Journal Pre-proof

Integrating Implicit Desire Fusion for Proactive Conversational Recommendation Systems with LLMs

Sitong Yan, Shengli Song, Chengen Lai, Chengen Lai, Guangneng Hu

PII: S1566-2535(26)00252-6
DOI: https://doi.org/10.1016/j.inffus.2026.104373
Reference: INFFUS 104373

To appear in: *Information Fusion*

Received date: 16 September 2025
Revised date: 3 February 2026
Accepted date: 10 April 2026

Please cite this article as: Sitong Yan, Shengli Song, Chengen Lai, Chengen Lai, Guangneng Hu, Integrating Implicit Desire Fusion for Proactive Conversational Recommendation Systems with LLMs, *Information Fusion* (2025), doi: https://doi.org/10.1016/j.inffus.2026.104373

Highlights

- Introduces desire-oriented CRS with precision, proactivity, and persuasion.
- Proposes DESIRE, a novel fusion-based framework for desire-oriented recommendation.
- Fuses implicit desire exploration with persuasive recommendations for proactive CRS.
- Demonstrates superior performance over existing LLM-based approaches across benchmarks

Integrating Implicit Desire Fusion for Proactive Conversational Recommendation Systems with LLMs

Sitong Yan[a], Shengli Song[a], Chengen Lai[a], Chengen Lai[a], Guangneng Hu[a]

[a]*Xidian University, Xi'an, 710126, China*

Abstract

Conversational Recommendation Systems (CRS) play a vital role in delivering personalized suggestions to improve user acceptance and engagement, and have been widely deployed in web services such as movie recommendations and e-commerce platforms. While prior research has largely emphasized recommendation accuracy and reactive intent elicitation through explicit questioning, the critical aspect of implicit desire remains underexplored. Moreover, existing LLM-based CRS frequently suffer from hallucinated responses and insufficiently persuasive explanations, further limiting their real-world applicability. To tackle these issues, we propose a new task, **desire-oriented CRS**, which highlights three key factors: *precision, proactivity, and persuasion.* Building on this concept, we introduce **DESIRE** (**DE**sire-fu**S**ed Conversat**I**on **RE**commendation),a fusion-based framework that integrates implicit multi-source desire signals with recommendation planning. DESIRE enables LLMs to infer dynamic user preferences and generate persuasive responses through effective desire reasoning fusion. Comprehensive evaluation on LastFM and Yelp datasets demonstrates that DESIRE significantly improves recommendation quality and dialogue persuasiveness, offering a robust foundation for fused desire recommendation interaction in next-generation AI agents.

Keywords: Conversational Recommendation System, Implicit Desire Fusion, Large Language Model(LLM), Human-computer Interaction

Email addresses: `styan@stu.xidian.edu.cn` (Sitong Yan), `shlsong@xidian.edu.cn` (Shengli Song), `laice@stu.xidian.edu.cn` (Chengen Lai), `jylee@stu.xidian.edu.cn` (Chengen Lai), `njuhgn@gmail.com` (Guangneng Hu)

1. Introduction

Conversational Recommendation System (CRS) engages users through natural language interactions to provide personalized recommendations, making it user-friendly for suggesting products or content. This approach is widely applied in web services, such as movie recommendations and the e-commerce platforms [1]. Traditional Web CRS primarily focuses on achieving high recommendation accuracy by reactively gathering user intents, categorized as **request-aware CRS**. However, users often express vague preferences or fail to specify the exact attributes of their desired items. In these cases, proactively deducing user desires becomes crucial for improving recommendation efficiency and enhancing user engagement, helping users navigate toward their multi-turn human-computer interactions. Addressing these challenges, **direct-ask CRS** systems have been developed to explore user preferences by explicitly asking attribute-related questions [2, 3, 4, 5]. While effective in preference elicitation, this approach often results in prolonged conversations, negatively impacting the overall user experience. Furthermore, prior research has largely overlooked the importance of offering persuasive, desire-oriented explanations alongside recommendations, which can significantly influence user decision-making and acceptance.

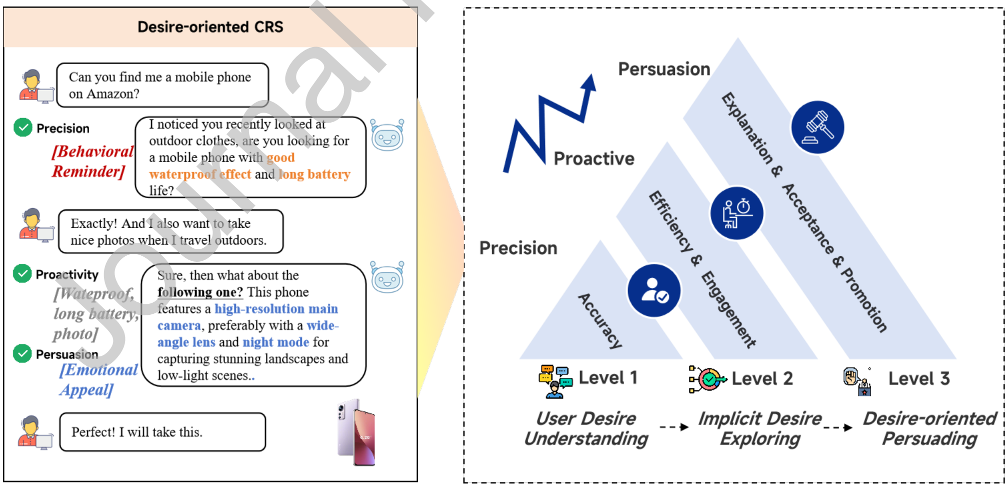

Figure 1: The successive levels and key factors of intelligent CRS.

To bridge this gap, we propose the concept of **desire-oriented CRS**, which operates with the following core goals: (1) **Precision:** the ability to

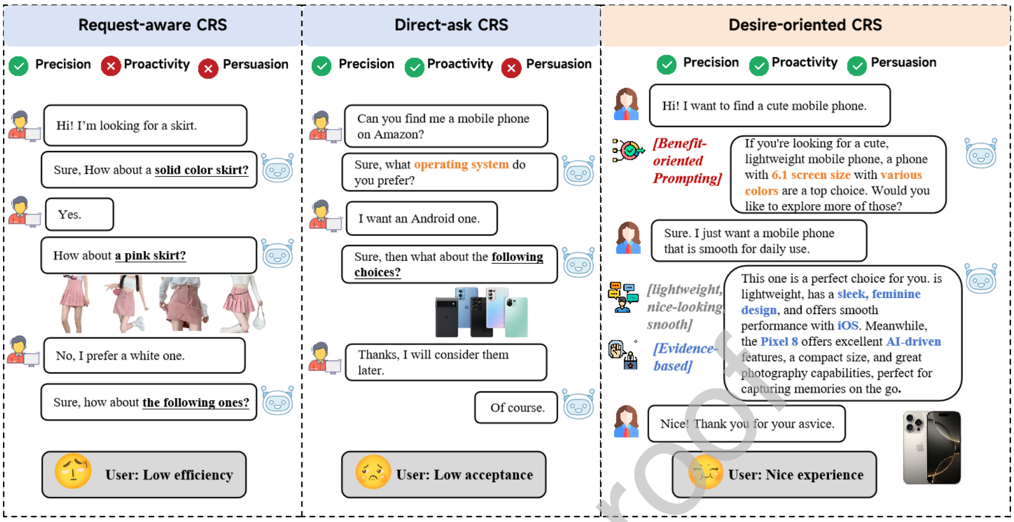

Figure 2: Illustrations of typical CRS (Request-aware CRS, Direct-ask CRS, and Desire-oriented CRS)

understand user desires and deliver highly accurate recommendations thoroughly; (2) **Proactivity:** the ability to proactively explore implicit user desires, engaging users in real-time and validating recommended options; and (3) **Persuasion:** the ability to identify key user desires and generate persuasive, desire-focused recommendation explanations. From a cognitive decision-making perspective, user preferences in conversational recommendation are often incomplete, unstable, and context-dependent. According to bounded rationality and information sufficiency theories, users may not be able to articulate their true intentions at the beginning of an interaction, and effective systems must actively elicit missing information to support satisfactory decisions. This observation provides a theoretical basis for desire-oriented CRS, where proactive questioning and clarification help reduce uncertainty and improve decision quality beyond purely reactive recommendation paradigms. These three factors represent successive levels of intelligence for Web CRS and collectively define a desire-oriented CRS, as depicted in Figure 1 and Figure 2.

