



PDF Download  
3742413.3789165.pdf  
23 March 2026  
Total Citations: 1  
Total Downloads: 0

 Latest updates: <https://dl.acm.org/doi/10.1145/3742413.3789165>

RESEARCH-ARTICLE

## Explanation Driving Exploration: Aligning Conversational Recommender Systems with Users' Exploratory Information Needs

MANVEER KALIRAI, University of Toronto, Toronto, ON, Canada

ANASTASIA KUZMINYKH, University of Toronto, Toronto, ON, Canada

Open Access Support provided by:

University of Toronto

Published: 22 March 2026

[Citation in BibTeX format](#)

IUI '26: 31st International Conference on Intelligent User Interfaces  
March 23 - 26, 2026  
Paphos, Cyprus

Conference Sponsors:  
SIGCHI  
SIGAI

# Explanation Driving Exploration: Aligning Conversational Recommender Systems with Users' Exploratory Information Needs

Manveer Kalirai  
Faculty of Information  
University of Toronto  
Toronto, Ontario, Canada  
manveer.kalirai@mail.utoronto.ca

Anastasia Kuzminykh  
Faculty of Information  
University of Toronto  
Toronto, Ontario, Canada  
Department of Computer Science  
University of Toronto  
Toronto, Ontario, Canada  
anastasia.kuzminykh@utoronto.ca

## Abstract

LLM-powered conversational recommender systems (CRSs) empower users to personalize recommendation services, giving them control over how recommendations are represented and explained. Explanations of why particular options are recommended are shown to be especially valuable when users explore unfamiliar items. While prior work on personalized explanations in recommender systems has focused predominantly on explanation style, there is still little understanding of what types of information explanations should contain to meaningfully support users' exploration. To allow CRSs to better align the explanations with users' informational needs, in this paper, we present the information composition for recommendation explanations. Informed by an exploratory interview-based user study, we propose four key informational dimensions: Essence, Experience, Exchange, and Entwinement. We then report a comparative evaluation showing that explanations structured along these dimensions are perceived as more supportive of engagement-related outcomes than baseline LLM-generated explanations. We conclude by outlining design implications for LLM-powered CRSs.

## CCS Concepts

• **Human-centered computing** → HCI theory, concepts and models; *User studies*.

## Keywords

conversational recommender systems, explanations, personalization, information exploration, information needs, leisure domains

## ACM Reference Format:

Manveer Kalirai and Anastasia Kuzminykh. 2026. Explanation Driving Exploration: Aligning Conversational Recommender Systems with Users' Exploratory Information Needs. In *31st International Conference on Intelligent User Interfaces (IUI '26)*, March 23–26, 2026, Paphos, Cyprus. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3742413.3789165>



This work is licensed under a Creative Commons Attribution 4.0 International License. *IUI '26, Paphos, Cyprus*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1984-4/26/03  
<https://doi.org/10.1145/3742413.3789165>

## 1 Introduction

Recommender systems help people explore vast content libraries and discover new material to consume. They generate personalized recommendations by inferring what might interest a user in a given moment, drawing on various types of knowledge about the user, the available items, and the surrounding context of interaction [37]. However, the usefulness of these recommendations depends on whether users can understand and evaluate why particular suggestions were made. Based on the representation of those recommendations alone, users often struggle to trace the system's underlying reasoning to evaluate how well the suggestions fit their current information needs and context [43, 44]. Explanations help bridge this gap by communicating the system's reasoning to the user in an intelligible, human-understandable way. Traditionally, explainable recommender systems have followed a one-size-fits-all approach, offering the same explanation to every user [9]. However, recent work increasingly emphasizes the importance of personalized explanations [30], looking into different dimensions of design choices (e.g., style, scope, format, and level of detail) [2, 9, 31].

