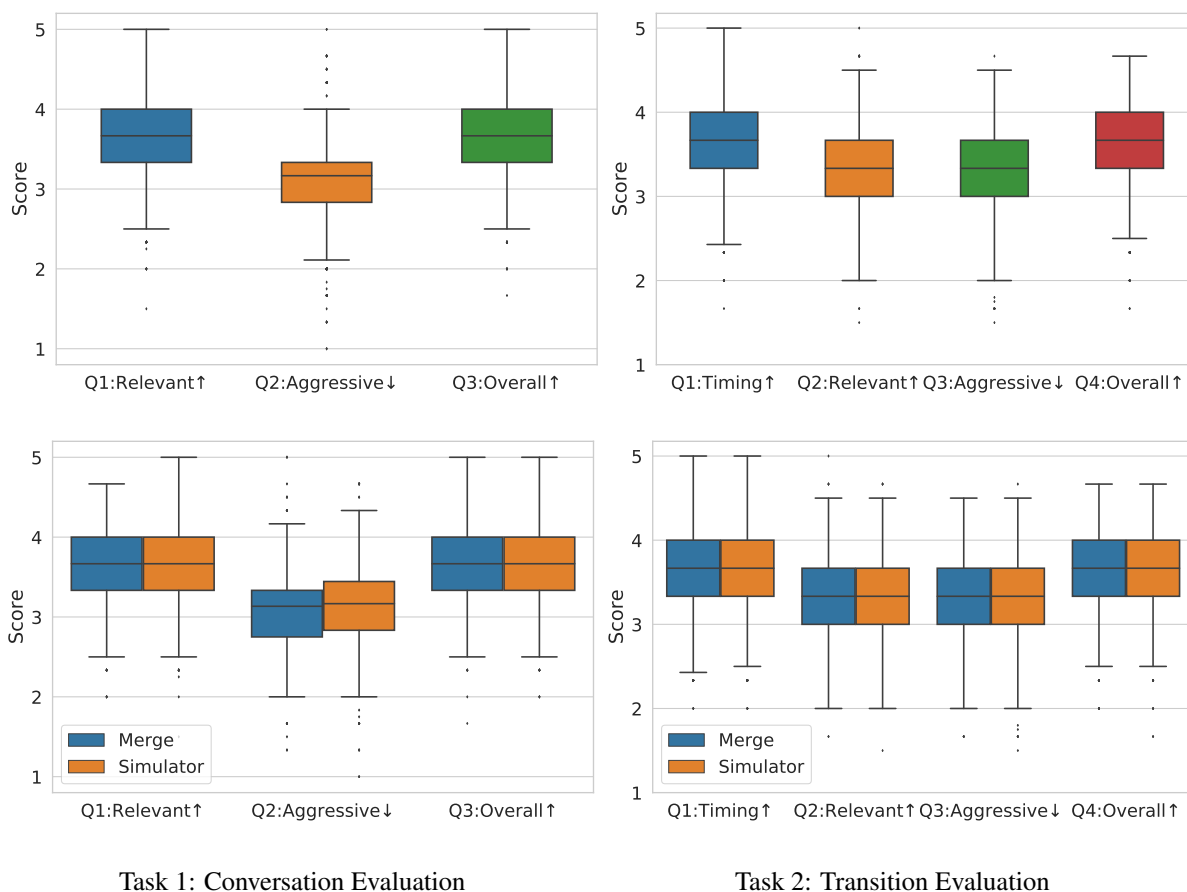


Figure 5: Score distribution of task 1 (left) and 2 (right). The top charts are averaged scores over three workers for all dialogues. The bottom charts are the separated averaged scores where TOD is from *Merge SGD* and *Simulators*. ↑ indicates the higher score the better of this aspect and vice versa for ↓.