The emergence of Large Language Models (LLMs) and their remarkable performance in tasks involving semantic understanding and response generation has opened new possibilities for developing CRS agents. Recent studies have shown that LLM-based agents can provide user-targeted recommendations [6], engage in proactive interactions [7], and generate persuasive explanations [8]. However, LLM-based Web CRS face significant challenges in desire-oriented applications, which can be categorized into (1) **Implicit Desire Exploration:** Proactively exploring user desires remains a major hurdle for LLM-powered CRS agents [6], given their inherently reactive nature. LLMs typically respond to user inputs rather than actively guiding the conversation. (2) **Persuasive Explanation Generation:** LLM-based CRS may sometimes produce misleading or deceptive responses, failing to deliver trustworthy and credible explanations [8].

Effective exploration must integrate heterogeneous signals, including dialogue context, long-term user preferences, external item knowledge, and temporal feedback, and reason about their complementary roles. This observation motivates a desire-centric formulation, where user desire is modeled as a multi-source information fusion problem and explicitly governs exploration, strategy selection, and response generation. To address the challenges above, we propose a novel framework called DESIRE(**DE**sire-fu**S**ed Conversat**I**on **RE**commendation), which fuses desire-oriented exploration with recommendation strategy planning to enhance LLM-based user-centric CRS. Inspired by the simple yet effective reasoning and action framework, ReAct [9], DESIRE prompts LLMs to infer users' implicit desires and interweave exploration responses with relevant recommendations. In this paper, we introduce a multi-source desire fusion formulation that grounds both exploration and persuasion decisions in a shared latent desire representation, offering a comprehensive assessment across three key dimensions: recommendation precision, dialogue proactivity, and response persuasion. In this paper, we focus on two web recommendation scenarios: music recommendation and e-commerce recommendation, which represent content-based recommendations and task-specific domains, providing a testbed to examine the reliance of LLMs on their internal knowledge. Experimental results demonstrate that DESIRE outperforms other PLM-based and LLM-based CRS baselines, effectively addressing the desire-oriented challenges in CRS across two datasets: LastFM and Yelp. Our contributions are as follows:

- We underscore the desire-oriented CRS with the three levels including precision, proactive and persuasion, emphasizing the challenges of web CRS with LLMs in task-specific recommendation scenarios.

- A novel desire-oriented fusion-based recommendation framework DE-

SIRE is proposed to enhance LLM's implicit desire exploration and persuasive recommendation, improving the precise recommendation with the guidance of external knowledge and key desire.

- Experimental results validate that our proposed DESIRE framework consistently and substantially outperforms existing LLM-based approaches with comprehensive user-centric metrics on LastFM and Yelp datasets, demonstrating the enhancement of LLMs on desire-oriented web CRS task.

2. Related Work

2.1. Conversational Recommendation System

To evaluate the performance of methods for CRS, some datasets that are based on movies and conversations annotated by crowdsource workers are released [10, 11], and a few other CRS datasets that involve e-commerce or Yelp review are adapted from non-conversational datasets [2, 5]. Early CRS methods only relied on basic methods of preference elicitation by the system and user-driven "critiquing" without utilizing natural language conversational [12]. Some methods relied heavily on pre-defined rules and had limited ability in modeling mixed-initiative interactions when conversational recommenders with natural language were first introduced. Subsequently, CRS utilizing model-based language generation and understanding emerged. However, these systems remained somewhat limited, primarily focusing on determining when to ask the user a question versus presenting recommendations, and deciding which questions to ask [13, 14]. Other studies have investigated learning more adaptable dialogue management modules in an end-to-end manner, typically by fine-tuning language models using dialogue data collected from crowdsourced workers. [15] integrates the recommender system and the dialog generation system to enhance the performance of the recommendation system. [10] proposed an end-to-end model to imitate the behavior of human players without considering the task goal itself. [11] model sub-components addressing different parts of the overall problem domain ranging from sentiment analysis and cold-start recommendation generation to detailed aspects of how natural language is used in this setting in the real world.

2.2. LLM-based Conversational Recommendation System

Recently, there has been a growing surge of research focused on using LLMs for training conduct conversational recommendation. A particularly relevant approach is zero-shot conversational recommenders with item-based [16, 17] or conversational inputs [18, 19], where an LLM generates an input based on a given description of a desired label, ensuring that the generated content matches the specified label. AI agents controlling pre-trained CRS or LMs for CRS tasks are also introduced as LLM-based methods for CRS [20, 21, 22]. Besides, user simulators evaluating interactive CRS systems has become a welcome issue [23, 24]. When ground truth labels do not support a fully differentiable loss function, some researchers demonstrate that it is still effective to fine-tune LLMs for language generation tasks using methods derived from reinforcement learning [25, 26]. Other works exploited reward signals inferred from the conversations for open-ended or task-based dialogue by using reinforcement learning to tune LLMs[27].

Recent advances in LLM-driven simulation enable scalable construction of conversational recommendation scenarios and support systematic evaluation of user-centric behaviors. While DESIRE itself does not rely on synthetic data generation as part of its method, this capability underpins the experimental setup and facilitates controlled analysis of desire-oriented recommendation strategies.

Besides, existing works have explored exploration, persuasion, or agent-based control in conversational recommendation, but they typically address these aspects in isolation or via heuristic coupling. DESIRE differs by explicitly modeling user desire as a multi-source fused signal that jointly governs exploration, persuasion, and generation.

3. Method

3.1. DESIRE: Overall Framework

Previous LLM-based CRS have focused on understanding explicit user intent and reactively providing item information as recommendations. However, they struggle with proactively exploring implicit desires and generating desire-aware persuasive explanations, resulting in low recommendation precision and efficiency. To address these limitations, we propose a dual-process framework named DESIRE (as shown as Figure 3), which is designed to address the intrinsic uncertainty of conversational recommendation by explicitly modeling desire exploration as a multi-source information fusion problem.

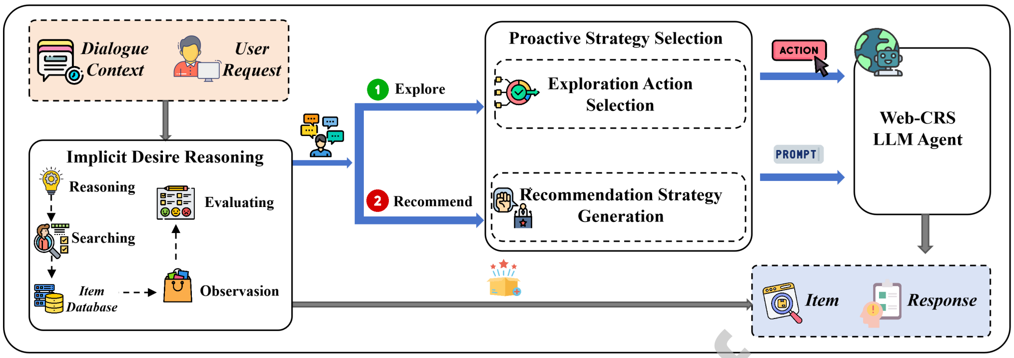

Figure 3: The overall framework of our method DESIRE, which unifies implicit desire reasoning and desire-aware proactive policy planning, both serving as tool-augmented modules to enhance Web-CRS LLM agents.

Taking dialogue history and the current user request as input, DESIRE first performs thorough implicit desire reasoning to intuitively predict user intent. In this process, the system needs to determine whether to proactively explore user implicit desire or perform credible recommendations with sufficient requests by searching from the item database. Based on this preliminary reasoning, the proactive policy planning module engages in deliberate action selection to further explore implicit desires or generate corresponding persuasion strategies for item promotion. Specifically, well-designed proactive exploration actions, such as direct inquiry, comparable questions, etc., along with persuasive strategies like evidence-based or emotional appeals, are employed for user-centric response generation. Empowered by these tool-augmented modules, the Web-CRS LLM agent not only identifies potential items but also delivers responses that are sensitive to implicit desires, resulting in more comprehensive and effective interactions.

3.2. Desire Signal Modeling and Fusion Formulation

In desire-oriented conversational recommendation, user intent is rarely conveyed through a single explicit signal. Instead, it emerges implicitly from multiple heterogeneous information sources that evolve over dialogue turns. We therefore formalize desire inference in DESIRE as a *multi-source information fusion* problem, where complementary desire signals are extracted, aligned, and fused to support proactive decision-making and grounded response generation.

Table 1: Summary of key notation.