The effectiveness of a personalized, user-centred explanation depends on the specific explanatory goal it is intended to achieve. In the recommender systems literature, seven goals are most commonly discussed: *transparency*, *scrutability*, *trust*, *effectiveness*, *persuasiveness*, *efficiency*, and *satisfaction* [2, 19, 43, 44]. Tintarev and Masthoff [44] observe that while additional goals have been proposed, the field has yet to converge on a unified framework for categorizing them. This suggests that the landscape of explanatory goals is still evolving, reflecting the expanding range of user needs and design priorities addressed by recommender systems. One emerging explanatory goal is *engagement*, defined by Musto et al. [29] as an explanation that helps users discover new information about an item. Engagement-oriented explanations have gained importance in light of concerns that recommender systems may reinforce consumption bubbles and restrict users' exposure to content beyond their prior consumption history [16, 17]. A growing body of work emphasizes the importance of explanations that support users' higher-order goals of *self-actualization* [22, 23, 42]. Self-actualization goes beyond simple information-finding to involve active discovery and exploration [44]. It is also intentional, giving users greater agency in directing their interactions towards specific goals [44]. Explanations that serve such aims help users

make more reflective and autonomous choices; enabling them to align their consumption with specific epistemic goals they might care about, such as broadening their horizons and discovering the unexplored [42, 44].

However, there remains limited understanding of how to structure the informational content of personalized explanations that fulfill the user-centred goal of engagement. In a comparative study of human- and algorithmically-generated explanations in a movie recommender system, Lu et al. [25] found that human-generated explanations elicited greater curiosity among participants. On closer analysis, they observed that one defining feature of human explanations is their integration of multiple information types rather than focusing on a single type. Humans typically combined three to five distinct kinds of information in their explanations, whereas system-generated explanations often included just one kind. The mixed composition of human explanations achieved higher informational accuracy because it was more likely to align with users' own reasoning. It still remains unclear what range and balance of information should be represented in explanations to meaningfully support engagement. This question is especially salient in current discussions of balanced explanations [16, 17]—those that neither push users too strongly towards a recommendation nor discourage them unnecessarily, but instead provide a fair and informative basis for exploration.

Questions about the informational balance of explanations now intersect with a broader shift in how explanations themselves are being produced. Large Language Models (LLMs) are increasingly embedded in conversational recommender systems, where users engage with recommendations and their explanations through interactive dialogue rather than passive consumption [26, 40]. These systems turn explanation into part of an ongoing conversational experience, allowing users to question, refine, or reinterpret recommendations in real time [51]. LLMs are also being used to enable users to create their own personalized conversational recommendation services from scratch, for example through custom GPTs on ChatGPT [50]. LLMs are particularly effective at producing fluent, natural language text, allowing explanations to feel more conversational and personally attuned. As explanation generation moves from deterministic to generative paradigms [48], earlier design knowledge on how to personalize explanations of recommendations may no longer suffice. Weisz et al. [47] argue that generative AI has inherently different properties (e.g., generative variability) which means that design principles developed for deterministic systems may not transfer directly to this new environment.

Collectively, this points to an open question about how explanatory information might be structured to reflect the breadth of users' exploratory informational needs and engagement goals, particularly as explanation design moves towards generative AI environments. Correspondingly, our two-part study first investigated **RQ1**: *What is the spectrum of information needs that people have about an item when exploring new, unfamiliar items they have not previously encountered?* We examined this question through an example recommendation scenario in a leisure domain—specifically, that of exploring books for pleasure reading. This scenario allowed us to ground and operationalize our investigation within a scoped yet representative recommendation context, while still capturing the

broader phenomenon of user information needs in content exploration. Through 30 semi-structured interviews, we identified four information dimensions that characterize this spectrum of user information needs: *Essence*, the fixed, intrinsic properties of an item; *Experience*, the relation of an item to the user's personal experience; *Exchange*, the relation of an item to other social actors across the user's various interpersonal spheres; and *Entwinement*, the way an item is situated in relation to other items.

