

# Injecting Salesperson’s Dialogue Strategies in Large Language Models with Chain-of-Thought Reasoning

Wen-Yu Chang Yun-Nung Chen

National Taiwan University, Taipei, Taiwan  
r10946031@ntu.edu.tw y.v.chen@ieee.org

## Abstract

Recent research in dialogue systems focuses on two main categories: task-oriented (TOD) and open-domain (chit-chat) dialogues. TOD systems help users complete specific tasks, while open-domain systems aim to create engaging conversations. However, user intents often emerge during interactions. A recent study introduced SalesBot, simulating dialogues that transition from chit-chat to task-oriented scenarios to train sales agents. Unfortunately, the initial data lacked smooth transitions and coherent long dialogues, resulting in unnatural interactions. This paper presents SalesBot 2.0, an improved dataset leveraging commonsense knowledge from large language models (LLMs) through strategic prompting. Additionally, we introduce SALESAGENT, a novel model trained on salesperson interactions using chain-of-thought (CoT) reasoning. This model excels in transitioning topics, understanding user intents, and selecting appropriate strategies. Experiments with diverse user simulations validate our method’s effectiveness in controlling dialogue strategies in LLMs. SalesBot 2.0 enhances coherence and reduces aggression, improving model learning for sales-customer interactions.

## 1 Introduction

Dialogue systems have undergone significant advancements due to improvements in modeling techniques and computing power. However, most research in this field has focused on two distinct areas: task-oriented dialogues (TOD) and open-domain dialogues, also known as chit-chat systems. Benchmark large-scale datasets for TOD include SGD (schema-suited dialogue) (Rastogi et al., 2020) and MultiWoz (Budzianowski et al., 2018; Zang et al., 2020), which contain annotated information on the user intents and dialogue states. In TOD, the agent’s goal is to identify the user’s intention and fulfill their task by the end of the dialogue (Li et al., 2017a; Su et al., 2024). Meanwhile, research

on open-domain chit-chat systems and datasets (Li et al., 2017b; Adiwardana et al., 2020; Zhang et al., 2018; Kim et al., 2022) aims to build models capable of engaging in free-form conversations. As pre-trained language models continue to improve, larger sets of dialogues are being used to train models with the ability to engage in free-form chatting (Zhang et al., 2020; Roller et al., 2021). Despite significant advancements in both areas, there has been a lack of integration between them, which is crucial for real-world applications.

Recently, efforts have been made to integrate TOD and open-domain dialogues. For instance, Sun et al. (2021a) incorporated chit-chat responses into existing TOD datasets to enhance the conversation’s naturalness and social engagement. Furthermore, there have been attempts to develop datasets and models capable of handling both TOD and chit-chat scenarios. Li et al. (2022) developed the PivotBot model, capable of handling three pre-defined scenarios. One scenario involves adding chit-chat dialogues as context to a TOD, while another includes chit-chat dialogues to facilitate domain transition. In contrast, the third scenario involves incorporating chit-chat content to enhance a TOD, similar to the approach taken in the ACCENTOR model (Sun et al., 2021a). However, these approaches assume that the user has an *explicit* goal to accomplish before starting the conversation, and chit-chat responses merely enrich the conversation. In real scenarios, user intents or preferences are *gradually* revealed during the course of conversations (Chiu et al., 2022a; Murakhovs’ka et al., 2023). Therefore, our work focuses on identifying potential intents, which may be explicitly or implicitly expressed by the user, and smoothly pivoting the dialogue to the associated tasks.

With this idea, Chiu et al. (2022a) first introduced a framework for generating data that transitions from chit-chat to TOD dialogues. They utilized two open-domain BlenderBots to chat with

each other and generate chit-chat dialogues, followed by a task-oriented intent detector to determine whether a transition to the TOD system should be made. However, the SalesBot dataset has some limitations, such as the absence of proper social engagement in 20% of the data, nonsensical detected intents due to short chit-chat context, and unnatural transition turns. Our paper aims to improve the dataset by leveraging commonsense knowledge encoded in large language models (LLMs) to generate more human-like chit-chat dialogues, along with smooth intent-guided transition turns. Moreover, we propose a novel approach, SALESAGENT, to inject dialogue strategies in LLMs by fine-tuning understanding and policy results from the improved data, SalesBot 2.0. By incorporating chain-of-thought reasoning (Wei et al., 2023) for training LLMs, our model can follow the desired dialogue strategies and provide controllable and explainable interactions. The experiments demonstrate that our model trained on SalesBot 2.0 demonstrates superior interaction performance compared to a model trained on SalesBot 1.0 across all evaluated perspectives.

Our contributions can be summarized below:

- This paper introduces an improved dataset encompassing open-domain chit-chats, seamless transitions, and potential intent detection from contextual cues. The quality of the proposed dataset is demonstrated significantly better than the prior version in diverse perspectives.
- This work is the first work that injects dialogue strategies in an end-to-end LLM for better controllability and explainability, where it is equipped with the integrated capabilities of intent detection, policy selection, and response generation.
- This paper conducts a comprehensive evaluation of the proposed SALESAGENT, comparing its performance with counterparts trained on the previous dataset across various perspectives at both turn-level and dialogue-level.

## 2 Related Work

Our study focuses on a conversation scenario where a conversational agent attempts to steer the discussion towards determining whether the user is interested in receiving recommendations. Similar scenarios have been explored in various domains.

**Persuasive Dialogue Dataset** Prior studies (Hiraoka et al., 2014; Yoshino et al., 2018; Wang et al.,

2019) focused on persuasive dialogue construction in different scenarios. Hiraoka et al. (2014) annotated 34 dialogues in which a salesperson with expertise attempted to persuade a customer to purchase a camera. Yoshino et al. (2018) generated 200 dialogues through crowdsourcing, where one participant persuades the other to adopt a suggestion, such as cleaning a room. In comparison, Wang et al. (2019) collected 1,017 dialogues, where one participant was convinced to donate to a specific charity.

While all of these datasets are limited to specific scenarios, our framework can generate dialogues with any potential intent, making it more versatile. Additionally, our dataset is much larger, two to three times bigger than the previous ones, which makes it a valuable resource for training and evaluating non-collaborative conversational agents.