450 *simulators* described in Section 2.3 can generate
 451 reasonable dialogues compared to *Merge SGD*, we
 452 visualize their individual scores as shown at the
 453 bottom of Figure 5. There is no significant differ-
 454 ence between two groups, and we further investi-
 455 gate their score distribution of each question
 456 shown in Figure 6. Both results tell that given the
 457 context of open-domain utterances, our TOD simu-
 458 lators are able to generate the suitable task-oriented
 459 dialogues with comparable quality to those from
 460 the publicly available benchmark TOD data—SGD.
 461 Consequently, our framework can be utilized to
 462 generate large-scale data cost-effectively and the
 463 generation quality is comparable with the current
 464 benchmark dialogue data.

465 In addition to the proposed framework and the
 466 released dataset, our collected human judgement
 467 also has the potential of providing valuable contri-
 468 butions to dialogue community or industrial prod-
 469 ucts. For example, each question along with its
 470 corresponding scores from crowdworkers can be
 471 treated as a research problem which may interest re-




Figure 6: The score distribution between *Merge SGD* and *TOD Simulators* in terms of their overall dialogue quality (Q3 of T1).

472 searchers, e.g. how human judges a dialogue is too
 473 aggressive. The release data along with the scores
 474 has a great potential for future research directions.

475 5 Related Work

476 Our work is related to dataset construction for build-
 477 ing persuasive dialogue systems that try to persuade

the participant to take a specific action. [Hiraoka et al. \(2014\)](#) annotated 34 dialogues, in which an experienced salesperson tries to convince a customer to buy a camera. [Yoshino et al. \(2018\)](#) requested crowdsourcing workers to generate 200 persuasive dialogues. In each dialogue, one participant persuaded another one to adopt his suggestion such as cleaning a room. [Wang et al. \(2019\)](#) collected 1017 dialogues, in which one of the participants was convinced to donate to a specific charity. We can see that the covered conversation scenarios in these datasets were strictly limited to specific tasks, while our scenarios are more general and can be easily extended to different cases. Also, our constructed dataset is about three times larger than the prior work, indicating the usefulness of the recent pre-trained paradigm.

The topic of conversational recommendation systems is also related to our work. A number of attempts have been made to collect training data for conversational recommendation systems. These studies ([Wu et al., 2019](#); [Zhou et al., 2020](#); [Xu et al., 2020](#)) first extracted a path consisting of an entity or attribute nodes from a knowledge base. Then they asked annotators to write conversational recommendation dialogues. The flow of mentioned topics in a dialogue should follow the extracted path. Similarly, [Liu et al. \(2020\)](#) also built a dataset by asking human workers to create dialogues based on a topic path. It should be noted that, in these datasets, the goal of such systems is to only make *entity* recommendations instead of *tasks*, while our work goes beyond them in naturally transferring from chit-chat to task-oriented dialogues and completing a task the user may want.

Another related work is generating a transition between two given open-domain utterances. [Tang et al. \(2019\)](#) proposed to generate the transition conditional on a specific word, because they want the generated transition can drive the conversation topic to the specified word. [Sevegnani et al. \(2021\)](#) collected a new dataset of human-created one-turn topic transitions. Each dialogue contains 2 utterances with different topics and 1 transition in the middle of them.

There are some recent studies trying to merge chit-chat and task-oriented dialogues, but the purposes of merged dialogues differ from ours. [Sun et al. \(2021\)](#) enhanced the utterances in task-oriented dialogues by appending chit-chat sen-

tences. They hope that the agent gains more social, personalized, and engaging communication skills. [Ennen et al. \(2021\)](#) proposed a dialogue system that can transfer the style of generated response from chit-chat to task-oriented styles. However, the system is a prototype model, there is still a large gap to properly bridge chit-chat and task-oriented dialogues. The motivation of our work is closely similar to that of the study by [Yu et al. \(2017\)](#). They manually created several task-oriented response generation strategies specifically designed for the movie promotion scenario. In addition, the expert knowledge was utilized to design reinforcement learning rewards that help their dialogue system to decide which action to take (i.e., continuing chit-chat or selecting a task-oriented strategy to reply).

