

# Wizard of Shopping: Target-Oriented E-commerce Dialogue Generation with Decision Tree Branching

Xiangci Li<sup>1\*</sup> Zhiyu Chen<sup>2</sup> Jason Ingyu Choi<sup>2</sup>

Nikhita Vedula<sup>2</sup> Besnik Fetahu<sup>2</sup> Oleg Rokhlenko<sup>2</sup> Shervin Malmasi<sup>2</sup>

<sup>1</sup> AWS AI Labs <sup>2</sup> Amazon.com, Inc.

lixiangci8@gmail.com

{zhiyuche, chojson, veduln, besnikf, olegro, malmasi}@amazon.com

## Abstract

The goal of conversational product search (CPS) is to develop an intelligent, chat-based shopping assistant that can directly interact with customers to understand shopping intents, ask clarification questions, and find relevant products. However, training such assistants is hindered mainly due to the lack of reliable and large-scale datasets. Prior human-annotated CPS datasets are extremely small in size and lack integration with real-world product search systems. We propose a novel approach, TRACER, which leverages large language models (LLMs) to generate realistic and natural conversations for different shopping domains. TRACER’s novelty lies in grounding the generation to dialogue plans, which are product search trajectories predicted from a decision tree model, that guarantees relevant product discovery in the shortest number of search conditions. We also release the first target-oriented CPS dataset *Wizard of Shopping* (*WoS*), containing highly natural and coherent conversations (3.6k) from three shopping domains. Finally, we demonstrate the quality and effectiveness of *WoS* via human evaluations and downstream tasks.

## 1 Introduction

E-commerce platforms use product search to help customers find items in billion-scale product catalogs. Current product search systems rely on a *search-and-refine* process, where customers iteratively refine keywords expressing their shopping needs and then assess the top- $k$  results for relevance. This process is customer-driven, and without sufficient expertise or patience, may lead to suboptimal purchase decisions or shopping journey abandonment. Conversational Product Search (CPS) (Zhang et al., 2018; Bi et al., 2019; Zou et al., 2022) address the limitations of traditional

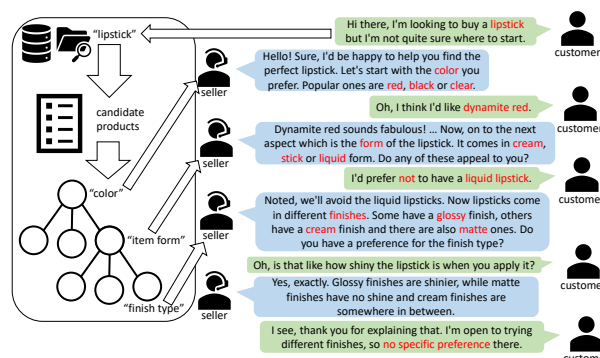


Figure 1: An example of generated shopping conversation using TRACER, which leverages the LLM with different roles and decision tree-based planning.

product search by adding a conversational layer into product search, which provides a more natural and delightful shopping experience than traditional systems.

Nonetheless, existing product search systems are not optimized for either natural language queries or real dialogues. One solution involves training a model to rewrite conversations into decontextualized queries (Yu et al., 2020; Wu et al., 2022). However, prior rewriting datasets do not cover the shopping domain. Another solution is to directly train CPS models with dialogues (Zhang et al., 2018; Bi et al., 2019; Zou et al., 2022), but this approach often relies on synthetic datasets with templated content, which limits the realism of the trained models. Consequently, the effectiveness of these models on real dialogues remains uncertain. Recently, Bernard and Balog (2023) released a small-scale dataset with expensive human annotation. Overall, the scarcity of datasets presents a significant challenge in developing intelligent CPS assistants.

In order to tackle the above challenge, we propose a novel approach, TRACER (Target-oriented e-commerce dialogue generation with decision tree branching) to simulate shopping conversa-

\* Work performed by the author as a PhD candidate at The University of Texas at Dallas before joining AWS AI Labs.

tions between two LLM agents. Specifically, we assign two distinct roles to each LLM: (1) a customer with specified shopping interests (a.k.a. the "Apprentice"); (2) a seller agent (a.k.a. the "Wizard") that has access to a product catalog. To overcome the known shortcomings of LLMs, such as hallucination, prompt brittleness, lack of controllability, and planning (Kaddour et al., 2023), two agents leverage dialogue plans predicted from a decision tree to ground dialogue generation. Since the decision tree is trained to select product attributes that maximally split the search space, customers are guaranteed to reach target products in the shortest number of search conditions, minimizing user effort (Al-Maskari and Sanderson, 2010). Human evaluations further validate that the generated dialogues are highly natural and coherent since LLMs can easily convert dialogue plans to natural conversations (see Figure 1 as an example). Our contributions are as follows<sup>1</sup>:

- A highly controllable, scalable, and easy-to-apply dialog generation approach TRACER for different shopping domains.
- Release of *Wizard-of-Shopping* (WoS) dataset which includes 3.6k shopping conversations across 3 different product domains.
- Demonstration of the utility of WoS by showing improvements in conversational query generation and product ranking tasks.

## 2 Related Work

**Conversational Product Search.** Zhang et al. (2018) proposed a seminal CPS system that unified conversational product search and recommendation, and actively asked questions to understand user needs. Their framework predicts both the true item and the next question at each conversation turn. Bi et al. (2019) extended this by predicting the probability of an aspect-value pair being positive or negative to enhance conversational product ranking performance. Zou et al. (2022) learned a joint representation of users, queries, items, and conversations to retrieve target items within a latent semantic space. Xiao et al. (2021) created an online shopping conversational dataset by transferring the utterance structures from movie domain conversations. They used proprietary search behavior data to supervise the construction of intent flow.

<sup>1</sup><https://github.com/jacklxc/Wizard-of-Shopping>

A fundamental issue in prior studies (Zhang et al., 2018; Bi et al., 2019; Zou et al., 2022) is that all the proposed CPS models are evaluated using slotted templates specifically designed for clarification questions. The weaknesses of such an evaluation are twofold. First, as the quality of simulated conversations is not assessed and may be of poor quality, the trained conversational product rankers might not generalize well to real conversations. Second, customers may lack sufficient product knowledge, making it unrealistic to expect them to always provide answers to the seller’s questions.

The paucity of datasets significantly hinders CPS research. Many datasets used in prior works (Xiao et al., 2021; Chen et al., 2020; Zhao et al., 2022) are privately owned, which prevents reproduction and further reusing of their work for future studies. On the other hand, Bernard and Balog (2023) released MG-ShopDial, a human-generated e-commerce dialogue dataset consisting of 64 conversations. We perform an in-depth comparison in §4.6. Moreover, *due to commercial interests and privacy concerns, no large-scale human-generated CPS dataset is publicly available, and such a dataset may remain unavailable in the foreseeable future.*

*Our work aims to address the data scarcity issue by guiding LLMs to generate natural CPS conversations.* We prompt LLMs to generate conversations grounded on sampled customer preferences and product information from a real product catalog. The customer is also allowed to ask sellers general knowledge questions, which is lacking in prior work (Zhang et al., 2018; Bi et al., 2019; Zou et al., 2022). Overall, our generated dialogues in the WoS dataset are more natural and diverse than prior simulated conversations. We believe that our proposed method and datasets will be a valuable contribution to CPS and its sub-tasks.

### **Difference with Prior Work in Conversational Recommendation & Search.**

Similar to previous CPS work, prior studies in conversational recommendation (CR) systems (Li et al., 2018; Hayati et al., 2020; Liu et al., 2020, 2021; Deng et al., 2021; Zhang et al., 2022b) explored the recommendation strategies by learning both the user and item representations. Indeed, the problem of learning a conversation strategy for the seller to find proper product aspects that shorten the product-seeking process seems to be shared between CPS and CR. A key difference between our work and prior studies is that *we do not build a new CPS or CR sys-*

tem. Instead, we focus on the creation of a dataset featuring natural conversations, accompanied by dialogue plans based on decision trees and LLMs. This dataset is designed to support downstream tasks such as query generation and product ranking. Notably, our product retrieval methods build upon traditional approaches like BM25, incorporating conversational layers to achieve CPS. This sets our work apart from previous CPS efforts (Zhang et al., 2018; Bi et al., 2019; Zou et al., 2022). Additionally, many CR works are in the movie domain (Li et al., 2018; Hayati et al., 2020; Liu et al., 2020, 2021) and there is no existing CPS dataset specifically tailored for the e-commerce domain. Our work addresses this gap by introducing a dataset, along with its corresponding dataset generation method, designed explicitly for the e-commerce domain.

### 3 Dialogue Generation Approach

As shown in Figure 1, the goal is to generate a shopping dialogue given a selected set of preferences that is derived from a product ( $p$ ), randomly sampled from a product catalog. TRACER generates dialogues in three steps:

1. Customer preference sampling (§3.1) assigns relevant product attributes from  $p$  into three distinct preference groups.
2. Dialogue planning (§3.2) fits decision trees to predict the sequence of product aspects to be addressed so that customers can find desired products with minimal effort.
3. Verbalization (§3.3) utilizes customer preference, planned search trajectories, and prompt engineering to generate natural and relevant dialogues.