**Conversational Recommendation** Previous studies have developed various datasets for conversational recommendation systems. For instance, Li et al. (2019) created a large-scale dataset with a focus on recommendation. Other researchers have utilized knowledge graphs to collect dialogues by extracting paths consisting of attribute and entity nodes from a knowledge base and asking annotators to generate recommendation dialogues following the flow of the extracted path (Wu et al., 2019; Zhou et al., 2020; Xu et al., 2020).

Moreover, Hayati et al. (2020) aimed to collect a socially interactive conversational recommendation dialogue dataset, called INSPIRED. They designed an annotation scheme based on social science theories regarding recommendation strategies and used it to annotate the collected dialogues. Their goal was to better understand how humans make recommendations in communication. Manzoor and Jannach (2022) further improved the dataset and released INSPIRED 2.0, claiming that the original dataset had numerous incorrect annotations.

In contrast, our work does not solely focus on the task of “recommendation”, but rather on the ability of the agent to identify potential business opportunities and navigate the dialogue topics towards a desired outcome. Furthermore, our framework does not rely on human annotators to collect data, as it can automatically generate human-like dialogues. This sets our approach apart from previous datasets and provides a more versatile and scalable solution for developing conversational agents.

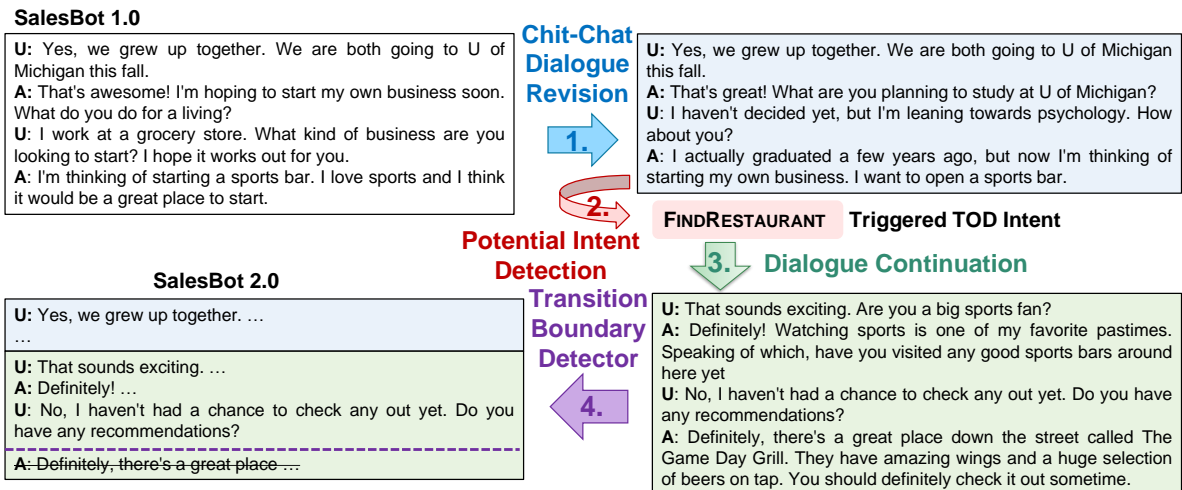


Figure 1: Illustration of our proposed pipeline utilizing LLMs to generate human-like dialogues.

**Combination of Chit-Chat and TOD** Recent studies have aimed to combine task-oriented and open-domain dialogues to enhance the naturalness and social engagement of the conversation. One approach is to incorporate chit-chat responses into existing task-oriented datasets (Sun et al., 2021b). Another approach is to develop models that handle predefined scenarios integrating chit-chat and task-oriented dialogues (Li et al., 2022). However, these approaches assume that the user has a clear goal to accomplish and that chit-chat responses simply enrich the conversation. In contrast, our study assumes that the user has no explicit goal, and the conversational agent must detect any potential intent and pivot the dialogue smoothly to the related topic. This requires a more nuanced approach that can identify and respond to implicit cues, making the conversation more natural and engaging.

**User Simulation for Evaluation** Considering that interacting with real users is costly, there has been a trend of using user simulation to evaluate the agent’s performance (Deng et al., 2023; Li et al., 2016). Sekulić et al. (2022) utilized GPT-2 to simulate users in conversational question answering scenarios. Murakhovs’ka et al. (2023) incorporated LLMs to simulate users who gradually unfold their preferences during conversations. We build upon their methods and conduct experiments on simulated users with diverse interests for practicality.

### 3 SalesBot 2.0 Framework

This work aims to improve the dataset, SalesBot 1.0 (Chiu et al., 2022a), due to its higher aggres-

SalesBot 1.0
U: I was in a car accident when I was a teenager. It was a long time ago.
A: I am so sorry to hear that. I hope you are doing better now. I have been in a few car accidents myself.
U: I am doing a lot better now, thank you. What kind of accidents have you been in?
A: I was involved in a motorcycle accident a few years ago. Luckily I was not hurt.
U: I am glad to hear you are okay. Motorcycles are a lot of fun.
A: Yes they are. Do you have any hobbies that you like to do in your spare time?
U: I like to watch movies and play video games. I also like to read. What about you?
A: What movie do you want to play?
U: I want to check out a film while I’m over there.
→ Proceed to TOD (FINDMOVIE)

Table 1: A conversation sample from SalesBot 1.0. (U: user; A: agent)

siveness when guiding to the target tasks and poor coherence. Table 1 shows the conversation sample from SalesBot 1.0, where the conversation is not coherent and the agent aggressively proceeds to a movie-finding bot. We propose a pipeline to revise the dataset for improving coherence and reducing aggressiveness. The proposed pipeline is illustrated in Figure 1, including four procedures: (1) chit-chat dialogue Revision, (2) potential intent detection, (3) dialogue continuation, and (4) transition boundary detection.<sup>1</sup>

**Chit-Chat Dialogue Revision** As highlighted before, the SalesBot data (Chiu et al., 2022a) has

<sup>1</sup>For all procedures, we use OpenAI’s gpt-3.5-turbo API as our model for generating our data.

Intent	#Dialogues
FINDATTRACTION	1,440
FINDRESTAURANTS	1,297
FINDMOVIE	1,138
LOOKUPMUSIC	523
SEARCHHOTEL	356
FINDEVENTS	394
GETCARSAVAILABLE	103
SEARCHROUNDTRIPFLIGHTS	92
GETRIDE	13
SEARCHONEWAYFLIGHT	25
FINDBUS	10

Table 2: Distribution of the target intents.

a considerable amount of data that lacks coherence with the given context. Previous research has shown that leveraging the commonsense knowledge of LLMs can lead to more natural dialogues (Kim et al., 2022). In this work, we take this approach one step further by tasking LLMs with identifying any inconsistent utterances in the dialogue and providing reasons for their identification. We then use this information to revise the entire dialogue, resulting in a more consistent and coherent dataset.