However, we can observe that the prior work in these studies heavily relied on human efforts (data collection, expert-created strategies, etc.). Therefore, it can be expensive and hard to extend their data or method the practical cases due to the requirement of larger-scale training data. Our proposed framework benefits from the pre-trained models and shows its outstanding conversational capability. The flexibility of extending to diverse cases is also validated, considering that all components inside the framework can be easily substituted by the updated models, and the generated data can be used by semi-supervised or unsupervised methods for cold-start scenarios.

6 Conclusion

This paper proposes a novel framework to generate dialogues that naturally transition from open-domain to task-oriented scenarios at a large scale without heavy human efforts. Our proposed chit-chat to task-oriented transition approach can capture the suitable timing when the user shows the implicit intents and generate the diverse and natural transition turn to trigger the task-oriented utterances. Our human evaluation shows that the automatically generated dialogues have a reasonable quality with natural conversation flows from a business point of view. The released dataset and framework empowers research community to easily obtain large-scale target dialogues and the human annotated scores can be utilized for related work. This paper has a great potential of guiding future research directions and benefiting the community of both research and industry.

References

- 578 Daniel Adiwardana, Minh-Thang Luong, David R. So,
579 Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang,
580 Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu,
581 and Quoc V. Le. 2020a. [Towards a human-like open-domain chatbot.](#)
582
- 583 Daniel Adiwardana, Minh-Thang Luong, David R So,
584 Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang,
585 Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu,
586 et al. 2020b. [Towards a human-like open-domain chatbot.](#) *arXiv preprint arXiv:2001.09977.*
587
- 588 Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang
589 Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ra-
590 madan, and Milica Gašić. 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*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
591
592
593
594
595
596
- 597 Philipp Ennen, Yen-Ting Lin, Ali Girayhan Ozbay, Fer-
598 dinando Insalata, Maolin Li, Ye Tian, Sepehr Jalali,
599 and Da shan Shiu. 2021. [Towards a universal nlg for dialogue systems and simulators with future bridging.](#)
600
601
- 602 Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation.](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
603
604
605
606
607
- 608 Takuya Hiraoka, Graham Neubig, Sakriani Sakti,
609 Tomoki Toda, and Satoshi Nakamura. 2014. [Reinforcement learning of cooperative persuasive dialogue policies using framing.](#) In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1706–1717, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
610
611
612
613
614
615
616
- 617 Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and
618 Yejin Choi. 2020. [The curious case of neural text degeneration.](#)
619
- 620 Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu,
621 Semih Yavuz, and Richard Socher. 2020. [A simple language model for task-oriented dialogue.](#) In *Advances in Neural Information Processing Systems*, volume 33, pages 20179–20191. Curran Associates, Inc.
622
623
624
625
- 626 Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang
627 Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset.](#) In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
628
629
630
631
632
- 633 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*, pages 1036–1049, Online. Association for Computational Linguistics.
634
635
636
637
638
639
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization.](#) In *International Conference on Learning Representations.*
640
641
642
- A. H. Miller, W. Feng, A. Fisch, J. Lu, D. Batra, A. Bordes, D. Parikh, and J. Weston. 2017. [Parlai: A dialog research software platform.](#) *arXiv preprint arXiv:1705.06476.*
643
644
645
646
- Mahdi Namazifar, Alexandros Papangelis, Gokhan Tur, and Dilek Hakkani-Tür. 2020. [Language model is all you need: Natural language understanding as question answering.](#)
647
648
649
650
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD.](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
651
652
653
654
655
656
657
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. [Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset.](#)
658
659
660
661
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, et al. 2021. [Recipes for building an open-domain chatbot.](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325.
662
663
664
665
666
667
668
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.](#)
669
670
671
- Karin Sevegnani, David M. Howcroft, Ioannis Konstas, and Verena Rieser. 2021. [OTTERS: One-turn topic transitions for open-domain dialogue.](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2492–2504, Online. Association for Computational Linguistics.
672
673
674
675
676
677
678
679
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost.](#) *CoRR*, abs/1804.04235.
680
681
682
- Eric Michael Smith, Mary Williamson, Kurt Shuster, Jason Weston, and Y-Lan Boureau. 2020. [Can you put it all together: Evaluating conversational agents’ ability to blend skills.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2021–2030, Online. Association for Computational Linguistics.
683
684
685
686
687
688
689