---

#### Algorithm 1 Decision Tree-Based Search Strategy

---

```

function DECISIONTREESearch( $RevPref$ )
   $P_o \leftarrow search(RevPref)$ 
   $X, Y \leftarrow MakeDataset(P_o)$ 
   $Tree \leftarrow DecisionTree.fit(X, Y)$ 
  return  $P_o, Tree$ 
end function

```

---

#### 3.1 Customer Preference Sampling

Each product is associated with a varying number of features or aspect-value pairs, e.g.  $\{(A_1, V_1), (A_2, V_2), \dots, (A_m, V_m)\}$ . For example, a tablet case can have associated features of  $\{(model, iPad), (color, blue), (material, TPU)\}$ . We randomly assign one of the following customer

---

#### Algorithm 2 Interactive and single-pass dialogue generation strategies from §3.3.

---

```

Require:  $LLM, preference$ 
 $\triangleright$  Stage 1: Customer starts with a product category.
 $Conv\_hist \leftarrow []$   $\triangleright$  Interactive generation
 $Utterances \leftarrow LLM.verbalize\_stage1(PC)$ 
 $Conv\_hist.extend(Utterances)$ 
 $\triangleright$  Stage 2: Product search.
 $Plan, Plan_{hist} \leftarrow [], []$ 
while  $|P_o|$  has not converged do
   $RevPref \leftarrow [PC, Plan_{hist}]$ 
   $P_o, Tree \leftarrow DecisionTreeSearch(RevPref)$ 
   $Plan \leftarrow Tree.traverse(preference)$ 
 $\triangleright$   $Plan = \{(A_1, V_1, I_1), \dots, (A_d, V_d, I_d)\}$ 
   $Plan_{hist}.extend(Plan)$ 
   $Utterances \leftarrow LLM.verbalize\_stage2(Plan, Conv\_hist)$ 
   $Conv\_hist.extend(Utterances)$ 
end while
 $\triangleright$  Stage 3: Product recommendation.
 $p \leftarrow RandomSample(P_o)$ 
 $Utterances \leftarrow LLM.verbalize\_stage3(p, Conv\_hist)$ 
 $Conv\_hist.extend(Utterances)$ 
 $\triangleright$  Single pass generation
 $Single\_pass\_conv \leftarrow LLM.verbalize(preference, Plan_{hist}, p)$ 
Output  $Conv\_hist, Single\_pass\_conv$ 

```

---

interest values ( $I$ ) to each aspect of the sampled product  $p$ :

- **Wanted:** customer is interested
- **Unwanted:** customer is NOT interested
- **Optional:** customer does not care much

The *wanted* aspect-value pairs come directly from  $p$ . As for an aspect-value pair  $(A_x, V_x)$  from  $p$  assigned with *unwanted*, we randomly replace  $V_x$  with a value  $V'_x \neq V_x$  as the customer's unpreferred value, where  $V'_x$  appears as a value for  $A_x$  in the product catalog, so that the *unwanted* pair becomes  $(A_x, V'_x)$ . For example, if  $A_x$ ="color" and  $V_x$ ="blue",  $V'_x$  can be "red". We use the Product Category ( $PC$ ) of product  $p$ , as well as the sampled *wanted*, *optional*, and *unwanted* features to simulate the *preference* of the customer agent:

$$preference = [PC, (A_1, V_1, I_1), \dots, (A_m, V_m, I_m)] \quad (1)$$

where  $I_i \in \{wanted, unwanted, optional\}$ . When  $I_i = optional$ , we set  $V_i$  as empty since the customer does not care about the value of the optional aspect.

#### 3.2 Dialogue Planning with Decision Tree

Given a sampled *preference*, dialogue planning as the next step determines the sequence of product aspects to be addressed in the conversation. Since

---

**Instruction**

You are a scriptwriter. For the aspects or aspect value pairs below, write a chat conversation between the customer and seller about the seller trying to narrow down the customer's need. **The seller knows well about the sold products, while the customer has limited knowledge about the products.** [...] The customer should speak out first and say he/she wants to buy  $\{\{ProductCategory\}\}$ . [...] You MUST cover ALL aspect-value pairs below, including the wanted, unwanted, and optional aspect-value pairs. **Make sure you reorganize and reorder the aspect value pairs and translate them into a natural conversation in a meaningful way.** [...]

---

**Wanted / Optional / Unwanted Features**

Aspect or aspect value pairs that the customer must say {he wants to have / is optional / he does not want to have} for the target product:

Aspect:  $\{\{Aspect\}\}$ , Value:  $\{\{Value\}\}$ ;

...

---

**Candidate Value Examples**

**Additionally, there are some common values for each aspect to be mentioned. The seller should use these typical values in his question:**

Aspect:  $\{\{Aspect\}\}$ , Typical Values:  $\{\{Value_1\}\}$ ,  $\{\{Value_2\}\}$ ,  $\{\{Value_3\}\}$ ;

...

---

**Final Instruction**

Output a conversation following all instructions above and make sure you double check if ALL aspect value pairs above are mentioned:

---

Table 1: Simplified prompt format for single pass dialogue generation approach. Prompts in bold are designed for improving conversation naturalness (§3.4).

previous studies on web search evaluation discovered that user effort is inversely correlated with user satisfaction (Al-Maskari and Sanderson, 2010), we claim that real customers will also want to minimize their efforts in CPS. To reflect this intuition in dialogue generation, we focus on controlling the “customer” LLM to be aligned with expected customer behavior by intelligently deciding what to ask at each turn.

An intuitive approach is to utilize the concept of purity in decision trees (Mehta et al., 1996; Quinlan, 2014). At each turn, given *partially* revealed preference  $RevPref^2$ , the decision tree selects the next product attribute that maximally divides the current search space. This is important because in real applications, customers are also expected to reveal their preferences step-by-step.

As Algorithm 1 shows, first, at each turn that requires search, a set of products ( $P_o$ ) that satisfy  $RevPref$  are retrieved from our search system. Note that our search system is a simple rule-based system that can filter out products based on  $RevPref$ . Then, product attribute features

<sup>2</sup> $RevPref$  is initialized with  $PC$  only and grows as more product features are revealed by the decision trees.

from  $P_o$  are used for constructing a temporary training dataset to fit a decision tree, and we use all aspects from product catalog as features for each product. If a product does not have certain aspects, corresponding values are treated as empty. For example, let’s consider a product from the *tablet case* category with three features. Input features ( $X$ ) can be represented as  $\{(model, iPad), (color, blue), (material, TPU)\}$  and the corresponding labels ( $Y$ ) are represented as a string “*model:iPad&color:blue&material:TPU*”. In other words, products associated with the same aspect-value pairs all belong to the same leaf node of the decision tree. As a result, the fitted *Tree* contains a structured collection of nodes whose attributes are the product aspect keys and split conditions are corresponding aspect values.

Lastly, given the fitted *Tree* and *preference*, the best dialogue plan (a sequence of aspect-value-interest tuples) is generated. Specifically, we traverse starting from the root and choose the next child based on the customer *preference*’s corresponding values and interest of the product aspect of the parent node. Any traversed aspect that does not appear in *preference* is *optional*, indicating the customer does not care about the product aspect. Note that decision tree is fitted individually for each aspect selection step and used together with additional strategies (§3.4) to increase the naturalness of generated dialogues. This search and planning process (Algorithm 1) iterates until every product in  $P_0$  satisfies *preference*, or until there are no more product aspects to discuss (Algorithm 2 Stage 2).

### 3.3 Verbalization

We feed the given *preference* and dialogue plan contexts through zero-shot prompting to generate dialogues, proposing two methods of verbalization.

**Interactive Generation.** We use different prompts for the seller and customer (Appendix Tables 15 and 16) agents so that they alternatively speak to each other utterance-by-utterance, like in the real world. The agents respond based on the previous conversation history. To ensure that the LLM-powered agents follow the instructions to adhere to the pre-determined features, we use a *dialogue state tracker* (Young et al., 2010) to track the customer’s *wanted*, *optional*, and *unwanted* features, that have been mentioned and are to be mentioned after each utterance. This is



implemented by a hybrid strategy of rule-based keyword matching and a GPT-4 functional call prompt similar to prompts in Table 16, except that the expected output is the updated inputs to the “update\_dialogue\_state” function in Table 16 instead of response utterances. This complex mechanism implies that the interactive approach is more prone to errors, as we show in §4.2.

**Single-Pass Generation.** Alternatively, we propose a simpler approach by feeding all product features pre-determined by the decision tree plan into a single LLM, and generating the entire conversation with a single pass (Table 1, see full prompts in Table 17). Despite we lose the interactive capacity, the generated conversations have fewer errors, as we show in §4.2.

### 3.4 Enhancing Conversational Naturalness

Our overall approach is described in Algorithm 2.

**How conversation starts.** When dialogue starts, similar to prior work (Zhang et al., 2018; Bi et al., 2019; Zou et al., 2022), each customer always starts with the target product category as an initial request. For example, a customer asks for a lipstick recommendation in Figure 1.

**How conversation ends.** When all remaining products from  $P_o$  satisfy *preference*, the seller agent randomly recommends one of the products to the customer. The conversation continues for up to three more turns to conclude the shopping journey. Note that we do not specifically define finding the exact target product as a stopping criteria. Instead, we consider search results converging to a “product family” that satisfies *preference* is close enough to satisfy customer’s shopping interest.

**Improving naturalness by adding hints.** Customers may not always be aware of available options for a given product aspect, which was not well-addressed in prior studies (Zhang et al., 2018; Bi et al., 2019; Zou et al., 2022). It is more user-friendly to provide customers with options when asking clarification questions. For example, if the aspect under consideration is the size of a product’s solid-state drive (SSD), a general customer may not know the typical storage size of SSDs. To mitigate this, we propose to always require the seller to provide hints based on up to 3 most frequent values associated with the product aspect being discussed (e.g., 256GB, 512GB, or 1TB for SSD size).

Similar to the proposal of product aspects by the decision tree, we also dynamically retrieve the most frequent values of each product aspect from the remaining candidate set of products. Providing typical aspect values as examples can help clarify the meaning of the product aspect, especially when it may not be known to customers, without requiring two additional turns of clarification questions.

### Improving naturalness by LLM knowledge.

The aspects may not always be intuitive to the customer, since customers have different levels of knowledge and may not understand certain technical or specific product aspects. LLMs are known to have world knowledge internally, which can be used to “smooth” out conversations given non-intuitive product features. We encourage LLM-powered agents to clarify the meaning of aspects that require additional knowledge to be understood. We also allow the seller agent to reorder the decision tree-selected product features to be discussed, if multiple features are provided at the same time (See bolded prompts in Table 1, Appendix Tables 15, 16 & 17). Note that the decision tree is used to globally select the best-k product aspects among all possible aspect values to be elicited from the customer. The LLM interprets and locally reorders selected product aspects to enhance the natural flow of the conversation.

## 4 Experiments and Results

### 4.1 Experimental Setup

We use a subset of the TREC product search dataset<sup>3</sup>, which is a cleaned derivative of the Amazon product dataset (McAuley et al., 2015) as our product catalog. As we explained previously, we do not have any available real-user behavior data. We experiment with the Home & Kitchen, Electronics, and Beauty & Personal domains, which contain 135k, 54k, and 47k products respectively. We remove non-intuitive product aspects, such as “ASIN”, “Date First Available” and “Is Discontinued By Manufacturer” and normalize certain aspects, such as “Brand Name” to “Brand” and “Colour” to “Color”. We also normalize the price range and customer review scores into ranges in the format of “between \$10 and \$20” and “higher than 4.5 stars” respectively. Finally, we discard products with fewer than two features, excluding customer reviews and price ranges.