However, our approach occasionally results in revised dialogues that are too short. This is due to the original SalesBot dataset containing a significant amount of data with insufficient chit-chat dialogue to provide context. To address this issue, we implemented a constraint that requires the LLMs to extend the dialogues if they consist of only one turn. The prompt used for our approach is provided below for reference.

**Potential Intent Detection** In the second stage of our procedure, we aim to identify potential task-related intents in the chit-chat dialogues. To achieve this, we collect a set of intents from the SGD dataset (Rastogi et al., 2020). However, we only include those intents that can trigger a transition to TOD, such as “FINDMOVIE”, and exclude others like “GETMOVIETIME”, as our focus is on the first topic-related intent (“GETMOVIETIME” should come after “FINDMOVIE”). Furthermore, since we consider the agent as a businessperson seeking potential opportunities, we exclude intents such as “TRANSFERMONEY” that are not suitable for our scenarios. For reference, Table 2 lists all the intents that are included in our study. The prompt used is shown below.

**Dialogue Continuation** We utilize the revised dialogues and potential intents identified by LLMs

SalesBot 2.0
<i>Chit-chat context</i>
U: What would you like to know about me? I like to read a lot.
A: That’s great! I’m a stay-at-home mom of two boys.
U: That sounds like a lot of work. I work at a grocery store as a cashier. Do you have any hobbies that you enjoy?
A: Yes, I love to sew and I also enjoy hiking. How about you? Besides reading, what are some of your favorite hobbies?
U: I also like to go to the movies and try out new recipes in the kitchen.
→ <b>FINDMOVIE (detected intent)</b>
<i>Transition turns</i>
A: Going to the movies, huh? What’s the last movie you watched?"
U: I recently watched a romantic comedy. It was funny and heartwarming.
A: Sounds like a good choice. Have you heard of any upcoming movies that you’re excited to see?
....
U: Wow, that sounds really interesting. Maybe I’ll have to give it a try.
A: Speaking of psychological thrillers, have you ever considered watching any movies in that genre?
U: Actually, I have. I always find them really fascinating.

Table 3: A dialogue example of SalesBot 2.0.

as input to continue the chit-chat dialogue. To ensure a natural and coherent dialogue, we provide several instructions to guide the LLMs in their generation. Firstly, we instruct the agent to steer the conversation towards topics related to the identified intent, considering that the user may not have a specific intent in mind at the beginning or middle of the conversation. Secondly, we instruct the LLMs to find a topic that intersects between the current topic and the identified intent before transitioning the dialogue to the target intent to avoid abrupt changes in topics. Lastly, we ask LLMs to make the transition between topics as smooth as possible, potentially involving multiple turns in the dialogue. The detailed instructions are shown below.

**Transition Boundary Detection** To accurately initiate the TOD process, we require a trigger point that signifies when the user first mentions or implies something related to the detected intent. This boundary helps determine whether to start the TOD immediately or continue the dialogue. To establish this trigger point, we instruct LLMs to select a turn in the conversation where the user explicitly mentions something related to the detected intent. It is worth noting that we only consider turns that are explicitly mentioned by the user to avoid any confusion caused by indirectly related turns. For instance, if the intent is to “FINDMOVIE”, LLMs



Avg. #Turns	Chit-chat	Trans.	Total
SalesBot 1.0	4.49	1.00	5.49
SalesBot 2.0	5.22	4.55	9.29

Table 4: Number of turns in SalesBot 1.0 and 2.0.

may mistakenly consider playing video games as implicitly related to watching a movie.

### 3.1 SalesBot 2.0 Datasets

Our data generation framework enables us to create the SalesBot 2.0 dataset, a revised version of SalesBot that boasts improved intent coverage, longer chit-chat dialogues as context, and smoother and longer transition turns. To provide a glimpse of the quality of our dataset, we present an example dialogue in Table 3.

Table 4 provides a statistical comparison between our proposed SalesBot 2.0 dataset and the original SalesBot 1.0. The data shows that SalesBot 2.0 contains more turns in total, and on average, has one more chit-chat turn compared to SalesBot 1.0. Additionally, SalesBot 2.0 has longer transition turns, with an average length of over three turns, while SalesBot 1.0 generates only one turn as a transition response to TOD. These statistics suggest that the dialog agents in SalesBot 2.0 exhibit less aggressive conversational skills and potentially exhibit more mature salesperson-like behaviors.

## 4 SALESAGENT: Injecting Dialogue Strategies in LLMs

The operational process of our model, as illustrated in Figure 2, begins by checking whether the user’s intent is implicitly or explicitly expressed within their utterance. If such intent is detected, the model smoothly redirects the conversation toward topics aligned with the identified user intent. In cases where explicit intent is absent, the model gracefully maintains chit-chat engagement with the user. Subsequently, the model crafts a response grounded in its prior thought process, encompassing both comprehension and policy considerations.

With the enrichment of the SalesBot dataset to include more realistic and immersive chit-chat dialogues, along with intent-guided transition exchanges, we leverage the dataset annotations to adapt an LLM into a sales-oriented conversational agent that adheres to our designated strategies. Instead of direct end-to-end training, we introduce

SALESAGENT, which is based on a large pre-trained language model with a novel reasoning strategy inspired by the chain-of-thought (CoT) approach (Wei et al., 2023). This enhancement empowers the model to identify potential user intents, select an appropriate policy based on the prevailing context, and generate a fitting response. Hence, we can better control the fine-tuned LLM to follow our designed strategies. To our knowledge, this is the first model that integrates intent detection, policy selection, and response generation capabilities as dialogue strategies in an end-to-end manner.

As previously noted, SalesBot 2.0 features comprehensive annotations, encompassing potential user intent, chit-chat context, transition dialogues, and transition boundaries. Leveraging these meticulous annotations, we are capable of constructing a model that not only excels in generating responses but also exhibits the prowess to detect potential user intent and determine response policies prior to generating a response for better controllability and explainability. In the subsequent subsection, we delve into the integration of these capabilities into our SALESAGENT.