Our second, evaluative study was guided by **RQ2**: *How do explanations structured along the four identified information dimensions perform in supporting user engagement across different content domains?* We conducted a within-subjects survey with 55 participants that extended this investigation beyond books to other popular leisure domains, including movies, music, and podcasts. Each participant was presented with several recommendation scenarios featuring fictional personas and was shown, side by side, two explanations for each item: one LLM-generated explanation structured according to our four information dimensions, and one LLM-generated explanation generated without reference to this structure. Participants rated which explanation they found more useful across a series of engagement-related statements. This design allowed us to assess whether explanations grounded in our proposed structure better supported engagement across different content domains. Across domains, the structured explanations received consistently higher engagement ratings. These findings suggest that organizing explanatory information along the four proposed dimensions can improve the engagement potential of LLM-generated explanations in exploratory recommendation contexts, particularly in leisure domains.

The implications of our findings point to the need for systematic approaches to structuring explanatory information in generative recommendation contexts. Rather than focusing solely on stylistic personalization, designers must consider how different kinds of information jointly shape users' exploration and sensemaking processes. This paper contributes (1) an empirically grounded framework of four information dimensions—*Essence*, *Experience*, *Exchange*, and *Entwinement*—that characterize users' exploratory information needs, and (2) an evaluation demonstrating that explanations structured around these dimensions enhance perceived engagement across multiple leisure domains. Together, these contributions advance the design of explanation facilities in conversational recommender systems towards supporting exploration as a process of sensemaking rather than mere justification.

## 2 Related Work

This section first reviews existing design knowledge on designing personalized explanations for recommender systems, and then highlights emerging developments that underscore the need to extend this knowledge towards supporting user-centred exploration and engagement.

### 2.1 Designing Personalized Explanations for Recommender Systems

Personalized explanations have long been presented as a way to make recommender systems more intelligible, yet most existing approaches optimize for presentation fit rather than information

relevance. Early frameworks, such as those introduced by Naiseh et al. [31] and Ain et al. [2], outlined the conceptual space of personalization—clarifying *what* to adapt, *how*, *for whom*, and *why*—but largely treated explanation generation as a static mapping between user traits and stylistic parameters. Empirical studies reinforced that user preferences vary with cognitive ability, domain familiarity, and situational urgency [3, 9], yet these adaptations typically adjust surface-level detail or tone rather than the composition of information itself. The result is a body of work where personalization enhances trust and satisfaction but does little to address the underlying informational needs that shape how users interpret, learn from, and act upon explanations.

A more recent wave of studies has begun to foreground the content logic of explanations. Research on open user models [14] and on-demand explanatory interfaces [13] emphasizes that personalization should reflect users' reasoning processes, organizing information in ways that align with how they make sense of a domain. Parallel technical advances such as logic-scaffolding [36] and prompt-aware frameworks [52] demonstrate that large language models (LLMs) can flexibly tailor explanatory narratives, but the criteria guiding what information to include remain opaque. Even multimodal explanation efforts [50] focus on diversity rather than informational intent. Collectively, these studies expose a persistent gap: we still lack a principled understanding of what informational components best fulfill users' engagement-oriented goals, especially in dynamic, conversational settings. This gap motivates our exploration of structured explanation design that explicitly organizes informational content around users' unfolding informational needs within LLM-powered recommender interactions.

## 2.2 Extending Explanation Design Towards Engagement Goals

Recent literature increasingly recognizes that explanations do more than justify algorithmic output—they mediate engagement by influencing how users explore, compare, and learn. Similar trends have been observed beyond recommender systems; for example, Brachman et al. show that developers value explanations that adapt to their contextual goals and cognitive state, underscoring that effective explanation design must align with users' situational information needs rather than static trust objectives [7]. Early work conceptualized explanation objectives mainly in terms of trust or transparency [16, 17], but empirical findings suggest these goals alone do not sustain user attention or foster understanding. Studies by Tran et al. [46] and Ahmad et al. [1] show that explanation design affects not only decision satisfaction but also the depth of interaction users pursue with recommended content. Balanced or two-sided explanations [16] and human-like peer explanations [25] reveal that engagement depends on how explanations frame contrast and uncertainty, inviting users to deliberate rather than merely accept. Such results position engagement as an explanatory goal in its own right—one that arises when users perceive explanations as resources for exploration and meaning-making, not simply as transparency artifacts.