<sup>3</sup>[trec-product-search.github.io](https://trec-product-search.github.io)

## 4.2 Human Evaluation

Since there is no gold conversation to evaluate against as we explained §2, we perform human evaluation at the conversation and utterance levels to examine the quality of the dialogues. The metrics are summarized in Table 2. For conversation-level evaluations, judges are asked to score each conversation on a 5-point Likert scale. The definition of each score is shown in Appendix Table 8. Note that the *overall naturalness* of 5 means the generated conversation is *indistinguishable from real-life human conversations*. We also evaluate the conversations on a finer-grained utterance level where judges identify all unsatisfactory, non-script following utterances, as a binary labeling task.

We evaluate conversations generated under 4 settings: GPT-4 single-pass generation, GPT-4 interactive generation, LLaMA-2-70B-Chat single-pass, and LLaMA-2-70B-Chat interactive generation. We begin by assessing inter-annotator agreement using Kendall’s  $\tau$ . To achieve this, we provide detailed annotation instructions to expert judges divided into 5 groups, with 3 judges per group. Each judge is assigned 12 examples, spread across 3 domains with 4 configurations per example. Therefore, we have evaluation results of 180 instances for 60 conversations, allowing us to calculate inter-annotator agreement. Furthermore, we assign two of the judges to work on 24 additional examples each. These examples are divided into 3 domains, with 2 examples per domain, and each example involves 4 configurations. Thus, we evaluate 108 dialogues in total.

## 4.3 Results

Table 3 shows the average scores of conversational-level human evaluation results, and Table 4 shows the average number of unsatisfactory utterances marked by the judges. As Table 3 and 4 show, we obtain a clear trend of scores across almost all metrics with moderate to high agreements (Average Kendall’s  $\tau = 0.51$  &  $0.44$  respectively) for both conversation-level and utterance-level evaluations. Unsurprisingly, GPT-4 outperforms LLaMA-2-70B-Chat, which is consistent with the existing leaderboard<sup>4</sup>.

The single pass generation approach consistently outperforms the interactive generation for both LLM types. This is because the single pass ap-

<sup>4</sup><https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>

Metric	Description
<i>Conversation-Level</i>	
Realism	Whether the conversation is likely to happen in the real world from a logical perspective.
Conciseness	Whether the conversation is too verbose.
Coherence	Whether the conversation is fluent and coherent with respect to the conversation history from the beginning.
Overall naturalness	Overall subjective impression of the conversation and whether it sounds natural.
<i>Utterance-level</i>	
Realism	Whether the utterance is likely to be spoken in the real world from a logical perspective.
Script-following	Whether the utterances are consistent with the product features and customer preferences.

Table 2: Human evaluation metrics.

Model	Strategy	Real	Concise	Coherence	Natural
GPT-4	Single pass	<b>4.3</b>	<b>4.8</b>	<b>4.7</b>	<b>4.2</b>
GPT-4	Interactive	3.3	2.9	3.4	2.9
LLaMA-2	Single pass	4.1	4.7	4.1	3.9
LLaMA-2	Interactive	2.5	3.0	3.3	2.5

Table 3: Conversational-level human evaluation results measured by Realism, Conciseness, Coherence and Naturalness. Average Kendall’s  $\tau = 0.51$ .

proach takes the global generation plan for both seller and customer agents and performs dialogue generation with a single pass, thus there are less likely to be contextual conflicts. On the contrary, as our error analysis (4.5, coherence) & Appendix Table 20 show, in the interactive approach, the LLM-powered agents have no access to the generation plan beyond the next utterance. Due to the alternative generation paradigm, there is a greater potential for errors. The only exception is for LLaMA-2-70B-Chat, where the single pass approach has a higher average number of non-script-following utterances than the interactive setting. This is likely because LLaMA-2-70B-Chat tends to follow instructions less closely than GPT-4 when crafting a coherent conversation.

## 4.4 Wizard-of-Shopping (WoS) Dataset

Since human evaluation results validate that TRACER with GPT-4 single pass generation can generate high-quality conversations, we propose to release an official, larger scale WoS dataset to (1) promote more research in CPS and (2) study the effectiveness of generated datasets in real downstream tasks. Table 5 shows the statistics of our

Model	Strategy	Unrealistic ↓	Not Following ↓
GPT-4	single pass	<b>0.8</b>	<b>0.1</b>
GPT-4	interactive	2.0	0.4
LLaMA-2	single pass	1.2	1.7
LLaMA-2	interactive	2.9	1.1

Table 4: Utterance-level human evaluation results in terms of the number of labeled unrealistic and non-script-following utterances. The lower, the better. Average Kendall’s  $\tau = 0.44$ .

<b>Total # conv</b>	3600
<b>Domains</b>	Electronics Home & Kitchen Beauty & Personal Care
<b># Conv per domain</b>	1200
<b>avg # utterance per conv</b>	19.7
<b>avg # search per conv</b>	2.2
<b>avg generation time per conv</b>	75.6s

Table 5: Statistics of *Wizard of Shopping* dataset.

*WoS* dataset. A typical example of a generated conversation is shown in Appendix Table 18. In contrast to prior conversational e-commerce datasets that are either large-scale but unnatural (Zou et al., 2022), or natural but small-scale (Bernard and Balog, 2023), our *WoS* dataset consisting of 3600 conversations is large-scale, and is designed to be human-like and realistic. Further, since our TRACER approach can be easily applied to other product domains to generate additional conversations, we will continue to expand the *WoS* dataset in terms of quantity and number of product domains<sup>5</sup>.

#### 4.5 Error Analysis of *WoS* Dataset

Despite the high naturalness and coherence of the conversations generated by TRACER, we identify the following errors in them. Appendix Table 9 shows the proportion of examples received low human evaluation scores.

**Verbosity.** In the interactive generation setting, the seller agent often “hurries” to include all the assigned product aspects in a single lengthy utterance. This behavior is primarily due to the seller agent’s inability to anticipate the potential for multiple turns in the conversation, such that it could gradually discuss one or two product aspects per utterance, as instructed in the prompt. Consequently, the interactive setting often produces much lengthier utterances compared to the single-pass generation, adversely impacting realism, conciseness, and naturalness scores. As the examples in Tables 19 &

<sup>5</sup>We will release the dataset and dialogue generation scripts upon acceptance.

20 show, the seller tends to produce overly lengthy utterances.

**Coherence.** We encourage the customer agent to spontaneously ask clarification questions in the conversation in order to improve the naturalness of the conversation. In interactive generation settings, when the customer asks clarification questions, the product value selection is usually left unfinished due to the customer’s confusion. In such cases, the seller agent usually also successfully responds to the customer’s questions. However, the seller does not give the customer a second chance to make a value selection response and instead directly moves to asking about new product aspects (See sentences in bold in Table 20). Consequently, this abrupt transition significantly undermines the coherence, realism, and naturalness of the conversation.

**Bad features.** The extracted aspect-value pairs from the product catalog are not perfect (Putthivithya and Hu, 2011; Raju et al., 2009; Yang et al., 2022), with some of them being very unnatural, despite the manual cleaning we performed and explained in §4.1. When such unnatural or unrealistic product aspects are selected by the decision tree, the LLMs are forced to verbalize them. In such cases, certain utterances may look awkward. For example, the inappropriately selected product aspects (in bold) in Table 21 make the conversation sound awkward.

**Not following instructions.** The LLMs may not always follow the prompts to verbalize selected product features and preferences faithfully. For example, the seller agent may ask about product aspects that have not been explicitly assigned from the catalog, based on the commonsense understanding of the LLM. Similarly, a customer agent may invent requirements that are not present in the assigned product values from the catalog. Additionally, the customer agent may not be precise in conveying their preferences, sometimes responding with statements like “this is an optional feature for me” instead of “I do not want this feature”. In Table 22, all the bolded unassigned product features are invented by the LLM. We quantify this issue in the “Not Following” column in Table 4.

#### 4.6 Comparison with Prior Dataset

This section compares *WoS* to MG-ShopDial (Bernard and Balog, 2023) which is the only publicly available e-commerce dialogue dataset.

**Annotation Protocol.** MG-ShopDial is annotated by crowdsourcers, with some annotators acting as shopping assistants and others as customers. Each role is given a checklist of actions to complete, ensuring that the annotated dataset encompasses a variety of intents. While *WoS* is initially generated by TRACER, we later demonstrate that through careful dialogue planning, our simulated dialogues can also encompass diverse intents. Additionally, because customer preferences are randomly sampled and vary in complexity, each dialogue may have a different “checklist”, resulting in greater diversity at the corpus level compared to MG-ShopDial.

**Product Catalog.** MG-ShopDial offers a curated catalog to annotators, which covers around 14 items per category. However, this limited selection may not accurately reflect the breadth of choices available in real-world shopping scenarios. As detailed in §4.1, we index over 236k product and the utterance generation in *WoS* depends on the most recent search results. Therefore, the dialogue generation process of TRACER closely resembles the iterative “search-and-refine” procedure of typical shopping experiences.

**Intent.** We use OpenAI GPT-4<sup>6</sup> to label each utterance with intents defined by Bernard and Balog (2023). See more setup details in Appendix D. As Table 12 shows, similar to MG-ShopDial, our *WoS* conversations exhibit diverse distributions of intents, which is consistent with the high naturalness observed in §4.3. Furthermore, MG-ShopDial has significantly more “Recommend” intent than *WoS*, while *WoS* has a dominating proportion of “Elicit preferences” intent. This correlates with our observation that the wizard (seller) in MG-ShopDial memorizes all candidate products due to a limited catalog size and starts the product recommendation stage early without explicit product search. In contrast, with the decision tree for dialogue planning, *WoS* conversations aim to narrow down product aspect values for product search from a real product catalog and adopt the strategy of not selecting any product until the search is converged. As a result, *WoS* conversations are significantly shorter than MG-ShopDial (19.7 vs. 34.3 utterances).

**Conversation styles.** Since MG-ShopDial is generated by humans with keyboard typing, their sentences are usually more brief (8.5 words per utter-

ance), with frequent absence of concluding punctuation. Specifically, only 256 of its 2,196 utterances conclude with a period, while 479 end with a question mark. In contrast, utterances in *WoS* tend to be lengthier, averaging 22.8 words per utterance, characterized by near-perfect grammar (as illustrated in Appendix Table 18). As a result, MG-ShopDial conversations tend to be more lengthy and casual than *WoS*.

**Expandability.** For integrating additional domains into MG-ShopDial, it is necessary to curate a new list of products and potentially retrain annotators. Our proposed TRACER can seamlessly extend to generating shopping dialogues for any new domains once the products are indexed.

## 5 Downstream Tasks

In this section, we show the effectiveness of the Wizard-of-Shopping dataset by applying it to two downstream tasks: Conversational Query Generation (CQG) and Conversational Product Ranking (CPR). For both tasks, we divided the *WoS* conversations into 3,000 for training, 300 for validation, and 300 for testing, with each split containing an equal number of dialogues from the three domains. More details are available in Appendix E.