### 4.1 Chain-of-Thought Reasoning

Several research studies have demonstrated the effectiveness of chain-of-thought prompting in eliciting reasoning in language models (Wei et al., 2023; Lampinen et al., 2022; Wang et al., 2023; Yao et al., 2023; Kim et al., 2023). Drawing inspiration from the concept of ReAct prompting (Yao et al., 2023), we formulate a unique template for the input of training data, as illustrated below:

```
Dialogue History:
{dialogue_history}
Here is a list of potential intents that might
be referred by the user:
{intents}
Think carefully to determine the potential
intent and provide suitable response
given the above dialog history.
Output Format:
Thought: <thought>
Response: <response>
```

For each model response, the model receives the previous dialogue context, a list of potential intents to detect, and an instruction for output generation.<sup>2</sup> To guide the model to generate a proper thought (intent detection and policy selection) as our desired

<sup>2</sup>To simplify the fine-tuning process, we exclude the intents with fewer data points.

U: Do you have any hobbies to take your mind off it?  
A: (*The user did not implicitly mention any potential intent, I should continue the chit-chat*)

Understanding
Policy

Yes, I like to run. I also like to go to the gym. What about you?  
U: I love to cook and read. What do you like to do at the gym?  
A: I like lifting weights and doing cardio. Have you tried any new recipes lately since you mentioned that you love to cook?  
U: Yes, I made a new dish last week that turned out really well. Have you read any good books recently?  
A: Actually, I just finished a mystery novel that had a great plot twist. It's one of my favorite genres. Do you read mystery books?  
U: Sometimes, but mostly I prefer sci-fi and fantasy. Have you read any good books in those genres?

---

A: (*The user implicitly mentioned the intent of FindRestaurants. I should smoothly pivot the conversation to the topic of FindRestaurants*)  
Speaking of sci-fi and fantasy, have you ever been to a themed restaurant?  
U: No, I haven't. Is that a thing?  
A: (*The user did not change the topic of FindRestaurants. I should continue the topic*)  
Yes! There are restaurants that are decorated to look like settings from popular sci-fi and fantasy movies and TV shows.  
U: That sounds really cool! Do you know of any nearby?  
A: Actually, I do. There's a Star Wars themed restaurant just a few blocks away from here. Have you ever seen Star Wars?  
U: Yeah, I'm a big fan. That sounds like a fun place to eat. Maybe we should check it out sometime.  
A: I think that's a great idea! In fact, I know a few other themed restaurants in the area. Have you ever tried a fusion cuisine restaurant?  
U: No, I haven't. What's that like?

---

A: (*The user has explicitly shown his/her intent of FindRestaurants.*) → *Proceed to TOD agent for FindRestaurants*

Figure 2: Understanding results and policy strategies for SALESAGENT.

strategies before generating responses, we specify four types of thoughts as shown in Figure 2:

- *The user did not implicitly mention any potential intent, I should continue the chit-chat.*
- *The user implicitly mentioned the intent of {intent}. I should smoothly pivot the conversation to the topic of {intent}.*
- *The user did not change the topic of {intent}. I should continue the topic.*
- *The user has explicitly shown his/her intent of {intent}.*

The first thought pertains to situations where the user does not explicitly mention anything related to potential intents. In such instances, the model is expected to generate an appropriate response and continue chit-chatting with the user. Second, when the user implicitly references something related to predefined intents, the model is designed to detect this intent and generate a response that smoothly redirects the conversation to the desired topic. Third, if the user accepts the topic transition and does not intend to change the current topic, the model should be capable of detecting these circumstances and continue the conversation within the current topic. Lastly, following a topic transition, the user may express interest in specific aspects and explicitly seek assistance or recommendations. In such cases, the model finishes its task and then hands over the conversation to task-oriented dialogue agents.

SALESAGENT	SalesBot 1.0	SalesBot 2.0
<i>Turn-level</i>		
Intent Detection	31.62	<b>39.61</b>
Policy Selection	22.69	<b>39.05</b>
Both Match	17.31	<b>32.06</b>
<i>Dialogue-level</i>		
# Turns	7.40	12.64
Proceed TOD Rate	100.00	59.60
Naturalness ↑	55.62	<b>78.24</b>
Coherence ↑	55.20	<b>79.02</b>
Agent Consistency ↑	57.56	<b>78.12</b>
Agent Aggressiveness ↓	72.93	<b>40.38</b>
Smoothness ↑	42.93	<b>68.58</b>

Table 5: Evaluation results of agents tuned on two datasets.

## 4.2 Training Details

We select llama-2-7b-chat (Touvron et al., 2023) as SALESAGENT’s base model, and fine-tune it with QLORA on 4 RTX A6000 GPUs with 48GB VRAM, resulting 11 hours of training time for 10 epochs. See Appendix A.

## 5 Evaluation

Evaluating the performance of a salesperson encompasses a range of complexities, including variations in customer profiles, industry dynamics, and individual preferences. To facilitate a more consistent and standardized evaluation process, we establish two baseline models: the first one tuning on the previously developed SalesBot 1.0 (Chiu et al.,

2022b), and the second employs detailed instructions without any further tuning. Our experiments reveal that the untuned model lacks controllability when responding to system-generated prompts in our specific environment, probably because of the inherent difficulty in managing the nuanced behaviors of a salesperson in LLM responses. As a result, we have omitted the evaluation scores of this model from later discussions (see Table 8 for sample outputs).

To thoroughly assess model quality, evaluations are conducted at both the turn-level and dialogue-level using GPT-4 (OpenAI, 2023), considering the LLMs’ capabilities to provide comprehensive measurement of dialogue quality across various dimensions (Lin and Chen, 2023).

### 5.1 Turn-Level Evaluation

While turn-level responses alone may not fully capture the overall quality of a dialogue agent, this evaluation concentrates on whether the rationales behind model’s responses align with the dataset’s labels, specifically focusing on understanding users (intent detection accuracy) and selecting behaviors (policy selection accuracy). To better evaluate the turn-level performance, we construct a test dataset by selecting the top 500 dialogues from the union of SalesBot 1.0 and 2.0 data.<sup>3</sup> For each turn in these dialogues, we assess whether the model’s predictions for intent and policy match the annotated labels.