690	Kai Sun, Seungwhan Moon, Paul Crook, Stephen	Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur	748
691	Roller, Becka Silvert, Bing Liu, Zhiguang Wang,	Szlam, Douwe Kiela, and Jason Weston. 2018. <i>Per-</i>	749
692	Honglei Liu, Eunjoon Cho, and Claire Cardie. 2021.	sonalizing dialogue agents: I have a dog, do you	750
693	<i>Adding chit-chat to enhance task-oriented dialogues.</i>	have pets too?	751
694	In <i>Proceedings of the 2021 Conference of the North</i>		
695	<i>American Chapter of the Association for Computa-</i>	Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen,	752
696	<i>tional Linguistics: Human Language Technologies,</i>	Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing	753
697	pages 1570–1583, Online. Association for Computa-	Liu, and William B Dolan. 2020. Dialogpt: Large-	754
698	tional Linguistics.	scale generative pre-training for conversational re-	755
		sponse generation. In <i>Proceedings of the 58th An-</i>	756
699	Alon Talmor, Jonathan Herzig, Nicholas Lourie, and	<i>Annual Meeting of the Association for Computational</i>	757
700	Jonathan Berant. 2019. Commonsenseqa: A ques-	<i>Linguistics: System Demonstrations</i> , pages 270–	758
701	tion answering challenge targeting commonsense	278.	759
702	knowledge. In <i>Proceedings of the 2019 Conference</i>		
703	<i>of the North American Chapter of the Association</i>	Kun Zhou, Yuanhang Zhou, Wayne Xin Zhao, Xi-	760
704	<i>for Computational Linguistics: Human Language</i>	aoke Wang, and Ji-Rong Wen. 2020. Towards topic-	761
705	<i>Technologies, Volume 1 (Long and Short Papers),</i>	guided conversational recommender system. In <i>Pro-</i>	762
706	pages 4149–4158.	<i>ceedings of the 28th International Conference on</i>	763
		<i>Computational Linguistics</i> , pages 4128–4139.	764
707	Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiao-		
708	dan Liang, Eric Xing, and Zhiting Hu. 2019. <i>Target-</i>	A Crowdsourcing Guideline	765
709	<i>guided open-domain conversation.</i> In <i>Proceedings</i>		
710	<i>of the 57th Annual Meeting of the Association for</i>	A.1 Task1: Conversation between A	766
711	<i>Computational Linguistics</i> , pages 5624–5634, Flo-	Salesperson and A Customer	767
712	rence, Italy. Association for Computational Linguis-		
713	tics.	In order to improve the skills to sell more prod-	768
		ucts, a beginner salesperson is learning dialogue	769
714	Xuwei Wang, Weiyang Shi, Richard Kim, Yoojung Oh,	strategies by reading prior conversations between	770
715	Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. <i>Per-</i>	customers and other salespeople. This beginner	771
716	<i>suasion for good: Towards a personalized persuasive</i>	salesperson needs your help to determine if a sales-	772
717	<i>dialogue system for social good.</i> In <i>Proceedings of</i>	person used a good dialogue strategy to conduct an	773
718	<i>the 57th Annual Meeting of the Association for Com-</i>	effective and strategic sales conversion.	774
719	<i>putational Linguistics</i> , pages 5635–5649, Florence,		
720	Italy. Association for Computational Linguistics.	In more detail, you will be presented with one	775
		conversation history between a salesperson and	776
721	Zhiguo Wang, Wael Hamza, and Radu Florian. 2017.	a customer. The salesperson may recommend a	777
722	<i>Bilateral multi-perspective matching for natural lan-</i>	movie, a song, attractions and so on for the cus-	778
723	<i>guage sentences.</i>	tom. Instead of recommending a product or ser-	779
		vice to the customer directly, the salesperson wants	780
724	Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu,	to make the recommendation more gradually and	781
725	Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang.	naturally by starting the conversation with chit-	782
726	2019. <i>Proactive human-machine conversation with</i>	chat.	783
727	<i>explicit conversation goal.</i> In <i>Proceedings of the</i>		
728	<i>57th Annual Meeting of the Association for Com-</i>	In this task, you need to rate the conversation	784
729	<i>putational Linguistics</i> , pages 3794–3804, Florence,	from the following 3 aspects:	785
730	Italy. Association for Computational Linguistics.		
		• How relevant is the recommended product or	786
731	Jun Xu, Haifeng Wang, Zhengyu Niu, Hua Wu, and	service to the conversation context?	787
732	Wanxiang Che. 2020. Knowledge graph grounded	• How aggressive is the salesperson’s commu-	788
733	goal planning for open-domain conversation gener-	nication strategy?	789
734	ation. In <i>Proceedings of the AAAI Conference on</i>	• Do you think the sales conversation is overall	790
735	<i>Artificial Intelligence</i> , volume 34, pages 9338–9345.	a good example of making sales recommenda-	791
		tions?	792
736	Koichiro Yoshino, Yoko Ishikawa, Masahiro	Questions	793
737	Mizukami, Yu Suzuki, Sakriani Sakti, and Satoshi		
738	Nakamura. 2018. <i>Dialogue scenario collection</i>	• How relevant is the recommended product or	794
739	<i>of persuasive dialogue with emotional expres-</i>	service to the conversation context?	795
740	<i>sions via crowdsourcing.</i> In <i>Proceedings of the</i>		
741	<i>Eleventh International Conference on Language</i>		
742	<i>Resources and Evaluation (LREC 2018)</i> , Miyazaki,		
743	Japan. European Language Resources Association		
744	(ELRA).		
745	Zhou Yu, Alan W Black, and Alexander I. Rudnicky.		
746	2017. <i>Learning conversational systems that inter-</i>		
747	<i>leave task and non-task content.</i>		