### 5.1 Conversational Query Generation (CQG)

Real product search conversations are likely to be verbose and redundant for the purpose of product ranking. Similar to the conversational search (Yu et al., 2020; Vakulenko et al., 2021; Chen et al., 2022; Wu et al., 2022) where a reformulation model is used to generate a more informative query from the dialogue history, we address *conversational query generation* (CQG) for product search. CQG aims to extract essential information such as the product category under discussion, desired features, undesired features, and optional product preferences from the customer. This extracted information can then be used as a query for a product search engine. And in fact, this is a reverse task of the LLM verbalization where we extract user preferences from the shopping dialogues. Similar to Dialogue State Tracking (Lu et al., 2023), CQG is helpful in tracing the interest of customers during the conversation. As Bi et al. (2019) suggests, utilizing positive and negative product features are crucial for training a conversational product ranking model. We train a conversational query generator on the *WoS* data to extract the product category and the

<sup>6</sup>[gpt-4-0125-preview](#)



Approach	F1	R-1	R-2	R-L
Baseline	0.008	0.137	0.047	0.087
D2Q	0.834	0.899	0.822	0.873
D2Q (GPT-4)	0.553	0.793	0.628	0.734

Table 6: Conversational Query Generation performance at the final turn measured by exact-F1, ROUGE-1, -2, and -L.

customer-preferred product features:

$$[PC, Wanted, Unwanted, Optional] \leftarrow CQG(Dialogue) \quad (2)$$

### 5.1.1 Approaches

**Baseline.** As stated in §2, to the best of our knowledge, *WoS* is the only CPS dataset that can be used for training downstream tasks. Therefore as a simple baseline, we directly use the full conversation history as the predicted query.

**Dialogue → Query (D2Q).** We use a seq2seq model, Longformer Encoder-Decoder (LED; Beltagy et al., 2020) fine-tuned on our *WoS* dataset to predict the product category and the *wanted*, *unwanted* and *optional* product features, given the full conversation history.

**D2Q (GPT-4).** For reference of LLM performance, we few-shot prompt OpenAI gpt-4-turbo-2024-04-09 with the same inputs and outputs as D2Q.

### 5.1.2 Experimental Results

Table 6 shows the performance of the CQG task. As expected, the weak baseline using conversation history performs poorly, while our fine-tuned query generator performs much better. Although GPT-4’s performance is slightly underestimated because few-shot demonstration examples do not show every edge case, D2Q by *WoS*-finetuned LED shows its superiority of both performance and cost in terms of both latency and price over GPT-4 as an example of LLM.

## 5.2 Conversational Product Ranking (CPR)

As the core component of CPS, CPR directly ranks the candidate products given the shopping conversation. We index or embed each product’s title and concatenate their aspect-value pairs according to the following approaches:

### 5.2.1 Approaches

**Baseline** We directly feed the full conversation history as the query to a BM25 ranker.

Approach	Ranker	MRR	H@1	H@10	H@100
Baseline	BM25	0.162	0.107	0.203	0.587
D2Q2P	BM25	<b>0.838</b>	<b>0.767</b>	<b>0.927</b>	<b>0.993</b>
D2Q2P	Roberta	0.675	0.583	0.780	0.937
D2Q2P (GPT-4)	BM25	0.763	0.680	0.903	0.903

Table 7: Downstream Conversational Product Ranking performance at the final turn.

**Dialogue → Query → Product (D2Q2P)** We use D2Q setting in §5.1.1 to generate queries given the full conversation history and apply the queries to a ranker, i.e.  $Products = Ranker(CQG(Dialogue))$ . We experiment with both a sparse ranker, BM25, and a dense ranker (fine-tuned RoBERTa; Liu et al., 2019).

**D2Q2P (GPT-4)** We use D2Q (GPT-4) setting in §5.1.1 and feed the predicted query to a BM25 ranker.

### 5.2.2 Experimental Results

Table 7 shows the mean reciprocal rank (MRR) and Hit@*k* of all methods. Similar to the trend observed for CQG, our D2Q2P approach significantly outperforms the baseline and GPT-4. Interestingly, the sparse BM25 ranker greatly outperforms the dense RoBERTa-based ranker. This is because the product representations include feature names that are lexically similar to the gold queries. Consequently, BM25 exhibits a strong performance in this ranking task. Moreover, the dense ranker may be under-fitted. The query generators and dense rankers fine-tuned on our *WoS* dataset significantly outperform the baseline that is not trained on our dataset.

## 6 Conclusions

We propose a method, TRACER, to automatically generate target-oriented shopping conversations without any human annotations. We leverage decision tree to explore the vast product search space, and construct a dialogue plan that minimizes the number of search steps required to retrieve a relevant product. The resulting corpus (*WoS*), generated using single-pass approach with GPT-4, not only achieved highly natural (4.2/5.0) and coherent (4.7/5.0) ratings from human annotators, but also showed substantial improvements when applied to two downstream tasks. By releasing our dataset and approach, we hope to expedite the research and development of intelligent CPS systems in future.

## Limitations

Despite the demonstrated potential of our TRACER approach and the *WoS* dataset, there are still some limitations that we leave for future work. First, TRACER does not consider more complicated shopping behavior such as comparison among similar products (Vedula et al., 2022). In the future, we could extend the ability of TRACER and allow customers to choose a similar product in the search results for comparison with assistance from a LLM agent. Second, the quality of *WoS* conversations relies on the quality of the product catalog. Since noise in the product catalog can directly propagate to the generated conversations. Methods proposed to curate and clean the product catalog (Ghani et al., 2006; Yang et al., 2022; Vedula et al., 2022) could be applied to further clean a product catalog. Finally, since the decision tree treats every product aspect and value as a categorical label, the semantics of product features are under-utilized. Future work may further leverage the zero-shot ability of LLMs to mine the semantics of the product aspects, so that similar product features may be merged to produce a more natural and efficient dialogue plan.

## Ethical Statements

The potential risk of our work could be biased simulations of real-world users' behavior, such that the learned downstream CQG and CPR models are biased. As a result, the end user's will of shopping might be distorted by the CPS system. However, we note that this risk is hypothetical.

## References

- Azzah Al-Maskari and Mark Sanderson. 2010. A review of factors influencing user satisfaction in information retrieval. *Journal of the American Society for Information Science and Technology*, 61(5):859–868.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Nolwenn Bernard and Krisztian Balog. 2023. Mghospital: A multi-goal conversational dataset for e-commerce. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan, SIGIR '23.
- Keping Bi, Qingyao Ai, Yongfeng Zhang, and W Bruce Croft. 2019. Conversational product search based on negative feedback. In *Proceedings of the 28th acm international conference on information and knowledge management*, pages 359–368.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. *MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling*. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Meng Chen, Ruixue Liu, Lei Shen, Shaozu Yuan, Jingyan Zhou, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2020. *The JDDC corpus: A large-scale multi-turn Chinese dialogue dataset for E-commerce customer service*. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 459–466, Marseille, France. European Language Resources Association.
- Zhiyu Chen, Jie Zhao, Anjie Fang, Besnik Fetahu, Oleg Rokhlenko, and Shervin Malmasi. 2022. *Reinforced question rewriting for conversational question answering*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 357–370, Abu Dhabi, UAE. Association for Computational Linguistics.
- Yang Deng, Yaliang Li, Fei Sun, Bolin Ding, and Wai Lam. 2021. Unified conversational recommendation policy learning via graph-based reinforcement learning. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1431–1441.
- Rayid Ghani, Katharina Probst, Yan Liu, Marko Krema, and Andrew Fano. 2006. Text mining for product attribute extraction. *ACM SIGKDD Explorations Newsletter*, 8(1):41–48.
- Ting Han, Ximing Liu, Ryuichi Takanabu, Yixin Lian, Chongxuan Huang, Dazhen Wan, Wei Peng, and Minlie Huang. 2021. Multiwoz 2.3: A multi-domain task-oriented dialogue dataset enhanced with annotation corrections and co-reference annotation. In *Natural Language Processing and Chinese Computing: 10th CCF International Conference, NLPCC 2021, Qingdao, China, October 13–17, 2021, Proceedings, Part II 10*, pages 206–218. Springer.
- Shirley Anugrah Hayati, Dongyeop Kang, Qingxi-aoyang Zhu, Weiyang Shi, and Zhou Yu. 2020. Inspired: Toward sociable recommendation dialog systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8142–8152.
- Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. Challenges and applications of large language models. *arXiv e-prints*, pages arXiv–2307.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. *Dense passage retrieval for open-domain question answering*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. *BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards deep conversational recommendations. *Advances in neural information processing systems*, 31.
- Siheng Li, Cheng Yang, Yichun Yin, Xinyu Zhu, Zesen Cheng, Lifeng Shang, Xin Jiang, Qun Liu, and Yujiu Yang. 2023. *AutoConv: Automatically generating information-seeking conversations with large language models*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1751–1762, Toronto, Canada. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019.