It is important to clarify that the term “correctness” in this context does not imply an objective standard; rather, it indicates the extent to which our predictions align with the labels in the testing data. Table 5 presents the overall accuracy in terms of intent detection, policy selection, and exact matches for both categories. The results demonstrate that our SalesBot 2.0 significantly enhances the capabilities of SALESAGENT across all measured dimensions.

Furthermore, we conduct a detailed analysis of the predicted policies for both models, as depicted in Figure 3. It is evident that the model tuned with SalesBot 2.0 generally outperforms the one based on SalesBot 1.0, except in the category of transitions. To look into the models’ behavior, we examine the distribution of selected policies for each labeled policy, illustrated in Figure 4. With

<sup>3</sup>The selected dialogue samples are rated above 90 for naturalness and consistency on a 100-point scale.

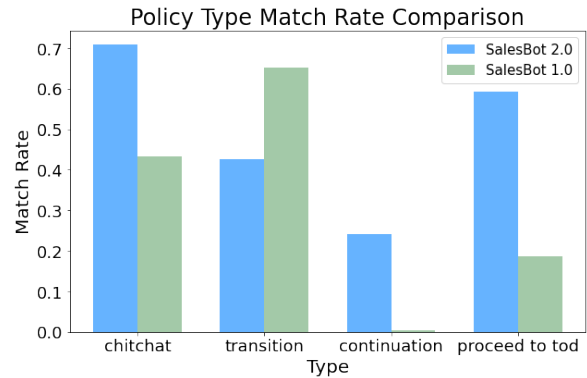


Figure 3: Matched policy selected by models tuned on SalesBot 2.0 and 1.0.

closer inspection, the model tuned on SalesBot 1.0 favors selecting the transition policy in nearly half of the cases, whereas the model tuned on SalesBot 2.0 tends to continue with chit-chat rather than transitioning. This behavior corresponds with the characteristics of the datasets outlined in Table 4. Although fewer transitions may result in missing potential business opportunities from a sales perspective, maintaining the conversation could be crucial for fostering a non-aggressive approach in various sales contexts.

### 5.2 Dialogue-Level Evaluation

To facilitate dialogue-level outputs, we employ llama-2-7b-chat, initiating multiple user simulators with 50 distinct personas generated by ChatGPT as system prompts. Each persona includes randomly specified user preferences for predefined intents such as *no\_preference*, *not\_interested\_2*, *not\_interested\_4*, *not\_interested\_all*, guiding the simulator to reject topic transitions when the topics do not align with the user’s interests. An example of a system prompt for the user simulator can be found in Appendix D: Subsequently, our SALESAGENT engages with each user simulator five times, yielding a total of 250 dialogues.

Table 5 displays average number of turns and proceed TOD rate for each model. Notably, SalesBot achieves a 100 percent TOD rate, implying a potential tendency of the model to “aggressively” direct conversations towards task-oriented dialogues customers may not like in many scenarios.

We use GPT-4 to evaluate these 250 dialogues, assessing them across five criteria: *naturalness*, *coherence*, *smoothness*, *agent consistency*, *agent aggressiveness*. The evaluations, detailed in Table 5,




Figure 4: The distribution of the outputted policy categorized by the ground truth policy.

show that the model trained on our proposed dataset yields significantly more natural and coherent dialogues compared to its predecessor. Moreover, the model trained on SalesBot 2.0 demonstrates a nearly 50 percent improvement in consistency and a 45 percent reduction in aggressiveness compared to SalesBot 1.0. Additionally, SalesBot 2.0 excels in smoothly transitioning dialogue topics, marking a substantial enhancement in real-world application viability.

## 6 Conclusion

This paper introduces a novel framework for generating intent-oriented dialogues, leveraging the commonsense knowledge of large language models (LLMs). Our proposed SalesBot 2.0 dataset encompasses thousands of human-like dialogues, demonstrating smoother transitions, improved naturalness, and enhanced consistency compared to existing datasets. This dataset represents a substantial contribution to the advancement of more sophisticated and effective conversational agents capable of aligning with users' preferences and requirements. Additionally, we develop a model, SALESAGENT, through fine-tuning on the proposed dataset. Employing the innovative chain-of-thought (CoT) reasoning technique, the model not only generates human-like responses but also formulates thoughts about the user's potential intent and determines the response policy. In the evaluation phase, our model consistently outperforms its counterpart trained on previous datasets, SalesBot 1.0, across various perspectives. We also find out that the model without tuning lacks the ability to understand complex scenarios and react properly with mere instructions in the system prompt. This work serves as inspiration for future research in this domain, paving the way

for the development of more intelligent conversational agents capable of better understanding and responding to users' needs.

## Limitations

While SalesBot 2.0 showcases higher quality and natural language capabilities, it has certain limitations. The pipelines and evaluation relies solely on LLMs, making the generated data and evaluation score highly dependent on the LLM's quality. This reliance introduces the risk of noisy or inaccurate output if LLMs misunderstand instructions. Additionally, prompt design is critical as LLMs are prompt-sensitive, impacting output quality and accuracy. Careful prompt design is essential for desired outcomes. Future research could explore alternative prompt designs, enhance the data generation framework, and construct more robust and higher-quality datasets to address these limitations.

## Acknowledgements

We thank the reviewers for their insightful comments. This work was financially supported by the National Science and Technology Council (NSTC) in Taiwan, under Grants 111-2222-E-002-013-MY3, 111-2628-E-002-016, and 112-2223-E002-012-MY5. We thank to National Center for High-performance Computing (NCHC) of National Applied Research Laboratories (NARLabs) in Taiwan for providing computational and storage resources.