Symbol	Description
t	Dialogue turn index
u_t	User utterance at turn t
r_t	System-generated response at turn t
C_t	Dialogue context up to turn t, defined as $C_t = \{u_1, r_1, \ldots, u_t\}$
S_i	The i-th raw desire signal source, $i \in \{1, 2, 3, 4\}$
H_u	Long-term preference history associated with user u
T_t	Temporal interaction records up to turn t
K^*	Retrieved item knowledge grounding recommendation and response generation
$f_i(\cdot)$	Source-specific encoding function that maps S_i to a latent representation
$\mathbf{d}_t^{(i)}$	Latent representation of the i-th desire signal at turn t
$\mathcal{D}_t$	Set of fused desire representations at turn t, i.e., $\mathcal{D}_t = \{\mathbf{d}_t^{(1)}, \ldots, \mathbf{d}_t^{(4)}\}$
a_t^{desire}	High-level desire decision at turn t, where $a_t^{\text{desire}} \in \{\text{explore}, \text{recommend}\}$
σ_t	Strategy selected at turn t conditioned on fused desire signals
e_t^a	Exploration action at turn t
e_t^p	Persuasive action at turn t
r_t^a	Response generated by exploration action
r_t^p	Response generated by persuasion action

3.2.1. Multi-source Desire Signals

At each dialogue turn t, DESIRE models user desire by integrating the following heterogeneous sources:

- **Dialogue Context (S_1).** The multi-turn conversational context $C_t = \{u_1, r_1, \ldots, u_t\}$, which provides implicit cues about user preferences, constraints, and evolving intent.

- **User Preference History / Persona (S_2).** Historical interaction patterns and latent preference profiles associated with the user (or simulator persona), capturing long-term or prior desires that may not be explicitly stated in the current turn.

- **Retrieved Item Knowledge (S_3).** Structured item attributes and

Table 2: Multi-source desire signals used in DESIRE. Each source is instantiated in a structured, data-level form consistent with the underlying datasets.

Data Source	Symbol	Example
Dialogue Context	S_1	C_t = [('user','need calm music'), ('sys','rec A'), ('user','no fast tempo'), ('sys','rec B')]
User Preference History	S_2	H_u = {'lastfm_hist':[('u42','a18',37),('u42','a07',21)], 'pref':{'genre':['jazz','ambient']}}
Retrieved Item Knowledge	S_3	K_{star} = [{'item_id':'i17','genre':'jazz','tempo':'slow'}
Temporal Interaction Signals	S_4	T_t = [{'turn':t,'item':'i88','fb':'reject'}]

factual descriptions K^* retrieved from the item database, serving as external grounding evidence for both desire inference and explanation generation.

- **Temporal Interaction Signals (S_4).** Turn-level behavioral dynamics, such as recent acceptance or rejection feedback and preference shifts, reflecting short-term desire evolution during the dialogue.

Each source provides a partial and noisy observation of the user's underlying desire. Individually, these signals are insufficient; jointly, they form a complementary basis for robust desire inference.

3.2.2. Desire Signal Representation

To enable fusion across heterogeneous modalities, each desire source S_i is transformed into a unified latent representation via a source-specific encoding function as Eq. 1.

$$\mathbf{d}_t^{(i)} = f_i(S_i), \quad i \in \{1, 2, 3, 4\}, \tag{1}$$

where $f_i(\cdot)$ denotes an LLM-guided structured reasoning process.

Specifically, dialogue context and temporal signals are summarized through instruction-constrained reasoning traces, user preference history is abstracted into preference distributions over item attributes, and retrieved item knowledge is represented as structured attribute–value pairs. These representations collectively form the desire signal set as Eq. 2.

$$\mathcal{D}_t = \{\mathbf{d}_t^{(1)}, \mathbf{d}_t^{(2)}, \mathbf{d}_t^{(3)}, \mathbf{d}_t^{(4)}\}. \tag{2}$$

3.2.3. Fusion Levels and Fusion Operator

Rather than performing a single monolithic aggregation, DESIRE conducts fusion at multiple complementary levels, consistent with established information fusion taxonomies.

Decision-level Fusion. Desire signals are fused to determine whether the system should further explore implicit user desires or proceed with item recommendation as Eq. 3.

$$a_t^{\text{desire}} = F_{\text{dec}}(\mathcal{D}_t), \tag{3}$$

where $a_t^{\text{desire}} \in \{\text{explore}, \text{recommend}\}$.

Strategy-level Fusion. Conditioned on the decision outcome, fused desire signals guide the selection of exploration actions or persuasive strategies as Eq. 4.

$$\sigma_t = F_{\text{str}}(\mathcal{D}_t \mid a_t^{\text{desire}}). \tag{4}$$

Generation-level Fusion. During response generation, fused desire representations are aligned with retrieved item knowledge to produce desire-aware and fact-grounded responses as Eq. 5.

$$r_t = F_{\text{gen}}(\mathcal{D}_t, K^*). \tag{5}$$

Across all levels, the fusion operator $F(\cdot)$ is implemented as a *gated, instruction-constrained LLM reasoning mechanism*, where different desire signals are explicitly exposed and adaptively weighted through structured prompts. This design allows the system to emphasize relevant sources while suppressing irrelevant or conflicting signals.

3.2.4. Fusion Objective and Benefits

The objective of desire fusion in DESIRE is threefold: (1) reducing uncertainty under incomplete or vague user input, (2) enhancing robustness by grounding generation in retrieved item knowledge to mitigate hallucinated attributes, and (3) enabling proactive exploration by modeling temporal and contextual desire evolution.

By formulating desire inference as a multi-source, multi-level fusion process, DESIRE provides a principled information fusion foundation for proactive and persuasive conversational recommendation, moving beyond purely reactive intent understanding.

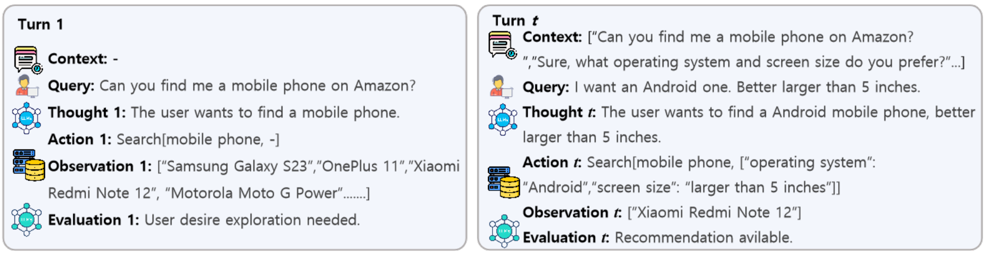

Figure 4: An example of Implicit Desire Reasoning Module with LLM-enhanced agent.

3.3. Implicit Desire Reasoning

Implicit Desire Reasoning module takes dialogue context C_t and user request q as initial input at each turn t, sequentially performing **Reasoning**, **Searching** and **Evaluating** procedure to determine whether the system should explore implicit user desire or recommend items. $\mathrm{LLM}_{\mathrm{DR}}$ (Desire Reasoning module) is responsible for synthesizing heterogeneous desire-related signals, including dialogue context, user preference history, retrieved item knowledge, and temporal interaction feedback. Its output is an explicit and structured desire representation that captures the user's latent intent and constraints, which serves as a shared intermediate state for downstream decision-making.

Inspired by the tool-augmented framework ReAct [9], which synergizes reasoning and acts as a uniform framework and generates observations by dynamically interacting with the environment, our method first generates reasoning thoughts to understand user intent with the prompt-based LLM agent and then searches in the item database. Specifically, desire reasoning LLM first generates user intent reasoning trace rt_t and search attribute sequence sa_t, which is contracted by item type and item attributes extracted by LLMs, as shown in Figure 3. In this module, we define action space with a single action *Search* and guide interaction with the item database. After searching in the item database, the item knowledge K^* is noted as current observation, composed of the relevant items knowledge. Finally, an LLM-enhanced agent serves as a two-classifier to evaluate whether the system should explore a more specific desire or provide accurate item recommendation, denoted as desire action da_t. The process of implicit desire reasoning is described in Eq. 6 and Eq. 7.

$$rt_t, sa_t = LLM_{DR}\left(C_t, q\right) \tag{6}$$

Journal Pre-proof

$$da_t = LLM_{EVA}\left(C_t, q, K^*\right) \tag{7}$$

To improve reproducibility, we explicitly clarify the inputs and outputs of the LLM-based components used in implicit desire reasoning. The desire reasoning module LLM_{DR} takes as input the dialogue context C_t and current user utterance q_t, and outputs (i) a structured reasoning trace summarizing inferred user intent, and (ii) a set of search attributes used to query the item database. The evaluation module LLM_{EVA} takes the dialogue context C_t, user request q_t, and retrieved item knowledge K^* as input, and outputs a binary decision indicating whether the system should further explore implicit desires or proceed with recommendation. Both modules operate under instruction constraints to ensure structured outputs and are detailed in Appendix A.