LLM-based conversational recommenders extend this potential by enabling iterative sensemaking. As Yun and Lim [50] observe,

users can now question, re-frame, and co-construct recommendations, transforming explanation from a one-way justification into an ongoing dialogue between user and system. This conversational flexibility, however, also introduces instability—LLM-generated explanations can vary in depth, focus, and epistemic framing across turns, leading to inconsistency in the informational composition that sustains engagement [26, 50]. Recent advances in explainable AI highlight this challenge: Lai et al. [24] advocate for selective explanations that curate which reasoning aspects are exposed to preserve cognitive coherence, while work on sensemaking interfaces [12, 41, 49] emphasizes the value of layered or structured representations that scaffold users' interpretive work. Informed by these insights, our study investigates how explanations in LLM-powered conversational recommender systems can be systematically structured to surface distinct informational dimensions that align with users' evolving curiosity and evaluative needs, thereby repositioning explanation design from persuasion toward sustained exploratory engagement.

In sum, existing work highlights the promise but also the uncertainty of how explanations can meaningfully engage users in LLM-powered recommendation. Our study addresses this by first uncovering the informational needs that shape engagement, and then examining how explanations that embody these needs perform in practice.

## 3 Method

To investigate how explanatory information can be structured to better support user engagement in generative recommendation contexts, we conducted a two-part study. The protocols for The first study employed semi-structured interviews ( $n=30$ ) to explore the breadth of users' informational needs during content exploration and identify a foundational set of information dimensions (**RQ1**). Building on these findings, the second study employed a within-subjects, scenario-based survey ( $n=55$ ) to evaluate how LLM-generated explanations structured according to these dimensions perform against those generated without such structure (**RQ2**). Together, these studies were designed to provide both conceptual and empirical insight into how explanation design can better support engagement in generative recommendation settings. The study protocols for both studies were approved by the University's institutional review board.

### 3.1 Study 1: Exploring User Information Needs

**3.1.1 Data Collection.** For **RQ1**—*What is the spectrum of information needs that people have about an item when exploring new, unfamiliar items they have not previously encountered?*—we conducted a user study involving semi-structured interviews ( $n=30$ ) with individuals who actively read for pleasure or leisure. We chose to focus on a single domain, i.e., book exploration for pleasure reading, to enable participants to reflect concretely on their information needs within a bounded and familiar context. This domain offers a strong case for studying exploratory behaviour: it involves navigating a vast, information-rich landscape where options (books) and decision-making information interact in complex ways, as demonstrated in prior library and information science research [11, 38, 39]. However, while this body of work highlights the richness of book

selection as an information activity, it does not clarify the specific types of information users actively seek, interpret, and weigh in this process. Books differ from many other cultural items in that their consumption can require a relatively greater substantial time investment, making it difficult to quickly assess preference. Unlike media such as songs, which can be sampled and revisited with minimal cost, evaluating a book often requires prolonged engagement [27]. For this reason, readers generally rely on a range of information signals when exploring unfamiliar books, heightening the importance of understanding their information needs in this context [11, 39].

**Participants.** To participate in the interview study, individuals required sufficient English proficiency to engage comfortably in daily social interactions and to understand and sign the consent form. We recruited a total of 30 participants (P1–P30; 17 women, 13 men; aged 19–41,  $M=28$ ,  $SD=5.3$ ). This sample size aligns with established practices in qualitative and interpretive HCI research. Recent work in the HCI literature on related topics of recommendation and information behaviour research similarly used interview study samples of comparable size to explore rich experiential and informational phenomena [6, 21, 28, 34].