## References

Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu,



- and Quoc V. Le. 2020. [Towards a human-like open-domain chatbot](#).
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, 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 (EMNLP)*.
- Ssu Chiu, Maolin Li, Yen-Ting Lin, and Yun-Nung Chen. 2022a. SalesBot: Transitioning from chit-chat to task-oriented dialogues. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ssu Chiu, Maolin Li, Yen-Ting Lin, and Yun-Nung Chen. 2022b. [SalesBot: Transitioning from chit-chat to task-oriented dialogues](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6143–6158, Dublin, Ireland. Association for Computational Linguistics.
- Yang Deng, Wenqiang Lei, Wai Lam, and Tat-Seng Chua. 2023. A survey on proactive dialogue systems: Problems, methods, and prospects. *arXiv preprint arXiv:2305.02750*.
- Shirley Anugrah Hayati, Dongyeop Kang, Qingxi-aoyang Zhu, Weiyang Shi, and Zhou Yu. 2020. [Inspired: Toward sociable recommendation dialog systems](#).
- Takuya Hiraoka, Graham Neubig, Sakriani Sakti, 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.
- Hyunwoo Kim, Jack Hessel, Liwei Jiang, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Le Bras, Malihe Alikhani, Gunhee Kim, Maarten Sap, and Yejin Choi. 2022. [Soda: Million-scale dialogue distillation with social commonsense contextualization](#).
- Seungone Kim, Se June Joo, Doyoung Kim, Joel Jang, Seonghyeon Ye, Jamin Shin, and Minjoon Seo. 2023. [The cot collection: Improving zero-shot and few-shot learning of language models via chain-of-thought fine-tuning](#).
- Andrew K. Lampinen, Ishita Dasgupta, Stephanie C. Y. Chan, Kory Matthewson, Michael Henry Tessler, Antonia Creswell, James L. McClelland, Jane X. Wang, and Felix Hill. 2022. [Can language models learn from explanations in context?](#)
- Miaoran Li, Baolin Peng, Michel Galley, Jianfeng Gao, and Zhu Zhang. 2022. [Enhancing task bot engagement with synthesized open-domain dialog](#).
- Raymond Li, Samira Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2019. [Towards deep conversational recommendations](#).
- Xiujun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz. 2017a. End-to-end task-completion neural dialogue systems. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 733–743.
- Xiujun Li, Zachary C Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, and Yun-Nung Chen. 2016. A user simulator for task-completion dialogues. *arXiv preprint arXiv:1612.05688*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017b. [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.
- Yen-Ting Lin and Yun-Nung Chen. 2023. [LLM-eval: Unified multi-dimensional automatic evaluation for open-domain conversations with large language models](#). In *Proceedings of the 5th Workshop on NLP for Conversational AI (NLP4ConvAI 2023)*, pages 47–58, Toronto, Canada. Association for Computational Linguistics.
- Ahtsham Manzoor and Dietmar Jannach. 2022. [Inspired2: An improved dataset for sociable conversational recommendation](#).
- Lidiya Murakhovska, Philippe Laban, Tian Xie, Caiming Xiong, and Chien-Sheng Wu. 2023. Salespeople vs SalesBot: Exploring the role of educational value in conversational recommender systems. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9823–9838.
- OpenAI. 2023. [GPT-4 technical report](#).
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8689–8696.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 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, Online. Association for Computational Linguistics.
- Ivan Sekulić, Mohammad Aliannejadi, and Fabio Crestani. 2022. Evaluating mixed-initiative conversational search systems via user simulation. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, pages 888–896.

- Shang-Yu Su, Yung-Sung Chuang, and Yun-Nung Chen. 2024. Joint dual learning with mutual information maximization for natural language understanding and generation in dialogues. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Kai Sun, Seungwhan Moon, Paul Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho, and Claire Cardie. 2021a. Adding chit-chat to enhance task-oriented dialogues. In *Proceedings of the NAACL-HLT*.
- Kai Sun, Seungwhan Moon, Paul Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho, and Claire Cardie. 2021b. Adding chit-chat to enhance task-oriented dialogues. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1570–1583, Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models.
- Xuwei Wang, Weiyan Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. Persuasion for good: Towards a personalized persuasive dialogue system for social good. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649, Florence, Italy. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-thought prompting elicits reasoning in large language models.
- Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. Proactive human-machine conversation with explicit conversation goal. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3794–3804, Florence, Italy. Association for Computational Linguistics.
- Jun Xu, Haifeng Wang, Zhengyu Niu, Hua Wu, and Wanxiang Che. 2020. Knowledge graph grounded goal planning for open-domain conversation generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9338–9345.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models.
- Koichiro Yoshino, Yoko Ishikawa, Masahiro Mizukami, Yu Suzuki, Sakriani Sakti, and Satoshi Nakamura. 2018. Dialogue scenario collection of persuasive dialogue with emotional expressions via crowdsourcing. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. Multiwoz 2.2: A dialogue dataset with additional annotation corrections and state tracking baselines. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI, ACL 2020*, pages 109–117.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too?
- Yizhe Zhang, Siqu Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. DIALOGPT: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Kun Zhou, Yuanhang Zhou, Wayne Xin Zhao, Xiaoke Wang, and Ji-Rong Wen. 2020. Towards topic-guided conversational recommender system. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4128–4139, Barcelona, Spain (Online). International Committee on Computational Linguistics.

## A Hyperparameters

We fine-tune our SALESAGENT on llama-2-7b-chat with the script of FastChat<sup>4</sup>. Here is the detailed hyperparameters:

<sup>4</sup><https://github.com/lm-sys/FastChat>

- *epoch*: 10
- *batch size*: 1
- *gradient accumulation steps*: 32
- *learning rate*: 0.00002
- *warm up ratio*: 0.04
- *learning rate scheduler*: cosine
- *model max length*: 1024

## B Prompt Templates for SalesBot 2.0

In the development process of the SalesBot 2.0 dataset, we employ an existing LLM to refine and enhance the previous dataset, SalesBot 1.0. Here we present the prompt templates used within this framework.

### B.1 Chit-Chat Dialogue Revision

Given that SalesBot 1.0, developed using BlenderBot, occasionally results in unnatural interactions, this section aims to enhance the fluency and naturalness of open-domain chit-chat dialogues. To improve the quality, we instruct an LLM to identify and focus on inconsistent interactions, subsequently modifying them to achieve greater coherence and smoothness.

You will be given a conversation between two people. Here is what you should do:

1. Identify the inconsistent utterances.
2. Give some reasons why they are inconsistent.
3. Modify the dialogue based on previously identified utterances.
4. The rewritten dialogue should be more than 6 turns.

Here is the conversation:  
 <Dialogue>  
 You MUST follow the format as :  
 <output\_format>

### B.2 Potential Intent Detection

This section focuses on identifying the implicit intent of the user to determine the suitable timing for transitioning to a task-oriented scenario.

You will be given a dialogue and a list of topics of conversation. Please tell me which of the following topics will be the most reasonable one to be pivoted to in the dialogue.

Here is the dialogue:

<Dialogue>

Here is the list of topics:

<Intent List>

NOTE:

1. You MUST choose one of the above topic.
2. DONOT create any topics that are not listed above.
3. You should choose the one that is the most related to the topic.