3.4. Proactive Strategy Selection

We emphasize that DESIRE does not perform explicit long-horizon planning in the classical reinforcement learning sense. Instead, it adopts a proactive strategy selection mechanism that reasons over candidate exploration and persuasion actions conditioned on the current dialogue state and fused desire signals. This design enables anticipatory decision-making while remaining compatible with LLM-based reasoning.

LLM_{EA} (Exploration and Action module) operates on top of the inferred desire representation and focuses on policy-level reasoning. It determines whether further desire exploration is needed and selects appropriate exploration actions or persuasive strategies accordingly. By decoupling desire inference from action selection, DESIRE achieves clearer modularity and more controllable decision behavior in conversational recommendation.

3.4.1. Exploration Action Selection

In this paper, several exploration strategies are designed to subtly guide users, making the chatbot a helpful assistant rather than just a sales tool, increasing the likelihood of uncovering and fulfilling latent needs.

- **Contextual Probing:** asks open-ended questions to gather more information about the user's preferences or current needs, even if the user has not provided specific details. For example, "Are you shopping for a special occasion, or just updating your wardrobe?"

- **Pattern Suggesting:** recommends products based on common purchasing patterns or trends that align with the user's vague requests. For example, "Many of our customers who browsed recently loved our fall collection. Would you like to check out some cozy sweaters or new boots?". Such action encourages exploration based on other user's behavior, making it easier for the user to narrow down their choices.
- **Preference Narrowing:** systematically narrows down the user's preferences by asking specific questions related to product categories, colors, or features. For example, "Are you looking for something in a specific color or style? Maybe a trendy piece or a timeless classic?"
- **Behavioral Reminder:** references the user's past browsing or purchase history to offer relevant suggestions, even when the user's current request is vague, which leverages past behavior to guide future purchases, enhancing personalization and relevance. For example, "I noticed you recently looked at summer dresses. Would you like to see more in that style, or are you thinking of something different this time?"
- **Benefit-oriented Prompting:** highlights the potential benefits or features of various products to spark interest and help the user identify needs they may not have realized, directing the user's attention to product features or benefits that align with common desires, nudging them towards a potential purchase. For example, " If you're looking for comfort, our memory foam shoes are a top choice this season. Would you like to explore more of those?"

Based on the above new-designed desire exploration actions, we employ an LLM agent as an exploration action selector, which takes dialogue context C_t and user request q as input and generates specific exploration action ea_t and exploration prompts ep_t, which is respresented as Eq. 8.

$$ea_t, ep_t = LLM_{EAS}\left(C_t, q, rt_t, K^*\right) \tag{8}$$

3.4.2. Persuasive Strategy Generation

In our work, we designed the persuasion strategy types following the theory of the Elaboration Likelihood Model [28, 8].This framework presents

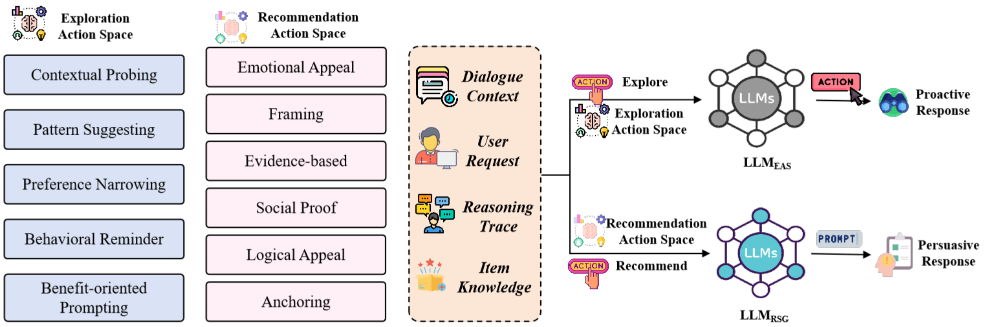

Figure 5: Proactive Strategy Selection Module.

a comprehensive approach to influence, analyzing distinct cognitive pathways individuals employ when responding to external triggers and how these mechanisms subsequently modify perceptions and drive actions. The model explores varied methods of information assimilation that ultimately reshape viewpoints and decision-making patterns.

Our persuasion strategy design is further grounded in established cognitive psychology frameworks. In particular, we draw inspiration from the Elaboration Likelihood Model (ELM), which characterizes persuasion as occurring through either a central route, based on careful evaluation of arguments, or a peripheral route, relying on heuristic cues. Strategies such as evidence-based reasoning and logical appeal correspond to the central route, while framing, emotional appeal, and social proof align with peripheral persuasion mechanisms. DESIRE does not enforce a fixed persuasion pathway; instead, it dynamically selects strategies conditioned on inferred user desire and interaction context, allowing the system to adapt its persuasive behavior in a cognitively plausible manner.

- **Logical Appeal** involves transparently presenting the system's reasoning process to influence users. For example, explaining how a product's features align with the user's specific needs helps the user understand why the product is recommended, fostering trust and encouraging them to accept the suggestion more readily.

- **Emotional Appeal** focuses on eliciting specific emotions or sharing impactful stories to build trust and deepen user connections. For instance, describing how a particular product (e.g., a kitchen appliance) could make a user's daily life more convenient or enjoyable taps into their emotions, making the recommendation feel more personal and meaningful, thus influencing their purchasing decision.

- **Framing** emphasizes the positive aspects of a decision in a credible way. For example, highlighting the benefits of purchasing a product (e.g., how buying a high-quality office chair can improve posture and comfort) enhances the perceived value of the decision, making the recommendation more attractive and appealing.

- **Evidence-based** Persuasion uses empirical data or objective facts to support a recommendation. For instance, showing verified customer ratings or industry certifications for a product (e.g., a high-performing smartphone or eco-friendly household items) reduces subjective bias and adds credibility, making the suggestion more persuasive and trustworthy.

- **Social Proof** relies on leveraging the behavior or endorsements of others to strengthen a recommendation. For example, displaying a product's high sales figures or positive customer reviews (e.g., showing how a popular pair of running shoes is favored by thousands of satisfied customers) capitalizes on the psychological tendency for people to follow the crowd, thus increasing the persuasive impact and credibility of the recommendation.

- **Anchoring** refers to presenting a credible, initial piece of information as a reference point to guide users in subsequent decisions. For example, first showcasing a product's best-selling status or award-winning design (e.g., presenting a best-in-class coffee machine) and then describing its features and benefits helps establish an anchor. This cognitive pattern demonstrates how primary disclosures disproportionately influence later reasoning processes and behavioral outcomes. Reliance on opening statements creates decision-making pathways that favor continuity over reevaluation, making the persuasion more effective.

As the conversation unfolds, we dynamically choose the most appropriate strategy to guide explanation generation at each turn, allowing for seamless adaptation to the evolving dialogue context. For any recommended item, the item information collection module retrieves detailed information from a credible source, such as an item database. With the conversation history C_t, user request q, and the retrieved item informationK^* as inputs, the strategy selector—driven by the LLM identifies the optimal strategy s from a set of desire-aware persuasive strategies and generate persuasive strategyra_t and explanations promptsrp_t, as described in Eq. 9.

$$ra_t, rp_t = LLM_{PSG}\left(C_t, q, K^*, rt_t\right) \tag{9}$$

At each dialogue turn t, DESIRE selects a strategy to balance effectiveness and efficiency. Formally, the objective is to maximize the probability of successful recommendation while minimizing unnecessary interaction turns, subject to grounding and consistency constraints. Rather than optimizing over a fixed long horizon, DESIRE adopts a receding-horizon decision process, where the dialogue state is updated after each turn based on the fused desire representation.