- Roberta: A robustly optimized bert pretraining approach.
- Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, and Wanxiang Che. 2021. Durecdial 2.0: A bilingual parallel corpus for conversational recommendation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4335–4347.
- 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.
- Bo-Ru Lu, Nikita Haduong, Chia-Hsuan Lee, Zeqiu Wu, Hao Cheng, Paul Koester, Jean Utke, Tao Yu, Noah A Smith, and Mari Ostendorf. 2023. Dialgen: Collaborative human-lm generated dialogues for improved understanding of human-human conversations. *arXiv preprint arXiv:2307.07047*.
- Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 43–52.
- Manish Mehta, Rakesh Agrawal, and Jorma Rissanen. 1996. Sliq: A fast scalable classifier for data mining. In *Advances in Database Technology—EDBT’96: 5th International Conference on Extending Database Technology Avignon, France, March 25–29, 1996 Proceedings 5*, pages 18–32. Springer.
- Duangmanee Putthividhya and Junling Hu. 2011. Bootstrapped named entity recognition for product attribute extraction. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1557–1567.
- J Ross Quinlan. 2014. *C4. 5: programs for machine learning*. Elsevier.
- Santosh Raju, Prasad Pingali, and Vasudeva Varma. 2009. An unsupervised approach to product attribute extraction. In *Advances in Information Retrieval: 31th European Conference on IR Research, ECIR 2009, Toulouse, France, April 6-9, 2009. Proceedings 31*, pages 796–800. Springer.
- Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2021. Question rewriting for conversational question answering. In *Proceedings of the 14th ACM international conference on web search and data mining*, pages 355–363.
- Nikhita Vedula, Marcus Collins, Eugene Agichtein, and Oleg Rokhlenko. 2022. What matters for shoppers: Investigating key attributes for online product comparison. In *European Conference on Information Retrieval*, pages 231–239. Springer.
- Chien-Sheng Wu, Steven C.H. Hoi, Richard Socher, and Caiming Xiong. 2020. *TOD-BERT: Pre-trained natural language understanding for task-oriented dialogue*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929, Online. Association for Computational Linguistics.
- Zequi Wu, Yi Luan, Hannah Rashkin, David Reitter, Hannaneh Hajishirzi, Mari Ostendorf, and Gaurav Singh Tomar. 2022. Conqrr: Conversational query rewriting for retrieval with reinforcement learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10000–10014.
- Liqiang Xiao, Jun Ma, Xin Luna Dong, Pascual Martínez-Gómez, Nasser Zalmout, Chenwei Zhang, Tong Zhao, Hao He, and Yaohui Jin. 2021. *End-to-end conversational search for online shopping with utterance transfer*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3477–3486, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Li Yang, Qifan Wang, Zac Yu, Anand Kulkarni, Sumit Sanghai, Bin Shu, Jon Elsas, and Bhargav KanaGal. 2022. Mave: A product dataset for multi-source attribute value extraction. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, pages 1256–1265.
- Steve Young, Milica Gašić, Simon Keizer, François Mairesse, Jost Schatzmann, Blaise Thomson, and Kai Yu. 2010. The hidden information state model: A practical framework for pomdp-based spoken dialogue management. *Computer Speech & Language*, 24(2):150–174.
- Shi Yu, Jiahua Liu, Jingqin Yang, Chenyan Xiong, Paul Bennett, Jianfeng Gao, and Zhiyuan Liu. 2020. Few-shot generative conversational query rewriting. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 1933–1936.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022a. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Yiming Zhang, Lingfei Wu, Qi Shen, Yitong Pang, Zhihua Wei, Fangli Xu, Bo Long, and Jian Pei. 2022b. Multiple choice questions based multi-interest policy learning for conversational recommendation. In *Proceedings of the ACM Web Conference 2022*, pages 2153–2162.
- Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W Bruce Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th acm international conference on information and knowledge management*, pages 177–186.



Nan Zhao, Haoran Li, Youzheng Wu, and Xiaodong He. 2022. [JDDC 2.1: A multimodal Chinese dialogue dataset with joint tasks of query rewriting, response generation, discourse parsing, and summarization](#). In [Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing](#), pages 12037–12051, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Jie Zou, Jimmy Huang, Zhaochun Ren, and Evangelos Kanoulas. 2022. Learning to ask: Conversational product search via representation learning. [ACM Transactions on Information Systems](#), 41(2):1–27.

## A More Related Work

### A.1 Dialogue Generation with LLMs

Collecting dialogue datasets for conversational applications requires significant effort from crowd workers (Budzianowski et al., 2018; Wu et al., 2020), incurring substantial financial and time costs. Hence, alternative simulation approaches such as DIALOGIC were developed for dialogue generation using GPT-3 (Brown et al., 2020) on the MultiWOZ dataset (Han et al., 2021). It successfully generated dialogues with near-human fluency, even when trained with as few as 85 initial seed dialogues. Similarly, Li et al. (2023) proposed to fine-tune OPT-13B (Zhang et al., 2022a) to generate conversations grounded on a given document. A semi-automated dialogue generation framework based on ChatGPT (Lu et al., 2023) was proposed to generate complex conversations in the insurance domain by generating sub-dialogues and using human feedback to ensure consistency. They demonstrated that incorporating the generated dataset into existing auto insurance claim call datasets can enhance the performance of Dialogue State Tracking (DST) (Young et al., 2010).

Unlike previous methods, we do not rely on initial seed annotations; we only leverage the zero-shot capabilities of LLMs to generate conversations that are grounded in a product catalog and a simulated user profile. Once a user profile is synthesized from the product information and elicited customer preferences, a decision tree is fitted to guide the dialogue generation, ensuring that the conversation concludes with the discovery of the target product. Similar to the DST task (Lu et al., 2023), we demonstrate that our generated dataset can be used to train a query generator that extracts preferences for product attribute values (§5).

## B Detailed Human Evaluation Metric Descriptions

### B.1 Conversation-level Evaluations

The metrics are summarized in Table 2. For conversation-level evaluations, judges are asked to score each conversation on a Likert scale of 1 to 5. The meaning of each of the scores is interpreted in Table 8.

**Realism.** To distinguish counter-intuitive conversations, *realism* examines whether the conversation is likely to happen in the real world from a logical perspective. We also consider *Factuality*,

which captures the commonsense-level factual errors, since they essentially reflect similar criteria. Meanwhile, we consider whether the speakers stick to their own role, i.e. the seller should ask questions or recommend products but not buy products, while the customer should focus on answering proposed questions to narrow down products to buy, instead of recommending products. We combine these three aspects into one score since they are highly related.

**Conciseness** examines whether the conversation is too verbose, such as a single utterance contains too many aspects to discuss.

**Coherence** measures whether the conversation is coherent to the conversation history from the beginning. Additionally, we also consider fluency, as it is also closely relevant. Specifically, we examine (1) whether there is repetition or forgetting about the context; (2) whether the agents ignore the previous utterance and proceed to the next topic; (3) whether the conversation gets stuck at some point and never proceeds.

**Overall naturalness** reflects the overall subjective impression of the conversation. Opposing to the *realism* which reflects whether the conversation is logical, the naturalness asks whether the judge subjectively feels humanness from the conversation.

### B.2 Utterance-level Evaluations

We further evaluate the conversations on a finer-grained scale. Judges are asked to list all unsatisfactory utterance IDs, as a binary classification task.

**Realism.** Similar to the conversation level, we examine the realism, factuality, and consistency of the role at the utterance level.

**Script-following.** Given the assigned product aspects and values, as well as the customer’s preference (*wanted*, *optional*, or *unwanted*), we evaluate whether the utterances are consistent with the assigned features. The utterances should cover all assigned features and should not alter the type of the customer’s preference (e.g. switching *unwanted* to *optional*). Further, agents should not invent new product features that are not assigned, which may create inconsistency with the search plan and the semantics of the conversation.

Metric	Score	Interpretation
Realism	1	Impossible to happen in the real world.
	2	Unlikely to happen in the real world, but still possible.
	3	May happen in the real world, but unrealistic.
	4	Likely to happen in the real world, with minor issue.
	5	Very likely to happen in real life.
Conciseness	1	Utterances are way too long and verbose. Not acceptable.
	2	Utterances are very long and verbose.
	3	Utterances are long but still acceptable.
	4	There is a space for simplifying utterances.
	5	Utterances cannot be further simplified.
Coherence & Fluency	1	The conversation is not complete till the end.
	2	Contains serious coherence or fluency problems, but the conversation is complete.
	3	Contains obvious coherence or fluency problems.
	4	Contains minor coherence or fluency problems.
	5	Coherent and fluent conversation without any flaw.
Overall naturalness	1	Very unnatural conversation. Difficult to read.
	2	Contains obvious naturalness issues.
	3	Obviously this is machine-generated, with minor naturalness issues.
	4	No obvious issues found, but it's not as natural as humans.
	5	Good conversation. Indistinguishable from real-life conversations.

Table 8: Interpretation of conversational-level human evaluation scores.

Model	Strategy	Realism $\leq 3$	Conciseness $\leq 3$	Coherence $\leq 3$	Naturalness $\leq 3$
GPT-4	Single pass	<b>0.12</b>	<b>0</b>	<b>0.02</b>	<b>0.17</b>
GPT-4	Interactive	0.48	0.64	0.50	0.69
LLaMA-2	Single pass	0.22	0.07	0.29	0.32
LLaMA-2	Interactive	0.68	0.73	0.46	0.78

Table 9: Proportion of conversations affected by low scores in realism, conciseness, coherence, and naturalness.

**Reasonable features.** The selected product features by the decision tree may contain inappropriate or unnatural product aspects and values to be discussed. For example, the customer is not supposed to know the value of the aspect “model number”. This issue usually originates from the quality of the product metadata, which is impossible to perfectly curate. Table 4 shows the results excluding utterances with unreasonable features to avoid confusion.

## C Large Language Models for Dialogue Generation

We experiment with two different LLMs: OpenAI GPT-4<sup>7</sup> and LLaMA-2-70B-chat. As the *chatbot-arena-leaderboard*<sup>8</sup> shows, GPT-4 was the best-performing LLM by then. However, as GPT-4 is proprietary, we also experiment with LLaMA-2-70B-chat, which is one of the best performing commercially available instruction-tuned LLMs. Due to the high GPU memory requirement of LLaMA-2,

we use a 5-bit quantized version.<sup>9</sup>

## D Intent Analysis of Wizard of Shopping vs. MG-ShopDial

**Experimental settings.** Extending 4.6, to compare the differences between utterances in *WoS* and MG-ShopDial (Bernard and Balog, 2023), we compare the intent distribution between these two datasets by prompting OpenAI GPT-4<sup>10</sup> (see prompts in Table 11) to label each utterance with intents defined by Bernard and Balog (2023) (Table 10). Similar to Bernard and Balog (2023), each utterance may correspond to multiple intent labels. We sample 100 *WoS* conversations from each of the three domains for analysis. As Bernard and Balog (2023) indicated, their agreements among crowd workers are relatively low (Fleiss  $\kappa=0.187$ ), thus we predict MG-ShopDial utterances with GPT-4 as well to make a fair comparison.

<sup>7</sup><https://platform.openai.com/docs/models, gpt-4-0613>

<sup>8</sup><https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>

<sup>9</sup>[https://github.com/abetlen/llama-cpp-python, https://huggingface.co/TheBloke/Llama-2-70B-Chat-GGML/resolve/main/llama-2-70b-chat.ggmlv3.q5\\_K\\_M.bin](https://github.com/abetlen/llama-cpp-python, https://huggingface.co/TheBloke/Llama-2-70B-Chat-GGML/resolve/main/llama-2-70b-chat.ggmlv3.q5_K_M.bin)  
<sup>10</sup>[gpt-4-turbo-2024-04-09](https://huggingface.co/gpt-4-turbo-2024-04-09)

Intent	Definition
Greetings	Indicates the beginning or end of the conversation
Interaction structuring	Utterances that make the conversation structured and natural (e.g., thanking, stalling)
Disclose	The client discloses information about what they are looking for
Clarification question	The agent asks a question to make sure it understands correctly a previous statement
Other question	Asks a question that is not a clarification question (e.g., factoid, follow-up questions)
Elicit preferences	The agent asks a question to find the client’s preferences (e.g., the color of an item, the budget)
Recommend	The agent recommends one or several items to the client
Answer	A participant gives an answer to the other participant’s information request
Explain	Provides an explanation to a previous statement (e.g., justifies suggestion or rejection of an item)
Positive feedback	Expresses positive feedback (e.g., confirmation, accept a recommendation)
Negative feedback	Expresses negative feedback (e.g., disagreement, rejection of a recommendation)
Other	Does not fit other labels

Table 10: Utterance intents defined by MG-ShopDial (Bernard and Balog, 2023).