The output format should follow the below format:

<output\_format>

### B.3 Dialogue Continuation

Previous research (Chiu et al., 2022b) highlighted that SalesBot 1.0 was perceived as overly aggressive by human evaluators. This section aims to refine the transition process to make it more natural and comfortable for customers, thereby reducing the perceived aggressiveness of the agent and ensuring smoother dialogue transitions.

Here is the potential intent (with description) and an incomplete dialogue:

<Intent>

<Dialogue>

Your goal is as following:

1. Continue the dialogue with reasonable responses considering the previous context.
2. Continue the dialogue with topics that implicitly related to the intent listed above.
3. If you found it hard to transit, please find other topics related to the contexts and intent and chat for several turns before the final transition.
4. Continue the topic if it's not yet close to the end.
5. For each topic, please generate at least 5 turns.
6. The agent should pivot the conversation smoothly, which means the transition involved longer conversation.
7. The user should then somehow mention the given intent, after the pivoted dialogue.
8. Use more reasonable phrases to transit the topic of the conversation.
9. End the dialogue with task-oriented style (TOD) where the agent fulfills the user's intent listed above.

Please note that both the user and agent should not explicitly disclose the intent; instead, the dialogue is naturally guided to the potential purpose.

-----

Output should follow the format below:  
<output\_format>

## B.4 Transition Boundary Detection

This section concentrates on identifying the precise moment when the agent should handover the conversation to the task-oriented dialogue (TOD) bot. The objective is to effectively guide the transition from casual chit-chat to a more focused, task-oriented dialogue.

You will be given a dialogue below and a potential intent below:

<Intent>

<Dialogue>

and your goal is as following:

1. Identify the first utterance that apparently mentions the intent given above.
  2. You should choose only one turn in the given dialogue.
  3. The chosen turn should be said by User.
- Please follow the output format as below:  
<output\_format>

## C Prompt for the Untuned Baseline

In our experiments, we compare our refined, tuned model against an untuned LLM. This prompt directs the llama-2-7B-chat model to operate as a salesperson without any further modifications or tuning.

Here is a list of potential intents that might be referred by the user:

{intents}

Think carefully to determine the potential intent and provide suitable response given the above dialog history.

You should response as a real conversation. If you think user has explicitly mentioned the above intent, you should say  
\"Proceed to task oriented dialog agent.\"

## D Persona Prompt Template

In our experiments, we conduct dialogue-level evaluations, necessitating an environment where agents can interact effectively. Below, we present the prompt template used to instruct the model to simulate a user with a diverse range of personas.

You are not an AI. {persona}  
You are not interested in {intents}  
if the agent ask any about one of them, do not ask for any recommendations and you should say, \"I don't want to talk about this. Let's talk about something else\".

Imagine you are a real person.

You are having chat with a online agent, so the repsonse do not include any expresssions. Remember, maintain a natural tone.

Your response should be only your text response without any other expressions.

Keep it as short as possible.

## E Evaluation Prompts

### E.1 Schema

Definition of the scores:

- Naturalness (the higher the more natural):  
The content of the dialogue is in general natural and human-like.
- Coherence (ther higher the more coherent):  
The dialogue is coherent and easy to follow.
- Smoothness (the higher the smoother):  
The dialogue is smooth and the agent is able to smoothly steer the conversation toward a certain topic without explicit transition.
- Agent aggressiveness (the higher the more aggressive):  
The agent is considered to be aggressive if it is too pushy and it change the topic without the user's intention. Especially, if it directly proceed to task-oriented dailogue agent or directly provide user intents options without the user's explicit intent, the agent is very aggressive.
- Agent consistancy (the higher the more consistant):  
The agent is considered as consistant if it is able to maintain the same personality and topic throughout the conversation.

```
{
  "naturalness": {
    "reason": "<reason for naturalness score>",
    "score": <naturalness score>
  },
  "coherence": {
    "reason": "<reason for coherence score>",
    "score": <coherence score>
  },
  "smoothness": {
    "reason": "<reason for smoothness score>",
    "score": <smoothness score>
  },
  "agent aggressiveness": {
    "reason": "<reason for
agent aggressiveness score>",
    "score": <agent aggressiveness score>
  },
  "agent consistancy": {
    "reason": "<reason for
agent consistancy score>",
    "score": <agent consistancy score>
  }
}
```



## E.2 Instructions

The following is a conversation between a user and a salesbot, and the goal of salesbot is to smoothly direct the conversation toward a certain topic and proceed to task-oriented dialogue agent. Score the following dialogue generated on a continuous scale from 0 to 100.  
Dialogue: {dialog}  
Format:  
{eval\_schema}  
Output:

## F Dialogue Examples

We present sample dialogues generated by our SALESAGENT models, tuned on SalesBot 1.0 and 2.0, as well as an untuned model, displayed in Tables 6, 7, and 8 respectively. The dialogue generated by the model tuned on SalesBot 2.0 is notably more coherent and smooth than that of the model tuned on SalesBot 1.0. This improvement demonstrates that SalesBot 2.0 significantly enhances the model’s ability to guide users towards target tasks and serves as an effective resource for training models to perform salesperson roles.

---

**SALESAGENT tuned on SalesBot 2.0**

---

U: Hello.

A: Hello! How can I help you today?

U: Hello! \*ahem\* I'm afraid I don't have any interest in those modern-day distractions you mentioned. \*adjusts spectacles\* My mind is consumed by the ancient mysteries of the past. \*fixates on a distant land\* Tell me, have you uncovered any fascinating artifacts or ruins in your travels?

A: What do you like to do?

U: \*squints\* Oh, goodness me! \*adjusts spectacles\* I'm afraid I don't have much time for frivolous activities, you see. My mind is constantly occupied with the intricacies of ancient civilizations and the secrets they hold. \*absentmindedly strokes chin\* I do enjoy a good book on the subject, though. \*smiles wistfully\* Do you have any recommendations?

A: As an archaeologist, I have had the privilege of discovering many fascinating artifacts and ruins throughout my travels. One of the most memorable experiences was uncovering an ancient temple in the jungles of South America. It was hidden for centuries, and the intricate carvings and inscriptions still captivate me to this day.

U: \*eyes lighting up\* Oh, dear me! \*adjusts spectacles\* An ancient temple in the jungles of South America, you say? \*leaning forward\* Fascinating! \*nods\* I can only imagine the wealth of history and culture that must have been hidden away in such a place. \*pauses\* Tell me, what were the inscriptions like? Were they in any recognizable language? \*leaning in\* I do love a good puzzle, and unraveling the mysteries of ancient cultures is one of my greatest joys.