3.5. LLM-based Web CRS Agent

As mentioned above, implicit desire reasoning and proactive policy planning agents are essential for supporting desire-oriented Web CRS tasks, while LLMs act as tool-augmented conversational agents, utilizing these tools to fulfill core CRS objectives. Upon receiving a new dialogue history C_t and user request q, the LLM-based agent uses these tools to determine the current exploration action a_t or recommendation explanation prompt p_t along with relevant item knowledge K^*. This information then guides the generation of a system response s^{t+1}_{system} and an item recommendation i through a prompting mechanism, as described in Eq. 10.

$$i, s_{t+1}^{\text{system}} = \begin{cases} LLM_{EAS}\left(C_t, q, a_t\right) & \text{if } da_t = \text{explore} \\ LLM_{PSG}\left(C_t, q, K^*, p_t\right) & \text{if } da_t = \text{recommend} \end{cases} \tag{10}$$

To mitigate hallucinated recommendations and unsupported attribute generation, DESIRE enforces grounding constraints during response generation. Specifically, all factual item attributes referenced in system responses must be supported by retrieved item knowledge K^*. The LLM is instructed to abstain from fabricating attributes that are not present in K^*, and to fall back to clarification or exploration actions when retrieved evidence is insufficient. These grounding constraints serve as a lightweight guardrail against hallucination and are consistently applied across all experimental settings, as detailed in Appendix A.

4. Experiment

To validate the effectiveness of DESIRE on desire-oriented CRS, our initial assessment focuses on the complete framework to analyze the comprehensive impact introduced by the tool-enhanced LLM-driven CRS. Subsequently, we conduct thorough component-wise analyses to explore the contribution of each distinct element. Given the architectural variations between models, PLM-based CRSs lack user adoption metrics, rendering Convincing Acceptance unquantifiable. The experimental investigation is structured around these core research inquiries (RQ):

- **RQ1:** How does DESIRE perform when benchmarked against current conversational recommendation systems, particularly in terms of desire reasoning capabilities?
- **RQ2:** How does the proactive user desire exploration and corresponding strategies contribute to Web CRS?
- **RQ3:** How much user engagement and acceptance have been improved with the incorporation of persuasive response generation?

4.1. Experiment Setup

4.1.1. User Simulator & Dataset

Manual evaluation of Conversational Recommendation Systems (CRS) often requires continuous user interaction, making it labor-intensive and typically feasible only in industrial lab settings [29]. Recent research, however, introduces interactive evaluation using LLM-based user simulations [23, 24], which has been experimentally shown to be a reliable alternative to human evaluators. In this paper, we adopt the simulation framework presented in [23, 24], specifically designed for two Web CRS datasets: LastFM[1] and Yelp[2]. The LastFM collection serves as a benchmark for assessing musical performer suggestions, whereas the Yelp data repository is utilized for evaluating item recommendations within e-commerce scenarios. LastFM consists of 1,801 users and 7,432 items, covering 33 attributes within 76,693 interactions. Yelps has a larger scale with 27,675 users, 70,311 items, 590

[1]https://grouplens.org/datasets/hetrec-2011/
[2]https://www.yelp.com/dataset/

item attributes and 1,368,606 interactions. Implementation details, including prompt templates and key hyperparameters, are provided in Appendix A to facilitate reproducibility.

4.1.2. Baselines

We compare DESIRE with SOTA PLM-based methods, i.e., EAR [2] and UNICORN [4]. To evaluate the effectiveness of our proposed implicit desire exploration module, we also compare DESIRE with other LLM-based CRSs, including InterCRS [23], ChatCRS [6] and MACRS [30]. GPT-3.5-Turbo is employed to construct LLM-enhance CRS in DESIRE and the baselines. Details of the baselines are shown as follows:

- **EAR [2]:** A three-stage methodology is formulated to optimize the synergistic coordination between dialogue and recommendation modules, adapting the reinforcement learning paradigm established in conversational recommendation systems.
- **UNICORN [4]:** It conceptualizes CRS decision-making through a consolidated policy optimization paradigm, implementing an adaptive graph neural mechanism that dynamically weights conversational states to determine action selection per dialogue turn.
- **InterCRS [23]:**It incorporates ChatGPT with zero-shot prompting and external recommendation models to explore the ChatGPT-based CRS. In this paper, we choose the InterCRS integrated with the recommendation model as the baseline.
- **ChatCRS [6]:** It decomposes the overall CRS problem into subcomponents handled by specialized agents for knowledge retrieval and goal planning, all managed by a core LLM-based conversational agent.
- **MACRS [30]:** It integrates two core mechanisms - collaborative agent behavior planning and feedback-anchored reflection processes - whose coordinated operation achieves significant user experience enhancement over baseline LLM direct deployment paradigms.

Unlike baselines that rely on fixed heuristics or reactive generation, DESIRE explicitly evaluates multiple candidate strategies at each turn and selects actions conditioned on multi-source desire signals. This distinction allows DESIRE to proactively explore or persuade based on anticipated user intent rather than solely reacting to surface utterances.

4.1.3. Evaluation Metrics

To realize user-aware CRS evaluation, recent studies have focused on incorporating LLM as a simulator and evaluator, along with system-centric and user-centric factors [24, 23, 31]. Following the previous evaluation frameworks, we focus on the following metrics that highly correlated with our proposed method, which highlights the necessity of incorporating user implicit desire exploration and generating desire-aware persuasion. Three evaluation dimensions can be categorized as: Precision, Proactivity and Persuasion.

- **Precision:** Web-based conversational recommendation systems must achieve dual operational imperatives: delivering recommendation accuracy while minimizing dialogue turns, both being critical determinants of user satisfaction. To quantitatively assess recommendation precision, we establish an evaluation framework employing success rate metrics (**SR@t**) [32].

- **Proactivity:** Proactivity demonstrates the ability of LLM-based Web CRS to efficiently explore user desire by incorporating various exploration strategies, which further enhances the conversation efficiency and the user-system interaction. Considering the deficiency of fine-grained annotated data for action selection accuracy evaluation, we examine the efficiency with the metrics of average turn (**AT**), which measures the typical number of interactions across all conversational exchanges, where the lower AT means an overall higher efficiency [33].

- **Relevance:** To evaluate the helpfulness of system proactive response, we utilize an LLM-based evaluator to examine the Web CRS from a new metric namely Relevance:1) Clarity of Expression: System responses must ensure linguistically accessible and structurally unambiguous formulations. 2) Information Sufficiency : Response generation requires optimal information density balancing, avoiding both data paucity and cognitive overload. 3) Preference Alignment : Dialogue acts should actively construct dynamic user preference profiles to facilitate recommendation relevance..

  Following the instruction of [8], we employ the metrics of **Persuasivenss** and **Convincing Acceptance** to evaluate the persuasion of recommendation explanation.

- **Persuasion:** Persuasiveness is achieved by instructing the LLM-based evaluator to score its watching intention, ranging from 1 to 5, where the evaluator rates its initial intention i_{pre} based solely on the item's name and is required to rate the intention i_{post} after reading the CRS explanation. Finally, the evaluator rates the 'true' intention i_{true} after seeing the full information about the item. And the Persuasiveness is calculated as follows.

$$\text{Persuasiveness } = 1 - \frac{i_{\text{true}} - i_{\text{post}}}{i_{\text{true}} - i_{\text{pre}}}. \tag{11}$$

  A higher Persuasiveness score means a stronger ability in arousing user's watching intention towards recommended items.

- **Convincing Acceptance:**

  Convincing Acceptance metric aims to assess dialogue-level credibility. It measures how often the CRS successfully convinces the simulator to accept a recommendation while maintaining high credibility. A higher Convincing Acceptance indicates a lower likelihood of users being misled by deceptive explanations.

$$\text{Convincing Acceptance}(d) = \mathbb{I}\big(i_d^{\text{rec}} = i_d^{\text{true}} \wedge \text{Cred}(d) \geq \tau\big), \tag{12}$$

  where $\text{Cred}(d)$ denotes a credibility score estimated by an automatic evaluator, τ is a fixed threshold, and $\mathbb{I}(\cdot)$ is the indicator function. The Convincing Acceptance score is computed by averaging Convincing Acceptance(d) over all dialogue sessions.

  If multiple recommendations are made within a dialogue, only the final accepted item is considered. Dialogues without explicit acceptance are treated as unsuccessful cases for both Persuasiveness and Convincing Acceptance. All metrics are computed at the dialogue level and then averaged across sessions to ensure consistency.

Following recent work in conversational recommendation, we adopt an LLM-based user simulator to enable scalable and reproducible multi-turn interactions. We emphasize that the simulator is not intended to faithfully model real user behavior, but rather to provide a consistent and controllable interaction environment. All compared methods interact with the same simulator under identical prompts, turn limits, and stopping criteria, ensuring fair relative comparison across methods.

Table 3: Our experiment results on the LastFM and Yelp datasets, demonstrate that DESIRE can effectively increase the user-centric recommendation success rate, especially attains remarkable performance utilizing GPT-3.5-Turbo and GPT-4-Turbo variants. The bold ones suggest the best performance and the underline notations represents the second best performance in different baselines, which aims at better comparsion of various methods. We conduct 5 random experiments and report the mean and standard deviation in Appendix C.