Instruction
Given an utterance with context in a conversation under e-commerce shopping scenario, identify all intents of the utterance defined below (format of "intent: explanation"):
Greetings: Indicates the beginning or end of the conversation
Interaction structuring: Utterances that make the conversation structured and natural (e.g., thanking, stalling)
Disclose: The client discloses information about what they are looking for
Clarification question: The agent asks a question to make sure it understands correctly a previous statement
Other question: Asks a question that is not a clarification question (e.g., factoid, follow-up questions)
Elicit preferences: The agent asks a question to find the client’s preferences (e.g., the color of an item, the budget)
Recommend: The agent recommends one or several items to the client
Answer: A participant gives an answer to the other participant’s information request
Explain: Provides an explanation to a previous statement (e.g., justifies suggestion or rejection of an item)
Positive feedback: Expresses positive feedback (e.g., confirmation, accept a recommendation)
Negative feedback: Expresses negative feedback (e.g., disagreement, rejection of a recommendation)
Other: Does not fit other labels
Examples (Utterance followed by intents):
{}
Contexts: {}
Utterance: “{}”
Answer all plausible intents names of the utterance above, separate with “;”. Do not return explanations.
Answer:

Table 11: Prompt format for predicting utterance intents.

## E Detailed Downstream Experiments

In this section, we extend §5 by comparing conversational-level and utterance-level generation approaches in addition to the weak baseline.

### E.1 Conversational Query Generation

#### E.1.1 Approaches

**Baseline.** As we explained in §2, to the best of our knowledge, our WoS dataset is the only CPS dataset that can be used for training downstream tasks. Therefore as a simple baseline, we directly use the full conversation history as the predicted query.

**Conversational-level.** We use a seq2seq model fine-tuned on our WoS dataset to predict the product category and the *wanted*, *unwanted* and *optional* product features, given the full conversation history.

**Utterance-level.** We assume that each utterance between the seller and customer encodes one or more product features. We use a seq2seq model to extract product features from each utterance and concatenate the features to be the final query.

**GPT-4** We few-shot prompt OpenAI gpt-4-0125-preview with the same inputs and outputs as the conversational-level setting above.

#### E.1.2 Experimental Settings

We fine-tune BART-base (Lewis et al., 2020) for the utterance-level approach and use LED-base (Beltagy et al., 2020) for the conversation-level approach due to the longer context required as input. We use a default learning rate of 5e-5 and train for 5 epochs for both models.

#### E.1.3 Experimental Results

Table 13 shows the performance of the CQG task. We report results on two settings: *wanted* feature-



Intent%	Annotated MG-ShopDial	Predicted MG-ShopDial	Wizard of Shopping
Answer	17.0	3.4	0.1
Clarification question	11.8	9.3	3.0
Explain	11.5	6.7	5.5
Interaction structuring	11.3	12.1	15.4
Other question	10.6	10.0	7.7
Positive feedback	8.3	12.1	7.1
Greetings	7.2	5.9	3.8
Recommend	7.2	<b>11.9</b>	<b>4.2</b>
Disclose	6.4	15.3	15.6
Elicit preferences	5.7	<b>9.3</b>	<b>34.2</b>
Other	1.7	0.7	0
Negative feedback	1.4	3.3	3.4

Table 12: Intent distribution of utterances in MG-ShopDial (Bernard and Balog, 2023) and Wizard of Shopping.

only setups concatenate PC and wanted features and all feature setups concatenate all features as defined in Eq. 2. As expected, the weak baseline using the conversation history performs poorly, while our trained query generator performs much better. When comparing the utterance-level approaches, LED outperforms BART, since the training data for BART is slightly smaller than LED’s. This is because we discard the input augmented conversation history that is too long for BART’s maximum context window (1024 tokens). Also, the LED-based utterance level approach under-performs the conversation level approach, presumably because integrating the results from multiple inferences passes is more error-prone.

We also experiment with GPT-4 as CQG for reference. Since the few-shot demonstration examples do not show every edge case, GPT-4’s performance is slightly underestimated. Nonetheless, D2Q by *WoS*-finetuned LED shows the superiority of both performance and cost over GPT-4.

## E.2 Conversational Product Ranking

### E.2.1 Approaches

**Baseline** Similar to the CQG task, we assume there is no conversational data available for training a ranker. Therefore, we directly feed the full conversation history as the query to a BM25 ranker.

**Dialogue → Product (D2P)** We directly embed the full conversation history to a dense ranker to rank the candidate products, i.e.  $Products = Ranker(Dialogue)$ .

**Dialogue → Query → Product (D2Q2P)** We leverage D2Q approach to generate queries given the full conversation history and apply the queries to a ranker, i.e.  $Products = Ranker(CQG(Dialogue))$ . We experiment with

both a sparse ranker, BM25, and a dense ranker.

**D2Q2P (GPT-4)** We use GPT-4 setting in §5.1.1 and feed the predicted query to a BM25 ranker.

### E.2.2 Experimental Settings

For the dense rankers, we fine-tune a Longformer-base (Beltagy et al., 2020) for the D2P setting, and a RoBERTa-based ranker<sup>11</sup> for the D2Q2P setting using a dense passage retrieval (Karpukhin et al., 2020) training approach<sup>12</sup>. For training both D2P and D2Q2P dense rankers, we use the ranking results given the corresponding gold queries as positive examples. The hard negative examples are constructed by removing the positive examples from the search results of the target product category.

### E.2.3 Experimental Results

Table 14 shows the mean reciprocal rank (MRR) and Hit@ $k$  of all methods. We observe that the D2P approach only slightly outperforms the baseline approach without training. We suspect that our DPR-Longformer is still under-fitted given 3000 conversations as the training set. On the contrary, by leveraging the query generators trained on our collected conversations, the ranking performance is greatly improved. Similar to the trend observed in the CQG task, the conversation-level approach outperforms the utterance-level approach that integrates multiple passes of inference outputs.

<sup>11</sup><https://huggingface.co/castorini/ance-msmarco-passage>

<sup>12</sup>[https://haystack.deepset.ai/tutorials/09\\_dpr\\_training](https://haystack.deepset.ai/tutorials/09_dpr_training)

Approach	QG	Feat.	F1	R-1	R-2	R-L
Baseline	-	+	0	0.056	0.020	0.048
Baseline	-	+/-/?	0.008	0.137	0.047	0.087
Utterance	BART	+	0.698	0.592	0.482	0.562
Utterance	BART	+/-/?	0.656	0.664	0.489	0.592
Utterance	LED	+	0.704	0.610	0.509	0.578
Utterance	LED	+/-/?	0.748	0.720	0.561	0.657
Conversation	LED	+	0.834	0.899	0.822	0.873
Conversation	LED	+/-/?	0.873	0.900	0.768	0.840
Conversation	GPT-4	+	0.553	0.793	0.628	0.734
Conversation	GPT-4	+/-/?	0.669	0.833	0.584	0.703

Table 13: Detailed Conversational Query Generation (QG) performance at the final turn measured by exact-F1, ROUGE-1, ROUGE-2, and ROUGE-L. We report the performance of only using the desired features (+) for downstream ranking, as well as using all features (wanted (+), unwanted (-), and optional (?)). Detailed version of Table 6.

Approach	QG	Ranker	MRR	H@1	H@10	H@100
Baseline	-	BM25	0.162	0.107	0.203	0.587
D2P	-	longformer	0.201	0.143	0.247	0.620
Utterance	BART	BM25	0.629	0.540	0.733	0.900
Utterance	BART	Roberta	0.217	0.163	0.270	0.477
Utterance	LED	BM25	0.667	0.593	0.740	0.900
Utterance	LED	Roberta	0.207	0.163	0.250	0.460
Conversation	LED	BM25	<b>0.838</b>	<b>0.767</b>	<b>0.927</b>	<b>0.993</b>
Conversation	LED	Roberta	0.675	0.583	0.780	0.937
Conversation	GPT-4	BM25	0.763	0.680	0.903	0.903

Table 14: Detailed downstream Conversational Product Ranking performance at the final turn. QG: query generator. Detailed version of Table 7.

---

**System Prompt**

You are acting as an Amazon seller to help narrow down product requirements for the customer that you are chatting with. Your output utterance is to clarify the customer's precise need by asking questions about the requested product features.

---

**Conversational history**

Conversational history:

seller: {{Utterance\_SI}}

customer: {{Utterance\_CI}}

.....

---

**Instructions about Features**

First of all, **if you are asked clarification questions by the customer, you must friendly address the customer's questions first.** In this case, you must give the customer an opportunity to answer unanswered questions before you proceed. This means you may disregard the following new features. Otherwise, you must ask the customer about the following specific aspect of the requested product category provided by the output of the "update\_dialogue\_state" function. In your response, pick one or two aspects to ask from "remaining\_feature", so that the conversation is most natural, e.g. asking about general aspects first and then diving into detailed aspects. Once these one or two aspects are picked, your questions must strictly cover all of the words in the aspects or aspect value pairs. **Integrate the examples and values provided naturally into your utterance.** Do not ask about other aspects not covered above. If there is a conflict between your commonsense and the following aspect value pairs, you must prioritize the following provided aspect. You must ask in a concise and natural way. Do not repeat what you, the seller, have already said before. Keep the question brief. Do not talk about the budget unless you are asked to. Do not repeat the product name/category over and over.

---

**Instructions**

What is the next utterance for you as a seller? Only reply with the response that you are supposed to make, not the customer's response. Don't start with your role, i.e. seller, just reply with the response. Remember, as a seller, you want to ask clarification questions to the customer. You never want to get products recommended or buy products from the customer. Be friendly to the customer, so that if the customer has any questions about the product or product aspect, answer patiently. **Treat the customer as if an old grandma has a limited understanding of the products.**

---

**Function Call Templates**

Function name: update\_dialogue\_state

Parameters: {

**remaining\_features:** {{Aspect<sub>1</sub> : Expl<sub>1</sub>, Expl<sub>2</sub>, Expl<sub>3</sub>}},...

    mentioned\_positive\_features: {{Aspect<sub>2</sub> : Value<sub>2</sub>}},...

    mentioned\_optional\_features: {{Aspect<sub>3</sub> : Value<sub>3</sub>}},...

    mentioned\_negative\_features: {{Aspect<sub>4</sub> : Value<sub>4</sub>}},...