Apart from being over the age of 18 and proficient in English, the only other requirement was that participants should engage in ‘pleasure reading’—i.e., self-directed reading motivated by personal interest, enjoyment, or curiosity rather than external obligations or requirements [32, 38]. Our recruitment approach used a combination of snowball sampling (i.e., asking participants to refer others who might be a good fit for the study) and purposive sampling within the researchers’ networks to ensure a balance of genders and ages. During recruitment, prospective participants were given an overview of the study’s purpose and high-level topics that would be covered in the interview. Those interested provided informed consent. Participants were scheduled based on their availability, and all participation was voluntary, without compensation.

**Procedure.** Each participant was invited to meet, one-on-one, with the primary researcher in a 45-minute to hour-long interview session conducted virtually over a video call on Zoom. With participant consent, each session was audio-recorded. The interviews were semi-structured in format, combining prepared questions with flexibility to allow for both consistency across interviews and open-ended exploration based on participants’ responses. This approach was well suited to the study’s goal of eliciting the range of information elements users consider when exploring unfamiliar items, examined here in the context of book search.

Interviews began with an open-ended prompt to ground discussion in participants’ experiences (e.g., “*Can you describe a recent book you chose to read?*”). Subsequent questions addressed how participants discover new books (e.g., “*How do you typically find books to read?*”), the types of information they consider when deciding whether to engage (e.g., “*What kinds of information do you consider when deciding if a book is a good fit?*”), and how these considerations vary across contexts or goals (e.g., “*Does what you look for in a book change depending on the situation?*”). Additional questions invited reflection on expectations and outcomes (e.g., “*How do you determine whether reading a book was worthwhile?*”). Throughout the interviews, the primary researcher used follow-up probes as needed to clarify responses, encourage elaboration, and elicit

concrete examples, enabling deeper exploration of influences that more structured methods (e.g., surveys or questionnaires) may not fully capture [15]. Interviews continued until thematic saturation was reached, at which point participants no longer introduced new types of information need or unique insights into their information experience.

**3.1.2 Data Analysis.** All 30 interviews were transcribed using Zoom’s automatic transcription service. Each transcript was reviewed against the corresponding audio recordings to ensure accuracy, then anonymized. Using Braun and Clarke’s qualitative thematic analysis approach [8], the primary researcher conducted two iterative rounds of initial coding on the transcripts to systematically identify emerging themes, capturing patterns in information needs and refining them through iterative analysis. Both researchers then collaborated to review, consolidate, and organize the codes, structuring them hierarchically to reflect relationships within and across categories. Considering these relationships, the researcher developed a two-level hierarchy, and mapped the identified thematic elements accordingly. This culminated in the development of a thematic framework consisting of four overarching dimensions of information, representing the highest hierarchical level. Within these dimensions, 14 constituent categories were identified, each encapsulating specific aspects of information. This thematic framework is outlined in Table 1.

## 3.2 Study 2: Evaluating Structured Explanations

**3.2.1 Experiment Design.** For **RQ2—How do explanations structured along the four identified information dimensions perform in supporting user engagement across different content domains?**—we conducted a controlled A/B survey experiment. To examine how well the four information dimensions identified in Study 1 generalized beyond the domain of books, we extended our evaluation to other popular leisure domains: movies, music, and podcasts. In total, we designed two recommendation scenarios per domain (2 scenarios  $\times$  4 domains = 8 scenarios in total).