Model		LastFM					Yelp				
		SR@15	AT	Relevance	Persuasiveness	Convincing Acceptance	SR@15	AT	Relevance	Persuasiveness	Convincing Acceptance
PLM-based Model	EAR	0.429	12.88	3.51	22.97	-	0.967	5.74	3.92	11.67	-
	UNICORN	0.535	11.82	3.87	21.45	-	**0.985**	5.33	4.07	10.76	-
LLM-based Model	InterCRS	0.431	13.45	4.12	68.51	56.25	0.561	7.26	4.51	53.76	40.57
	ChatCRS	0.479	12.91	4.21	60.31	52.70	0.756	6.61	4.69	65.19	48.21
	MACRS	0.544	14.47	4.62	65.49	59.87	0.791	6.25	4.74	73.87	61.46
	DESIRE(LLaMA-7B)	0.562	13.22	4.68	68.25	62.89	0.835	5.95	4.87	72.76	66.59
	DESIRE (GPT-3.5-Turbo)	0.583	12.76	4.70	70.28	69.41	0.937	5.78	4.89	79.16	75.96
	DESIRE (GPT-4-Turbo)	**0.601**	**11.78**	**4.75**	**72.56**	**70.98**	0.974	**5.21**	**4.91**	**82.27**	**79.03**

4.1.4. Latency and Resource Analysis

DESIRE follows a modular LLM-based architecture, where desire reasoning, strategy selection, and response generation are implemented as sequential prompting steps without introducing additional model parameters or training overhead. Therefore, the computational cost of DESIRE is dominated by the underlying backbone language model.

To ensure a fair comparison, DESIRE and all LLM-based baselines are evaluated using the same backbone models (e.g., GPT-3.5 or LLaMA), identical decoding configurations (temperature, top-p, and maximum token length), and the same stopping criteria. Under these controlled settings, DESIRE exhibits comparable inference latency and GPU usage to existing LLM-based conversational recommendation baselines. In contrast, PLM-based methods are computationally lighter but lack the capability to support multi-step desire reasoning and persuasive dialogue generation.

4.2. Main Results

The experimental results highlight significant differences between the PLM-based and LLM-based models, as well as insights into the performance of different LLM-based models, particularly the DESIRE models. The experimental results in Table 1 are discussed as follows:

4.2.1. PLM-based vs. LLM-based Models.

The LLM-based models consistently outperform PLM-based models across all metrics on both datasets. For instance, on the LastFM dataset, the best PLM-based model, UNICORN, achieves an SR@15 of 0.535, whereas the top LLM-based model, DESIRE (GPT-4-Turbo), reaches an SR@15 of 0.601, showcasing a significant improvement. A similar pattern is observed on the Yelp dataset, where PLM-based models score around 0.985 (UNICORN), but LLM-based models like GPT-4-Turbo achieve near-perfect scores (0.974). Furthermore, in terms of relevance, persuasiveness, and user acceptance metrics, LLM-based models outperform PLM-based models, particularly in user-centric metrics such as convincing acceptance, where LLM-based models perform remarkably well. This demonstrates that LLM-based models are better at understanding and persuading users, leading to higher acceptance rates.

We observe that PLM-based methods achieve very high SR@15 on Yelp, while LLM-based conversational recommenders exhibit a wider and sometimes lower performance range. This discrepancy is expected due to fundamental paradigm differences. PLM-based models are optimized for item ranking and retrieval, which directly aligns with the SR@15 objective under fixed candidate sets. In contrast, LLM-based methods must jointly handle multi-turn dialogue understanding, desire inference, strategy selection, and response generation, where recommendation success is influenced by additional sources of uncertainty. Moreover, while Yelp provides strong structured signals that benefit retrieval-based matching, conversational settings introduce evolving or under-specified user intent, which poses greater challenges for LLM-based CRS.

4.2.2. Comparision between LLM-based Models.

Among the LLM-based models, there are noticeable differences. DESIRE (GPT-4-Turbo) consistently outperforms all other models across on LastFM and Yelp. On LastFM, it achieves the highest values in all metrics, including SR@15 (0.601), relevance (4.75), and persuasiveness (72.56), indicating superior performance in recommendation tasks. Additionally, DESIRE (GPT-4-Turbo) demonstrates outstanding performance on the Yelp dataset, achieving the best scores in every metric, such as persuasiveness (82.27) and convincing acceptance (79.03). In contrast, models like InterCRS and ChatCRS, while performing better than PLM-based models, lag behind in comparison to more advanced models like DESIRE. This suggests that more advanced architectures and training strategies, as used in GPT-4-Turbo, yield signifi-

Table 4: Ablation study results on the LastFM and Yelp datasets with GPT-3.5-Turbo.

Model	SR@15	AT	Relevance	Persuasiveness	Convincing Acceptance
LastFM					
DESIRE(w/o desire reasoning)	0.327	12.85	4.68	69.51	52.77
DESIRE(w/o exploration action)	0.491	14.78	4.32	69.90	63.79
DESIRE(w/o persuasive strategy)	0.578	12.95	4.54	65.78	60.21
DESIRE	**0.583**	**12.76**	**4.70**	**70.28**	**69.41**
Yelp					
DESIRE(w/o desire reasoning)	0.712	7.98	4.79	76.55	52.98
DESIRE(w/o exploration action)	0.869	8.12	4.65	77.21	68.56
DESIRE(w/o persuasive strategy)	0.913	6.65	4.82	75.13	67.81
DESIRE	**0.937**	**5.78**	**4.89**	**79.16**	**75.96**

cantly better results in recommendation success and user acceptance.

4.2.3. Comparison between different LLM base of DESIRE.

The DESIRE model is evaluated in three versions: LLaMA-7B, GPT-3.5-Turbo alongside GPT-4-Turbo. Notably, GPT-4-Turbo demonstrates superior performance consistently, outperforming alternatives in every dataset and evaluation criterion.particularly on the Yelp dataset where it achieves an SR@15 of 0.974, relevance of 4.91, and convincing acceptance of 79.03. While GPT-3.5-Turbo also performs well, it slightly lags behind GPT-4-Turbo, particularly in persuasiveness (79.16 vs. 82.27) and convincing acceptance (75.96 vs. 79.03). The LLaMA-7B version, while still competitive, falls behind the GPT-based versions, especially in user engagement metrics like convincing acceptance, where it scores 66.59 on Yelp, compared to 70.98 for GPT-4-Turbo. This indicates that the latest iteration of GPT models offers significant improvements in both recommendation accuracy and user engagement over earlier versions or alternative architectures like LLaMA.

4.3. In-depth Analysis

In this section, we present ablation analyses conducted across two datasets using GPT-3.5-Turbo to evaluate the specified research inquiries, as detailed in Table 2.The experimental outcomes conclusively validate the beneficial effects of motivational analysis, anticipatory investigation, and convincing re-