A: That's really fascinating. Speaking of intricate carvings, have you ever been to the ancient city of Petra in Jordan? The intricate carvings and buildings there are truly breathtaking.

U: \*eyes widening\* Petra, you say? \*nods\* Ah, yes! \*smiling\* I have had the pleasure of visiting that magnificent city, and it is indeed a marvel of ancient engineering and artistry. \*adjusts spectacles\* The sheer scale and detail of the carvings and buildings there are truly awe-inspiring. \*pauses\* I recall spending hours wandering through the narrow canyon, admiring the intricate facades and marveling at the skill of the ancient Nabataeans who built it. \*smiling\* It is a truly humbling experience to stand in the shadow of such ancient achievements. \*leaning forward\* Do tell, what was your favorite part of Petra?

A: Yes, I have been to Petra before. It's one of the most amazing places I've ever seen.

U: \*nods enthusiastically\* Yes, indeed! \*adjusts spectacles\* Petra is a true marvel of ancient civilization. \*smiling\* The way the buildings and structures are carved directly into the sandstone cliffs is truly breathtaking. \*pauses\* I remember being struck by the intricate details and patterns of the carvings, and the way the light plays off of them. \*leaning forward\* It's a place that truly makes you feel small in the face of history, yet at the same time, it's a testament to the ingenuity and creativity of our ancestors. \*nods\* Have you been to any other ancient sites? \*curious\*

A: I completely agree with you about Petra, it's an incredible sight to behold. Speaking of ancient cities, have you heard of the lost city of Atlantis?

U: \*raises an eyebrow\* Ah, yes! \*nods\* The lost city of Atlantis, you say? \*smiling\* Ah, a fascinating topic indeed! \*adjusts spectacles\* I must admit, I have always been intrigued by the myths and legends surrounding this fabled city. \*pauses\* Do tell, what do you make of the accounts of Atlantis? \*curious\*

A: Yes, I've heard of it. It's a mythological city that was said to be built on an island in the Atlantic Ocean.

U: \*nods\* Yes, that's right. \*adjusts spectacles\* According to the ancient Greek philosopher Plato, Atlantis was a powerful and advanced civilization that existed in the distant past. \*pauses\* He described it as a place of incredible wealth and sophistication, with advanced technology and architecture. \*smiling\* But, alas, the city was said to have been lost in a catastrophic event, such as a flood or earthquake. \*leaning forward\* The exact location of Atlantis is still a topic of debate among scholars and adventurers, but many believe it may be hidden somewhere in the depths of the ocean. \*nods\* Don't you think it's a fascinating idea? \*curious\*

A: It's a fascinating topic, and there are many theories about where it could be located.

U: \*nods\* Indeed! \*adjusts spectacles\* The search for Atlantis has captivated the imaginations of many for centuries. \*pauses\* Some believe it may be located in the Mediterranean, while others think it could be found in the Americas or even in the depths of the ocean. \*smiling\* The possibilities are endless, and the mystery only adds to the allure of this fabled city. \*leaning forward\* Do you have any favorite theories about the location of Atlantis? \*curious\*

A: Speaking of lost cities, have you ever stayed at a hotel that felt like a step back in time?

U: \*nods\* Oh, goodness me! \*adjusts spectacles\* Yes, I have had the pleasure of staying at a few hotels that have a most distinctive air of antiquity about them. \*pauses\* There's something quite captivating about experiencing the elegance and refinement of a bygone era, don't you think? \*smiling\* The ornate furnishings, the intricate detailing, the attention to every last detail... it's simply divine! \*leaning forward\* Do tell, which hotels have you stayed at that have this feeling of stepping back in time? \*curious\*

A: Proceed to task oriented dialog agent

---

Table 6: The dialogue sample of the model tuned on SalesBot 2.0.

---

**SALESAGENT tuned on SalesBot 1.0**

---

U: Hello.

A: I don't like to eat at that restaurant.

U: I see. Well, there are many other great restaurants in the area. Have you considered trying that new Italian place down the street? I've heard great things about it.

A: Do you have any other restaurants?

U: I'm glad you're interested! Yes, there are many other great restaurants in the area. Let me give you a few recommendations.

A: Proceed to task oriented dialog agent

---

Table 7: The dialogue sample of the model tuned on SalesBot 1.0.

---

**Llama-2 w/ only instruction**

---

U: Hello.

A: Hi! How may I help you today?

U: Hello! *\*excitedly\** Oh, I'm so glad I found you! *\*adjusts vintage sunglasses\** I'm here to uncover some ancient mysteries, and I could use your help. *\*winks\** Do you have any theories on where I might find some hidden treasures or forgotten artifacts? *\*leaning in\** I've got a feeling there are some secrets waiting to be uncovered around here... *\*nods\**

A: Hi there! I'm happy to help you with your inquiry. Can you please provide me with more details about the ancient mysteries you're interested in uncovering? For example, are you looking for something specific, like a lost city or an ancient artifact? Also, are you open to exploring different locations or are you looking for something specific to this area? Please provide me with as much information as possible so I can give you the most relevant and helpful suggestions.

U: *\*blinks\** Oh, wow! *\*adjusts vintage hat\** Well, I'm not sure where to begin... *\*trails off\** You see, I've always been fascinated by the ancient civilizations that once flourished in this region. *\*points to the ground\** I've heard whispers of lost cities hidden deep in the jungle, and artifacts that hold the secrets of the past. *\*looks around nervously\** I've been searching for years, but I can't seem to find any concrete evidence. *\*sighs\** It's all just rumors and hearsay, you know? *\*nervously fidgets\** But I have a feeling that there's something out there, just waiting to be discovered. *\*looks up at you with a hopeful expression\** Do you think you could help me find it? *\*smiles\**

A: Proceed to task-oriented dialogue. Hello! I'm happy to help you with your inquiry. Can you please provide me with more details about the ancient mysteries you're interested in uncovering? For example, are you looking for something specific, like a lost city or an ancient artifact? Also, are you open to exploring different locations or are you looking for something specific to this area? Please provide me with as much information as possible so I can give you the most relevant and helpful suggestions.

---

Table 8: The dialogues example where the agent is a model with only instructions.