796	– 1: Not at all (it is impossible for me to find the relevance between the recommended item and the context)	can be ignored; the beginner salesperson should keep this example in his mind)	845
797			846
798		– 5: Very good (This is a perfect example of making recommendations; the beginner salesperson should keep it deeply in his mind)	847
799	– 2: Less than neutral (it is a bit hard for me to find the relevance between the recommended item and the context)		848
800			849
801			850
802	– 3: Neutral (With some effort, I can find a reasonable relevance between the recommended item and the context)		851
803		A.2 Task2: Transition between Chitchat and Task-Oriented	852
804			
805	– 4: Relevant (I can easily find that the recommend item has obvious relevance with the context, even though the recommended item is not perfectly matching the context)	In order to improve the skills to sell more products, a beginner salesperson is learning dialogue strategies by reading prior conversations between customers and other salespeople. This beginner salesperson needs your help to determine if a salesperson used a good dialogue strategy to conduct an effective and strategic sales conversion.	853
806			854
807			855
808	– 5: Very Relevant (the recommended item is perfectly matching the context)		856
809			857
810			858
811			859
812	• How aggressive is the salesperson’s communication?	You will be presented with a conversation between a salesperson and a customer. The salesperson may recommend a movie, a song, attractions and so on for the customer. Instead of recommending a product or service to the customer directly, the salesperson wants to make the recommendation more gradually and naturally by starting the conversation with chit-chat. Once the salesperson thinks it is the right time, he will say something (named transition in this task) to change the conversation from chit-chat to recommendation-making.	860
813			861
814	– 1: Not aggressive at all (the conversation flows very naturally and smoothly from chit-chat to making recommendations; If I was the customer, I feel very comfortable when the salesperson is making recommendations)		862
815			863
816	– 2: Less than neutral (The flow of the conversation is generally natural and smooth, although there are few imperfections)		864
817			865
818	– 3: Neutral (The salesperson starts to recommend an item; It is ok to me)		866
819			867
820	– 4: Aggressive (The salesperson suddenly starts to recommend an item; this makes me a bit uncomfortable)		868
821			869
822	– 5: Very aggressive (The salesperson suddenly starts to recommend an item; this makes me very uncomfortable)		870
823			
824		In this task, you will need to rate the transition from the following 4 aspects:	871
825			872
826		• Is it the right time to make the transition?	873
827		• Is the transition relevant to the conversation context?	874
828		• Is the transition aggressive?	875
829		• Is the transition overall good?	876
830			877
831	• Is the sales conversation overall a good example to the beginner salesperson?	Questions	878
832			
833	– 1: Not at all (This example is really very bad; the beginner salesperson should not spend time on learning this example)	• Is it the right time to make the transition?	879
834		– 1: Very bad time (This is definitely not the right time to do it. It is highly likely that the customer will find you very annoying)	880
835	– 2: Less than neutral (This example is not good; it would not be a pity if the beginner salesperson skips it)		881
836		– 2: Bad time (This is not a good time to make the transition. It may cause negative customer feelings)	882
837	– 3: Neutral (This is not a bad example; the beginner salesperson may learn some useful dialogue skills from it, but not very much)		883
838		– 3: Neutral (I don’t think making the transition at the time is good, but it is ok to me to continue the conversation if I was the customer)	884
839	– 4: Good (This is a good example of making recommendations; the imperfections		885
840			886
841			887
842			888
843			889
844			890