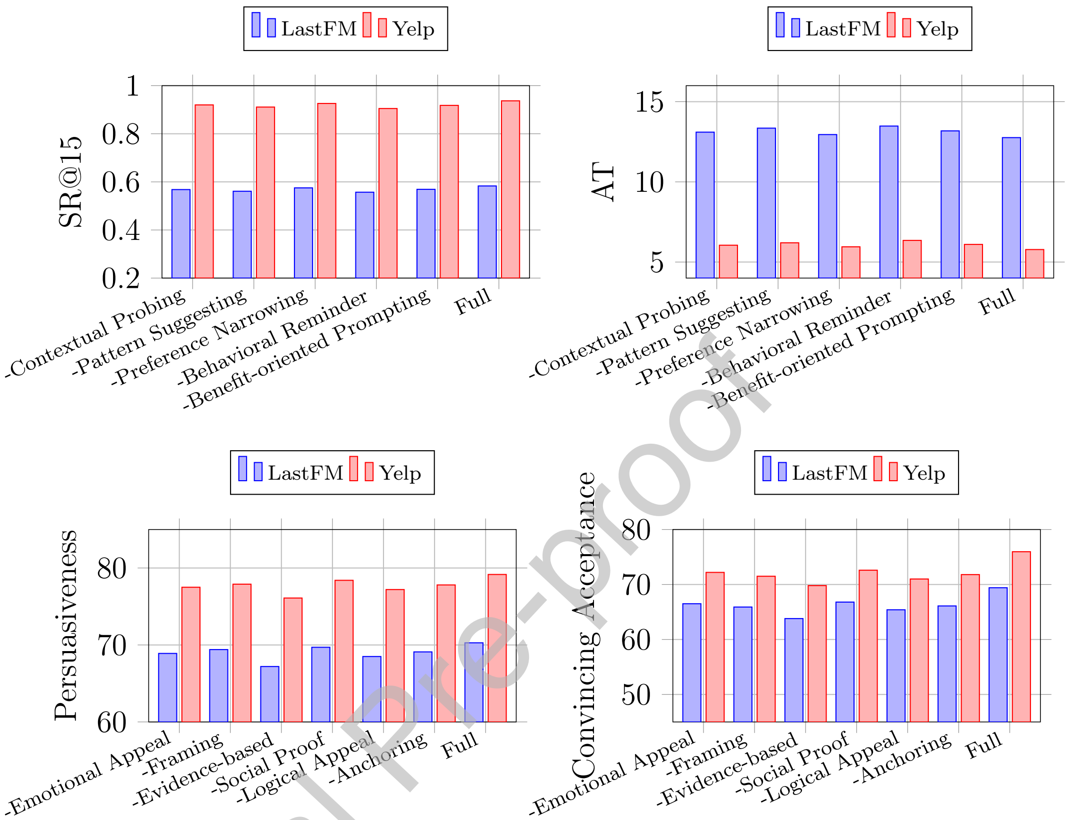

Figure 6: Fine-grained ablation visualization.

ply formulation within dialogue-driven recommendation frameworks. When combining all three mechanisms, DESIRE exhibits dominant results across every significant performance indicator, indicating these approaches are fundamentally important for enhancing both suggestion precision and participant interaction in online conversational recommendation platforms. The details of the analysis is illustrated as follows:

4.3.1. RQ1: Precise Recommendation

Table 2 shows that the full DESIRE model, which incorporates desire reasoning, outperforms the variant without desire reasoning across both datasets. For instance, on the LastFM dataset, the full DESIRE model achieves an SR@15 of 0.583, significantly higher than the 0.327 achieved by the model without desire reasoning. Similarly, on the Yelp dataset, DE-

SIRE reaches an SR@15 of 0.937, compared to 0.712 for the variant without desire reasoning. This demonstrates that the inclusion of desire reasoning improves recommendation precision, helping DESIRE provide more targeted and relevant suggestions. In terms of persuasiveness and user acceptance, the full model consistently outperforms the desire-reasoning-absent variant, indicating that understanding user desires enhances the effectiveness of recommendations in terms of both relevance and convincingness.

4.3.2. RQ2: Proactive Recommendation

The proactive exploration of user desires is shown to contribute significantly to recommendation success. On both datasets, the model without exploration actions performs noticeably worse than the full DESIRE model. On LastFM, the absence of exploration results in an SR@15 of 0.491, compared to the 0.583 of the full model. On Yelp, the difference is smaller but still significant (0.869 vs. 0.937). These results indicate that the proactive desire exploration improves the system's ability to identify the user's needs more accurately, leading to better recommendation outcomes. Furthermore, the improvement in convincing acceptance from 63.79 (w/o exploration) to 69.41 (full model) on LastFM, and from 68.56 to 75.96 on Yelp, suggests that exploring users' preferences contributes to a stronger user connection with the recommendations, enhancing both user engagement and acceptance.

4.3.3. RQ3: Persuasive Recommendation

Incorporating persuasive strategies into the response generation process significantly improves user engagement and acceptance. Without the persuasive strategy, the model's performance in terms of convincing acceptance drops to 60.21 on LastFM and 67.81 on Yelp, compared to 69.41 and 75.96 for the full model, respectively. Additionally, persuasiveness itself is enhanced with the full model, as shown by the higher scores on both datasets (70.28 on LastFM and 79.16 on Yelp). These results indicate that persuasive response generation is critical for increasing user trust and encouraging them to accept the recommendations. By making the responses more compelling and aligned with user desires, The implementation of compelling techniques substantially enhances not only engagement standards but also operational efficiency. These approaches demonstrate measurable improvements in user-system dynamics while optimizing performance metrics across multiple evaluation parameters.

Table 5: Human evaluation results on the sampled dataset from LastFM and Yelp based on pairwise win rates. For each dialogue, annotators compare DESIRE with a baseline and judge which response is better along different evaluation dimensions. Win / Loss / Tie indicate the percentage of dialogues where DESIRE is preferred, not preferred, or judged as equal. ‡/† denote p-value $< 0.1/0.05$ (statistical significance test)

DESIRE vs.	**InterCRS**			**ChatCRS**			**MACRS**		
	Win	**Loss**	**Tie**	**Win**	**Loss**	**Tie**	**Win**	**Loss**	**Tie**
Persuasiveness	**49.2**‡	36.7	14.1	**61.3**	24.5	14.4	37.8	**41.9**	20.3
Trustworthiness	**62.1**‡	20.6	17.4	**58.4**‡	19.8	21.7	**47.4**†	32.8	19.8
Groundedness	**51.9**†	31.2	16.9	**42.2**	40.6	17.2	**37.4**	32.5	30.1
Overall	**64.1**‡	23.8	12.1	**56.2**†	31.9	11.9	**49.5**‡	30.6	19.9

4.4. Human Evaluation

To complement automatic evaluation, we conduct a human study based on pairwise comparison. We randomly sample 100 dialogue sessions from each dataset. For each session, annotators are presented with the same dialogue context and two anonymized system responses, one generated by DESIRE and the other by a baseline model. Annotators are asked to judge which response is better overall in terms of persuasiveness and trustworthiness, or select a tie if they are indistinguishable. Each comparison is evaluated by three annotators, and the final decision is determined by majority vote. Table 5 reports the resulting win, tie, and loss rates.

Overall, DESIRE consistently outperforms all baselines across evaluation dimensions, demonstrating clear advantages in human-perceived response quality. In terms of *persuasiveness*, DESIRE achieves significantly higher win rates against InterCRS and ChatCRS, indicating that explicit multi-source desire modeling leads to more convincing and preference-aligned responses. DESIRE also shows strong performance on *trustworthiness*, with win rates exceeding 47% across baselines, highlighting the benefit of grounding responses in retrieved item knowledge to reduce unsupported claims. For *groundedness*, DESIRE maintains consistent improvements, particularly over InterCRS, while achieving comparable or higher win rates against ChatCRS and MACRS, suggesting better factual alignment without sacrificing conversational naturalness.

Across both datasets and all evaluation dimensions, DESIRE consistently achieves higher win rates than strong baselines, with particularly pronounced advantages in persuasiveness and groundedness. These results validate the effectiveness of DESIRE beyond automatic metrics and alleviate concerns regarding LLM-based evaluation bias.

5. Conclusion

In this paper, we have explored the critical role of desire-oriented Conversational Recommendation Systems in improving user engagement and acceptance in personalized recommendations. Based on the new task, we introduced the novel concept of desire-oriented CRS, centered around three key factors: precision, proactivity, and persuasion. Based on this concept, we developed DESIRE, a unified framework that fuses desire exploration with recommendation strategy planning. This framework enables LLMs to dynamically infer user preferences and generate both relevant recommendations and persuasive explanations through interactive dialogues. Additionally, we employed a comprehensive user-centric evaluation system to assess CRS performance across all three dimensions. Our extensive experiments on two web CRS datasets, LastFM and Yelp, demonstrate the effectiveness and superiority of DESIRE over existing PLM and LLM-based baselines.

Desire-oriented conversational recommendation has broader societal implications. By explicitly modeling user desires, systems like DESIRE can reduce irrelevant recommendations, improve transparency through grounded explanations, and support more informed user decisions, thereby enhancing user agency. At the same time, the persuasive nature of such systems raises ethical concerns regarding potential manipulation or bias reinforcement.

6. Limitations and Future Directions

Cold-start and Preference Sparsity. DESIRE may be less effective in cold-start scenarios where little or no prior user interaction history is available. Although dialogue-based exploration and external knowledge grounding can partially compensate by eliciting preferences through interaction, additional turns may be required to reach confident recommendations. Incorporating cross-session memory or population-level priors is a promising direction for future work.

Multilingual, Cross-cultural, and Cross-domain Generalization. Generalizing DESIRE across languages, cultures, and domains remains challenging. User desire expression and persuasive effectiveness can vary significantly across cultural and linguistic contexts, potentially affecting both interpretation and response generation. While DESIRE is model-agnostic and compatible with multilingual LLMs, adapting exploration and persuasion strategies to different norms, as well as extending to domains such as healthcare or travel with domain-specific constraints, warrants further study.