}

---

Table 15: Prompt format for interactive seller utterance generation. Prompts in bold are designed for improving conversation naturalness (§3.4).

<b>System Prompt</b>
You are a <b>lazy</b> customer on Amazon chatting with a seller.
<b>Persona</b>
<b>You have a limited understanding of the product to buy, so ask the seller about what the aspects mean when encountering a product aspect that an old grandma may not understand.</b>
<b>Conversational history</b>
Conversational history:
customer: {{ <i>Utterance_CI</i> }}
seller: {{ <i>Utterance_SI</i> }}
.....
<b>Instructions</b>
Given the utterance you just made, make an update to the specific aspects provided by the output of the <code>update_dialogue_state</code> function. For the aspect or aspect value pairs under “remaining_positive_features” below, you were asked to explicitly say that you like them in a <b>natural way</b> , so here you must pass these required aspect or aspect value pairs covered by your utterance under the “mentioned_positive_features” argument of function <code>update_dialogue_state</code> . Similarly, for the aspect or aspect value pairs under “remaining_optional_features” below, you were explicitly asked to express that they are optional or you are not concerned about them in a natural way, so here you must pass these optional aspect or aspect value pairs covered by your utterance under the “mentioned_optional_features” argument of function “ <code>update_dialogue_state</code> ”. Finally, for the aspect or aspect value pairs under “remaining_negative_features” below, you were explicitly asked to express that you don’t want them in a natural way, so here you must pass these negative aspect or aspect value pairs covered by your utterance under the “mentioned_negative_features” argument of functional <code>update_dialogue_state</code> .
<b>Function Call Templates</b>
Function name: <code>update_dialogue_state</code>
Parameters: {
remaining_positive_features: {{ <i>Aspect</i> <sub>1</sub> : <i>Value</i> <sub>1</sub> }},...
remaining_optional_features: {{ <i>Aspect</i> <sub>2</sub> : <i>Value</i> <sub>2</sub> }},...
remaining_negative_features: {{ <i>Aspect</i> <sub>3</sub> : <i>Value</i> <sub>3</sub> }},...
mentioned_positive_features: {{ <i>Aspect</i> <sub>4</sub> : <i>Value</i> <sub>4</sub> }},...
mentioned_optional_features: {{ <i>Aspect</i> <sub>5</sub> : <i>Value</i> <sub>5</sub> }},...
mentioned_negative_features: {{ <i>Aspect</i> <sub>6</sub> : <i>Value</i> <sub>6</sub> }},...
}

Table 16: Prompt format for interactive customer utterance generation. Prompts in bold are designed for improving conversation naturalness (§3.4).



<b>Instruction</b>
<p>You are a scriptwriter. For the aspects or aspect value pairs below, write a chat conversation between the customer and seller about the seller trying to narrow down the customer's need. <b>The seller knows well about the sold products, while the customer has limited knowledge about the products.</b> This means the customer should not voluntarily propose what he wants, but passively respond to the aspect-values that the seller proposes. Therefore, the seller proactively asks clarification questions while the customer clarifies the needs. The customer should speak out first and say he/she wants to buy <code>{{ProductCategory}}</code>. The seller asks about the aspect and the customer makes clarification on whether that aspect is something he/she wants, or it is optional. You MUST cover ALL aspect-value pairs below, including the wanted, unwanted, and optional aspect-value pairs. <b>Make sure you reorganize and reorder the aspect value pairs and translate them into a natural conversation in a meaningful way.</b> Start by talking about general product features and gradually transit into product-specific features. The seller should not ask too many aspects in one question. Instead only mention one or two aspects per question. Do not invent new aspects that are not mentioned below to talk about. Your script must be coherent with the conversation history, and there might be more scripts to be continued in the future, so don't terminate the conversation. Only output utterances, no non-verbal actions.</p>
<b>Wanted / Optional / Unwanted Features</b>
<p>Aspect or aspect value pairs that the customer must say {he wants to have / is optional / he does not want to have} for the target product:  Aspect: <code>{{Aspect}}</code>, Value: <code>{{Value}}</code>;</p>
<b>Candidate Value Examples</b>
<p><b>Additionally, there are some common values for each aspect to be mentioned. The seller should use these typical values in his question:</b>  Aspect: <code>{{Aspect}}</code>, Typical Values: <code>{{Value<sub>1</sub>}}</code>, <code>{{Value<sub>2</sub>}}</code>, <code>{{Value<sub>3</sub>}}</code>;</p>
<b>Format Instruction</b>
<p>Output in the format of  customer: AAAAA  seller: BBBBB  customer: CCCCC  seller: DDDDD</p>
<b>Final Instruction</b>
<p>Output a conversation following all instructions above and make sure you double check if ALL aspect value pairs above are mentioned:</p>

Table 17: Prompt format for single pass dialogue generation approach. Prompts in bold are designed for improving conversation naturalness (§3.4).

<b>Product Category</b>	lipstick
<b>Wanted Features</b>	color: dynamite red
<b>Optional Features</b>	skin type; customer review; finish type; price; special feature
<b>Unwanted Features</b>	brand: Gocheaper
<b>Speaker</b>	<b>Utterance</b>
customer	Hi there, I'm looking to buy a <b>lipstick</b> but I'm not quite sure where to start.
seller	Hello! Sure, I'd be happy to help you find the perfect lipstick. Let's start with the <b>color</b> you prefer. Popular ones are <b>red, black or clear</b> .
customer	Oh, I think I'd like <b>dynamite red</b> .
seller	Dynamite red sounds fabulous! Now let's talk about the <b>brand</b> . Some popular ones include <b>Revlon, NYX Professional Makeup, and Maybelline New York</b> . I noticed that there's also a brand named <b>Gocheaper</b> . Any preference regarding the brand?
customer	Oh, I <b>don't want</b> the brand <b>Gocheaper</b> for sure. As for the other brands, I have limited knowledge about them so I have no preference.
seller	That's perfectly fine. Now, on to the next aspect which is the <b>form</b> of the lipstick. It comes in <b>cream, stick or liquid</b> form. Do any of these appeal to you?
customer	I'd prefer <b>not</b> to have a <b>liquid lipstick</b> .
seller	Noted, we'll avoid the liquid lipsticks. Now lipsticks come in different <b>finishes</b> . Some have a <b>glossy</b> finish, other have a <b>cream</b> finish and there are also <b>matte</b> ones. Do you have a preference for the finish type?
customer	<i>Oh, is that like how shiny the lipstick is when you apply it?</i>
seller	<i>Yes, exactly. Glossy finishes are shinier, while matte finishes have no shine and cream finishes are somewhere in between.</i>
customer	I see, thank you for explaining that. I'm open to trying different finishes, so <b>no specific preference</b> there.
seller	You're welcome, that's great to know! Lipsticks can also have some <b>special features</b> like being <b>portable, long-lasting, or durable</b> . Is there anything among these that you think is important?
customer	Portable, long-lasting and durable all sound nice but <b>not mandatory</b> for me.
seller	Good to know. The lipsticks are also categorized based on <b>skin type</b> . There's the <b>all, normal and dry</b> skin type. What's your skin type?
customer	<b>I don't think that's something that would affect my choice, as I don't know my skin type.</b>
seller	No problem, we can skip that. Lastly, <b>customer review</b> and <b>price</b> are often considered. What are your thoughts on these aspects?
customer	Well, <b>I don't really look at customer reviews and I'm flexible regarding the price.</b>
seller	It sounds like you're quite flexible, which makes it easier to find the right product for you. Given your preferences, might I suggest Maybelline New York Color Sensational Red Lipstick Matte Lipstick, Dynamite Red? This is based on your preference for dynamite red color and it's not liquid. The brand is well established and it's a fairly popular one.
customer	That sounds perfect. Thank you for your help!

Table 18: A single-pass generation conversation from GPT-4. The dialogue plan is displayed at the top of the table, and the pre-determined customer preference is reflected in the conversation with **bold**, and the clarification question and answer are *italicized*.

<b>Product Category</b>	cable organizer bags & cases
<b>Wanted Features</b>	color: black
<b>Optional Features</b>	form factor; size; price; material type
<b>Unwanted Features</b>	brand: Sunsir
<b>Speaker</b>	<b>Utterance</b>
customer	Can you recommend a good cable organizer bag or case for me? I prefer one with a lot of compartments.
seller	Sure, I'd be happy to assist you in finding the perfect cable organizer. Could you please let me know your preferred color for the organizer? Would you like a black, grey, or maybe even a pink one? And how about the form? Do you prefer it to be in the form of a bag, case, or a sleeve?
customer	I'll go with black as my color preference. And what do you mean by form factor?
seller	Sure, I understood your color preference. With regards to the form factor, it refers to the design and physical shape of the cable organizer. For example, some are designed like a bag with a zipper, some come in the form of a hard case, while others are like a sleeve. Could you please tell me about your brand preference? There are several well-known brands such as Bagsmart, Bubm, and Inateck.
customer	I don't have a specific brand preference, but I don't want Sunsir.
seller	I see, let's narrow down the options further. Regarding the size of the cable organizer, are you looking for a large one that can store many items, a small, compact one for easier carry, or a one-size fits all type of organizer?
customer	I'm flexible with the size.
seller	Great! Now, considering the material of the cable organizer, do you have a preference? Some people like nylon ones for the durability, while others may choose polycarbonate or plastic options for their lightweightness. Do any of these resonate with your needs?
customer	What do you mean by material type?
seller	By "material type", I meant the fabric or substance that the organizer is made from. Cable organizers can be made from a variety of materials. Here are a few examples: 1. <b>**Nylon**</b> : This is a type of synthetic fabric that's known for its strength and durability. It's also generally resistant to dirt and easy to clean. 2. <b>**Polycarbonate or Plastic**</b> : Some organizers are designed with a hard exterior for added protection. These types are lightweight, and plastic ones are particularly good for water resistance. 3. <b>**Leather**</b> : There are also organizers made from leather or PU leather. They offer a sleek, stylish look while also providing a good level of protection. Each material comes with its own advantages. Nylon is usually longer-lasting, polycarbonate or plastic is lighter and waterproof, and leather gives a classy touch. Do any of these materials align with your preferences?
customer	I think I would like an organizer made of nylon for its durability. Is there any specific one you would recommend that meets my criteria? It should be black, not in the Sunsir brand, flexible in size, and made of nylon.
seller	Based on your preferences for a black, nylon-made cable organizer that's not from the Sunsir brand, I have a few suggestions for you: 1. The first is the [BAGSMART Universal Travel Cable Organizer](https://www.amazon.com/BAGSMART-Universal-Organizer-Electronics-Accessories/dp/B017SKRWL4). It's a great choice with several compartments to store your items such as a 9.7 inch iPad, Kindle, and Power Adapter. It's black with a hint of red color that gives it an edgy look. 2. The second option is [Inateck Electronics Organizer](https://www.amazon.com/Inateck-Electronics-Organizer-Splashproof-Waterproof/dp/B07ZV2CZSR). It's also a black, splashproof and waterproof travel cable organizer that should serve your needs well. 3. The third one is [Luxtude Electronics Organizer](https://www.amazon.com/Luxtude-Electronics-Organizer-Portable-Essentials/dp/B086M63FXV). This Travel Cable Case is a portable tech bag perfect for cable storage, cord storage and storing electronic accessories/USB/SD/Charger. It's also black, aligning with your color preference. Please let me know if you need more information about these products or if there's anything else I can assist you with.
customer	The BAGSMART Universal Travel Cable Organizer sounds like a good fit for me. I like the number of compartments it has and the aesthetic design. I'll go ahead and add it to my cart. Thank you very much for your assistance!