891	– 4: Good time (it is a good time to make the transition, but maybe it will be perfect if the transition is made earlier or later)	– 2: Less than neutral (This transition is not good; It would not be a pity if the beginner salesperson skips this example)	936
892			937
893		– 3: Neutral (This is not a bad transition; the beginner salesperson may learn some useful dialogue skills from it, but not very much)	938
894			939
895	– 5: Very good time (it is a perfect time to make the transition)		940
896			941
897	• Is the transition relevant to the conversation context?	– 4: Good (This is a good example of making a transition; the imperfections can be ignored; the beginner salesperson should keep this example in his mind)	942
898			943
899	– 1: Not at all (it is impossible for me to find the relevance between the transition and the context)	– 5: Very good (This is a perfect example of making a transition; the beginner salesperson should keep it deeply in his mind)	944
900			945
901	– 2: Less than neutral (it is a bit hard for me to find the relevance between the transition and the context)		946
902			947
903	– 3: Neutral (With some effort, I can find a reasonable relevance between the transition and the context)	• Which transition of the following do you think is the best?	948
904		– <i>transition 1</i>	949
905	– 4: Relevant (I can easily find that the transition has obvious relevance with the context, even though the transition is not perfectly matching the context)	– <i>transition 2</i>	950
906		– <i>transition 3</i>	951
907	– 5: Very Relevant (the transition is perfectly matching the context; it is hard for me to find a better transition)	– <i>transition 4</i>	952
908			953
909	• Is the transition aggressive?		954
910	– 1: Not aggressive at all (the conversation flows very naturally and smoothly from chit-chat to making the transition; If I was the customer, I feel very comfortable when the salesperson is doing it)		955
911			956
912	– 2: Less than neutral (The flow of the conversation is generally natural and smooth, although there are few imperfections)		
913	– 3: Neutral (The salesperson starts to make the transition; It is ok to me)		
914	– 4: Aggressive (The salesperson suddenly starts to make the transition; this makes me a bit uncomfortable)		
915	– 5: Very aggressive (The salesperson suddenly starts to make the transition; this makes me very uncomfortable)		
916			
917	• Is the transition overall good?		
918	– 1: Not at all (This transition is really very bad; the beginner should not spend time on leaning this transition)		
919			
920			
921			
922			
923			
924			
925			
926			
927			
928			
929			
930			
931			
932			
933			
934			
935			

B Generated Dialogue Samples

Task-oriented Simulators

Sales: Oh, I'm sorry to hear that. Do you have any other family members that you are close to?

User: Yes, I have two brothers and two sisters. We are all very close. How about you?

Sales: I don't have any siblings, but I have a lot of cousins that I'm very close to.

User: That's great. I'm glad you have a good relationship with your cousins. What do you like to do for fun?

Sales: I like to play video games and watch movies. What about you, what do you enjoy doing?