Ethical Considerations and Deployment Risks. Proactive desire exploration and persuasive recommendation raise ethical and privacy concerns in real-world deployment. Overly aggressive or opaque persuasion may influence user choices in unintended ways. DESIRE mitigates these risks through explicit desire representations and external knowledge grounding, which improve transparency and reduce unsupported persuasion. Nevertheless, responsible deployment practices, including user consent, data minimization, and human oversight, remain essential.

Model Bias and Fairness Considerations. The backbone language models used in DESIRE (e.g., GPT-3.5/4 and LLaMA) may inherit biases from pretraining data, potentially affecting recommendation behavior. As current benchmarks lack reliable demographic annotations, our evaluation focuses on effectiveness and persuasion rather than fairness. Nevertheless, DESIRE is model-agnostic and compatible with fairness-aware models or debiasing strategies, and systematic bias analysis is left as future work.

Declaration of interests

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Guangneng Hu reports financial support was provided by National Natural Science Foundation of China. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] C. Gao, W. Lei, X. He, M. de Rijke, T.-S. Chua, Advances and challenges in conversational recommender systems: A survey, AI open 2 (2021) 100–126.

[2] W. Lei, X. He, Y. Miao, Q. Wu, R. Hong, M.-Y. Kan, T.-S. Chua, Estimation-action-reflection: Towards deep interaction between conversational and recommender systems, in: Proceedings of the 13th International Conference on Web Search and Data Mining, 2020, pp. 304–312.

[3] W. Lei, G. Zhang, X. He, Y. Miao, X. Wang, L. Chen, T.-S. Chua, Interactive path reasoning on graph for conversational recommendation, in: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining, 2020, pp. 2073–2083.

[4] Y. Deng, Y. Li, F. Sun, B. Ding, W. Lam, 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, 2021, pp. 1431–1441.

[5] Y. Zhang, X. Chen, Q. Ai, L. Yang, W. B. Croft, Towards conversational search and recommendation: System ask, user respond, in: Proceedings of the 27th acm international conference on information and knowledge management, 2018, pp. 177–186.

[6] C. Li, Y. Deng, H. Hu, M.-Y. Kan, H. Li, Chatcrs: Incorporating external knowledge and goal guidance for llm-based conversational recommender systems (2025) 295–312.

[7] Y. Deng, L. Liao, L. Chen, H. Wang, W. Lei, T.-S. Chua, Prompting and evaluating large language models for proactive dialogues: Clarification, target-guided, and non-collaboration, in: Findings of the Association for Computational Linguistics: EMNLP 2023, 2023, pp. 10602–10621.

[8] P. Qin, C. Huang, Y. Deng, W. Lei, T.-S. Chua, Beyond persuasion: Towards conversational recommender system with credible explanations, arXiv preprint arXiv:2409.14399 (2024).

[9] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. R. Narasimhan, Y. Cao, React: Synergizing reasoning and acting in language models, in: The Eleventh International Conference on Learning Representations.

[10] D. Kang, A. Balakrishnan, P. Shah, P. A. Crook, Y.-L. Boureau, J. Weston, Recommendation as a communication game: Self-supervised bot-play for goal-oriented dialogue, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 1951–1961.

[11] R. Li, S. Ebrahimi Kahou, H. Schulz, V. Michalski, L. Charlin, C. Pal, Towards deep conversational recommendations, Advances in neural information processing systems 31 (2018).

[12] K. McCarthy, Y. Salem, B. Smyth, Experience-based critiquing: Reusing critiquing experiences to improve conversational recommendation, in: International Conference on Case-Based Reasoning, Springer, 2010, pp. 480–494.

[13] Y. Sun, Y. Zhang, Conversational recommender system, in: The 41st international acm sigir conference on research & development in information retrieval, 2018, pp. 235–244.

[14] X. Zhang, H. Xie, H. Li, J. CS Lui, Conversational contextual bandit: Algorithm and application, in: Proceedings of the web conference 2020, 2020, pp. 662–672.

[15] Q. Chen, J. Lin, Y. Zhang, M. Ding, Y. Cen, H. Yang, J. Tang, Towards knowledge-based recommender dialog system, arXiv preprint arXiv:1908.05391 (2019).

[16] D. Di Palma, Retrieval-augmented recommender system: Enhancing recommender systems with large language models, in: Proceedings of the 17th ACM Conference on Recommender Systems, 2023, pp. 1369–1373.

[17] S. Dai, N. Shao, H. Zhao, W. Yu, Z. Si, C. Xu, Z. Sun, X. Zhang, J. Xu, Uncovering chatgpt's capabilities in recommender systems, in: Proceedings of the 17th ACM Conference on Recommender Systems, 2023, pp. 1126–1132.

[18] Z. He, Z. Xie, R. Jha, H. Steck, D. Liang, Y. Feng, B. P. Majumder, N. Kallus, J. McAuley, Large language models as zero-shot conversational recommenders, in: Proceedings of the 32nd ACM international conference on information and knowledge management, 2023, pp. 720–730.

[19] S. Sanner, K. Balog, F. Radlinski, B. Wedin, L. Dixon, Large language models are competitive near cold-start recommenders for language-and item-based preferences, in: Proceedings of the 17th ACM conference on recommender systems, 2023, pp. 890–896.

[20] Y. Feng, S. Liu, Z. Xue, Q. Cai, L. Hu, P. Jiang, K. Gai, F. Sun, A large language model enhanced conversational recommender system, arXiv preprint arXiv:2308.06212 (2023).

[21] Y. Liu, W. Zhang, Y. Chen, Y. Zhang, H. Bai, F. Feng, H. Cui, Y. Li, W. Che, Conversational recommender system and large language model are made for each other in e-commerce pre-sales dialogue, in: Findings of the Association for Computational Linguistics: EMNLP 2023, 2023, pp. 9587–9605.

[22] X. Huang, J. Lian, Y. Lei, J. Yao, D. Lian, X. Xie, Recommender ai agent: Integrating large language models for interactive recommendations, ACM Transactions on Information Systems 43 (4) (2025) 1–33.

[23] X. Wang, X. Tang, W. X. Zhao, J. Wang, J.-R. Wen, Rethinking the evaluation for conversational recommendation in the era of large language models, in: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 2023, pp. 10052–10065.

[24] C. Huang, P. Qin, Y. Deng, W. Lei, J. Lv, T.-S. Chua, Concept–an evaluation protocol on conversation recommender systems with system-and user-centric factors, arXiv preprint arXiv:2404.03304 (2024).

[25] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al., Training language models to follow instructions with human feedback, Advances in neural information processing systems 35 (2022) 27730–27744.

[26] N. Stiennon, L. Ouyang, J. Wu, D. Ziegler, R. Lowe, C. Voss, A. Radford, D. Amodei, P. F. Christiano, Learning to summarize with human

feedback, Advances in Neural Information Processing Systems 33 (2020) 3008–3021.

[27] Y. Chow, A. Tulepbergenov, O. Nachum, M. Ryu, M. Ghavamzadeh, C. Boutilier, A mixture-of-expert approach to rl-based dialogue management, arXiv preprint arXiv:2206.00059 (2022).

[28] J. T. Cacioppo, R. E. Petty, C. F. Kao, R. Rodriguez, Central and peripheral routes to persuasion: An individual difference perspective., Journal of personality and social psychology 51 (5) (1986) 1032.

[29] C. Huang, P. Qin, W. Lei, J. Lv, Reduce human labor on evaluating conversational information retrieval system: A human-machine collaboration approach, in: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 2023, pp. 10876–10891.

[30] J. Fang, S. Gao, P. Ren, X. Chen, S. Verberne, Z. Ren, A multi-agent conversational recommender system, arXiv preprint arXiv:2402.01135 (2024).

[31] Y. Jin, L. Chen, W. Cai, X. Zhao, Crs-que: A user-centric evaluation framework for conversational recommender systems, ACM Transactions on Recommender Systems 2 (1) (2024) 1–34.

[32] X. Wang, D. Wang, C. Xu, X. He, Y. Cao, T.-S. Chua, Explainable reasoning over knowledge graphs for recommendation, in: Proceedings of the AAAI conference on artificial intelligence, Vol. 33, 2019, pp. 5329–5336.

[33] Q. Shen, L. Wu, Y. Zhang, Y. Pang, Z. Wei, F. Xu, B. Long, J. Pei, Multi-interest multi-round conversational recommendation system with fuzzy feedback based user simulator, ACM Transactions on Recommender Systems 2 (4) (2024) 1–29.