Table 19: Bad conciseness example by GPT-4 interactive generation. The seller's utterances are verbose.

<b>Product Category</b>	fitted sheets
<b>Wanted Features</b>	included components: 1 fitted sheet; color: natural; number of pieces: 1
<b>Optional Features</b>	material; brand; fabric weight; size
<b>Unwanted Features</b>	None
<b>Speaker</b>	<b>Utterance</b>
customer	Hello there, I am interested in buying fitted sheets. Could you please recommend some for me?
seller	Could you please specify the size of the sheets you're looking for? For instance, do you need queen, twin, or full size? Also, could you share your preference for the material - would you prefer microfiber, cotton, or other fabric types?
customer	I'd like one set of fitted sheets. As for the size and material, I don't have any particular preferences. However, could you please explain what you mean by "microfiber" material? I am not very familiar with this term.
seller	Certainly! Microfiber is a type of synthetic material known for its softness, durability, and resistance to wrinkles and stains. <b>For the fitted sheets, would you prefer a pack with multiple pieces or just one sheet?</b>
customer	Just one sheet is enough for me.
seller	Given that you are looking for just one fitted sheet, would you have any color preferences or special design requirements for the sheet? Also, do you have any specific washing instructions in mind, such as machine wash or hand wash for your sheets?
customer	As for the color, I would prefer something dark, maybe black or navy. As for washing, I think being machine washable would be most convenient for me.
seller	I see that you have an inclination towards dark colors and machine washable fitted sheets. Moving forward, are you interested in any particular brand like Bare Home, Elegant Comfort, or Ambesonne? Also, do you have any preference for the weight of the fabric? For example, would you prefer 130, 110, or 100 grams per square meter fabric weight?
customer	I don't really care about the brand or the weight of the fabric. I just need it to include one fitted sheet. Can you help me understand what "fabric weight" means?
seller	Sure, I'd be happy to explain. The fabric weight is usually measured in grams per square meter (GSM). It's an indicator of the fabric's thickness and heaviness. A higher GSM means a thicker and heavier fabric. <b>Now, you mentioned that you only need one fitted sheet.</b> So, in your case, we want to ensure that the set you buy includes only one fitted sheet, instead of multiple pieces or additional items like pillowcases. Does that sound about right to you? Would you prefer a set that exclusively includes one fitted sheet, without any additional components?
customer	Yes, that's right. I only need a set that includes one fitted sheet without any additional components.
seller	I see that you prefer dark colors. However, I'd like to clarify if you'd only like black or navy, or if you're open to other dark colors like grey or dark blue as well? Also, would you be interested in a white color in case we don't have the dark color you prefer in stock?
customer	I am open to other dark colors. However, I don't want a white one. I prefer the sheet to have a more natural look.
seller	Great, I understand that you are looking for a dark-colored, machine-washable fitted sheet with a natural look. Taking into account your preferences, I'd like to recommend the BIOWEAVES 100% Organic Cotton Twin Fitted Sheet. This sheet is made from 100% organic cotton with a 300 thread count. It has a sateen weave, which gives it a soft, luxurious feel. The sheet comes in an undyed natural color which should fit your description of a 'natural look'. It is also GOTS certified and comes with a deep pocket design. Furthermore, it's machine washable for easy care. Priced at \$32.99, it's a great value for the quality. Would you like to proceed with this one?
customer	Sure, that sounds perfect. I am happy with the BIOWEAVES 100% Organic Cotton Twin Fitted Sheet you recommended. But I am not certain about what 'GOTS certified' and 'sateen weave' mean. Could you explain these terms before I make a purchase?

Table 20: Bad coherence and conciseness example by GPT-4 interactive generation. Bolded sentences are those the seller does not give the customer a second chance to make a value selection response and directly jumps into asking about new product aspects.

<b>Product Category</b>	countertop blenders
<b>Wanted Features</b>	price: between \$50 and \$100; is dishwasher safe: yes; blade material: stainless steel
<b>Optional Features</b>	country of origin; power source; number of speeds; department; color; voltage; style; wattage; included components; customer review; capacity
<b>Unwanted Features</b>	brand: Kacsoo
<b>Speaker</b>	<b>Utterance</b>
customer	Hi, I'm looking to purchase a countertop blender. Can you help me out?
seller	Of course! Let's start with a couple of general features. Could you specify the price range you're hoping for the blender?
customer	Yes, I'd like it to be between 50 and 100.
seller	Great! Now, blenders also come with different types of blades. We usually recommend stainless steel as they're the most durable. Would you like a blender with a stainless steel blade?
customer	Yes, that works for me. I want my blender to have a stainless steel blade.
seller	I see. There are many brands that offer these features. However, I need to confirm, would you prefer a Kacsoo blender?
customer	No, not really. I'd like to consider other brands instead.
seller	No problem, we have plenty of other brands to choose from. Now, did you have a preference for the country of origin for the blenders? We have blenders made in Mexico, the USA, and China.
customer	Not really, the country of origin is optional for me.
seller	I understand. Let's move on to the power source. Would you like a blender that's powered by <b>AC, electricity, or corded electricity</b> ?
customer	That doesn't really matter to me. It's also optional.
seller	Fair enough. And what about the number of speed settings? We have blenders with up to 16 speed settings, as well as those with fewer options.
customer	That's also optional for me.
seller	Alright, I understand. How about the <b>department</b> and color? Some customers prefer their blenders match their kitchen's aesthetics.
customer	Both the department and the color are optional to me.
seller	Got it. Moving on, the voltage and style can also vary among blenders. Do you have a preferred <b>voltage</b> and style?
customer	Those aspects are optional for me.
seller	Understood. Now, what about the wattage and the components included with your blender?
customer	Those aspects aren't really important to me. They are also optional.
seller	Noted. And lastly, how about customer reviews and capacity? Some people prefer larger blenders and highly-rated products.
customer	Those are optional aspects for me as well.
seller	Thank you for all this information. Our blenders also come with different dishwasher safety levels. Is a dishwasher safe blender important to you?
customer	Yes, I'd like my blender to be dishwasher safe.
seller	Alright, considering all your preferences, I would like to recommend the "Oster 6812-001 Core 16-Speed Blender with Glass Jar, Black". It's priced at \$59.68, has a durable stainless steel blade, and also dishwasher safe. The other aspects that you mentioned as optional all meet your standards as well.
customer	That sounds perfect. Thank you for your help.

Table 21: Bad feature example by GPT-4 single-pass generation. The inappropriately selected product aspects in bold make the conversation awkward.



<b>Product Category</b>	cooling pads
<b>Wanted Features</b>	manufacturer: Kootek
<b>Optional Features</b>	None
<b>Unwanted Features</b>	color: white
<b>Speaker</b>	<b>Utterance</b>
customer	Hi, I'm looking for a cooling pad for my laptop.
seller	Great! Let's start by narrowing down the options. Can you tell me what kind of manufacturer are you interested in?
customer	Umm...I'm not sure. I've heard of Kootek, is that a good brand?
seller	Absolutely! Kootek is one of our best-selling brands for cooling pads. Let's focus on their products for now. How about the color? Do you have a preference?
customer	Ah, well...I don't want white. Is that an option?
seller	Yes, we have cooling pads in different colors. How about red or blue?
customer	Hmm...I think I like the sound of <b>blue</b> .
seller	Great! Now, let's talk about <b>size</b> . What size is your laptop?
customer	It's a <b>15-inch laptop</b> .
seller	Perfect! We have a Kootek cooling pad that fits 12-17 inch laptops. Would you like to have <b>adjustable mounts for the laptop stand height angle</b> ?
customer	Sure, that sounds useful.
seller	Great! Do you need any extra features like <b>USB ports or LED lights</b> ?
customer	LED lights would be cool, but I'm not sure about USB ports. Can you explain what they do?
seller	Of course! USB ports allow you to connect other devices to your laptop while it's on the cooling pad. For example, you could plug in a mouse or an external hard drive.
customer	Oh, that's handy. Yeah, I think I want USB ports then.
seller	Great! Lastly, would you like a <b>chill mat</b> that has a <b>quiet fan or a more powerful fan</b> ? Or would you prefer a balance between the two?
customer	Hmm...I don't know much about fans. Can you explain the difference?
seller	Sure thing! A quieter fan will provide a gentle breeze and might be better suited for everyday use, while a more powerful fan will provide a stronger airflow and might be better for gaming or heavy video editing.
customer	Oh, I see. I think I'll go with the <b>balance between the two</b> then.
seller	Excellent choice! Based on your preferences, I recommend our Kootek Laptop Cooling Pad 12"-17" Cooler Pad Chill Mat 5 Quiet Fans LED Lights and 2 USB 2.0 Ports Adjustable Mounts Laptop Stand Height Angle, Blue. It fits your budget and has all the features you're looking for. Would you like to take a look at it?
customer	Yes, that sounds great. Can I get a link to the product page?
seller	Of course! Here's the link. Let me know if you have any other questions or need further assistance.
customer	Thanks for your help! I think I'm good to go.
seller	You're welcome! Enjoy your shopping experience, and don't hesitate to reach out if you need anything else!

Table 22: Not instruction-following example by LLaMA2-70B-Chat single-pass generation. The only assigned required feature is the manufacturer of Kootek, and the only unwanted feature is white color.