User: I like playing video games as well. I also like to go to the movies. I love going to the theater. - **FindMovies**

Sales: *Sure, what kind of movies do you like to watch? Can you see [movie_name]? - [Transition]*

User: I would love to see Captain Marvel. I'd like to watch it at Century at Tanforan.

Sales: I found [count] movie you might like. How about [movie_name]?

User: Yeah, Hotel Mumbai sounds good.

Sales: You want to play [title] without subtitles, is that right?

User: Yes, that's right.

Sales: Okay, so you want me to play the movie you requested. I'll start the movie now.

User: Yes, please do.

Sales: Can you please confirm that you want to play [title] without subtitles?

Merge SGD

User: I like to drive around in my black car and listen to music. What do you like to do?

Sales: I love going to the beach and swimming. What kind of music do you listen to?

User: I listen to a lot of rap and hip hop. Do you have a favorite rapper? - **LookupMusic**

Sales: *I'm a top notch rapper and I'm a rock musician. Are you in the mood for the music. - [Transition]*

User: I'm in the mood for some music. Can you find songs from the album Camila.

Sales: What about the song [song_name] from the album [album] by [artist]?

User: I'm not in the mood for that one, do you have a different song?

Sales: What about the song [song_name] from the album [album] by [artist]?

User: What is the genre of the song?

Sales: Real Friends is a [genre] song.

User: Sounds good.

Sales: Do you want to hear the song now?

User: Yes, I would love to.

Sales: You want to play [song_name] on your [playback_device]?

User: No, I would prefer to play it on my kitchen speaker.

Sales: You want to play [song_name] on your [playback_device]?

User: Yes, that's right.

Sales: The song is now playing.

User: Thank you very much.

Sales: Is there anything else I can help you with?

User: No, that is all. Thank you for your time.

Sales: Enjoy your music. Have a wonderful day.

Merge SGD

User: London is a great place to visit, but I wouldn't want to live there. Too many people.
Sales: I know what you mean. London is the most populous city in the United Kingdom.
User: It's also the most densely populated metropolitan area in the world.
Sales: Yes, it is. It has a population of 8,537,673.
User: Wow, that's a lot of people. Do you live in the city or the suburbs?
Sales: I live in London, the capital of the UK. It's a great city to live in.
User: I've never been there but I've always wanted to go. I've heard it's the cultural, financial, and media capital of Europe.
Sales: Yes, that is true. It is also one of the world's most visited cities as well.
User: I would love to go there and see all of the museums and historical sights. -
FindAttractions
Sales: *Sure. What city? What kind of attraction are you looking for? - [Transition]*
User: I'm looking for attractions to visit
Sales: In which city do you want to search for?
User: I would like attractions in San Francisco
Sales: I found [count] attractions, for example [attraction_name], a [category]
User: Is it a child-friendly place? Is entry free?
Sales: The entry is free and it's a good place for kids
User: Very good! Can you give me any other suggestions? Any attraction is fine
Sales: I found [count] attractions. I suggest [attraction_name], a [category]
User: This is fine too
Sales: You need anything else?
User: Yes, I'd like to watch a movie there, preferably a 3d show
Sales: There are [movie_name], [movie_name] and [movie_name].
User: I'd like other proposals
Sales: There are [movie_name], [movie_name] and [movie_name]
User: I like IT two! What online movies do you propose?
Sales: What kind of movie do you prefer?
User: I love Suspense movies
Sales: I found [count] movies: do you prefer [title], [title] or [title]?
User: I prefer The Art of Self-Defense! I want to watch it now with English subtitles
Sales: Playing [title] with [subtitle_language] subtitles: do you confirm?
User: No, I prefer with French subtitles, thanks!
Sales: Playing [title] with [subtitle_language] subtitles: do you confirm?
User: Now it's correct, I confirm.
Sales: The movie is starting!
User: Who is the main actor?
Sales: The mai actor is [starring].
User: Thank you! It's all for now!

Table 4: These are more examples generated by different strategies in our framework. The text has the underline is the detected intent. The sentence in *italics* is the transition.