Collecting participant-specific data at the level required to instantiate all four information dimensions across multiple domains and scenarios would have required extensive elicitation of personal, social, and contextual information. This would potentially increase participant burden, introduce additional ethical considerations around disclosure. We therefore used fictional personas to standardize the informational context across participants. Fictional personas allowed us to represent a broader range of user information than is typically available in benchmark datasets or personal platform logs, in a way that aligns with how LLM-powered conversational recommender systems are envisioned. These personas were created using ChatGPT 5. First, we prompted the model to generate a blank version of an expanded user-data schema in YAML format—a human-readable hierarchical template that simulated the kind of data a conversational AI might use for personalization. Then, we asked the model to populate this schema to produce eight diverse user profiles. Each persona was randomly assigned to a content domain: Karen→Books, Antoine→Books, Mike→Movies, Sara→Movies, Liam→Music, Emily→Music, Brian→Podcasts, and David→Podcasts. Using each completed profile, ChatGPT was prompted to recommend one item within that domain.

**Table 1: Information dimensions and their corresponding information categories**

Information Dimension	Information Category
<b>Essence</b>	Item Metadata
	Item Subject Matter
	Item Stylistic Expression
	Item Format & Form Factor
<b>Experience</b>	Item Access to and Availability
	Alignment with Time Capacity
	Alignment with Effort/Skill Capacity
	Alignment with Affective/Sensory Capacity
	Alignment with Values/Identity
<b>Exchange</b>	Personal and Direct Social Exchange
	Community and Shared Interest Exchange
	Public and Cultural Exchange
<b>Entwinement</b>	Relationships to Similar Items of the Same Kind (Same Media Type)
	Relationships to Different Kinds of Items (Different Media Types)

For each scenario, we generated two explanations using ChatGPT 5: one *structured* and one *unstructured*. Drawing on Silva et al.'s [40] approach to explanation generation, the unstructured condition used a generic prompt asking the system to explain why [PERSONA] should [READ or WATCH or LISTEN TO] [RECOMMENDED ITEM—BOOK or MOVIE or ALBUM or PODCAST], given their data profile. The structured condition used the same persona and item but instructed the model to organize its reasoning according to the four dimensions identified in Study 1: *Essence*, *Experience*, *Exchange*, and *Entwinement*. Both conditions included formatting constraints (e.g., second-person phrasing and 100-word length). The full prompts and corresponding outputs used in the experiment are provided in the Appendix.

The study followed a within-subjects design where each participant was shown all eight recommendation scenarios across the four domains, presented in randomized order to mitigate order effects. In the survey, participants first read a short persona description (250–300 words), generated by ChatGPT based on the aforementioned schema, and were then told which item the system had recommended to that persona. Two explanations—one structured and one unstructured—were shown side by side, labelled “A” and “B.” The assignment of structured versus unstructured explanations to labels “A” and “B” was counterbalanced such that each condition appeared in each position in equal proportions across participants. Participants rated, on a 7-point scale, which explanation they believed better supported the persona on four engagement-related statements: (1) Which explanation helps [PERSONA] discover more information about the item? (2) Which explanation helps [PERSONA] understand the item in a way that feels more thoughtful or meaningful to them? (3) Which explanation would make [PERSONA] more likely to stay interested or keep looking for related things? (4) Which explanation gives [PERSONA] a broader sense of what the item relates to or represents?

The 7-point comparison scale ranged from “A much more” to “B much more” (i.e., A/B slightly, moderately, much more, with “about the same” as the midpoint).

**3.2.2 Data Collection & Participants.** Data collection was conducted through the research crowdsourcing platform *Prolific*. We recruited a total of 55 participants located in Canada and the United States who reported fluency in English. Recruitment was designed to achieve a balanced gender distribution across participants. Eligible participants were compensated at an average rate of £6.98 per hour. After completing the initial screening on the *Prolific* platform, participants were redirected to the main survey hosted on *Qualtrics*.

## 4 Findings

### 4.1 Study 1: Exploring User Information Needs

Through our analysis, we first characterized the different types of information needs that emerge during content exploration as “information dimensions.” We identified four information dimensions: *Essence* (the fixed, intrinsic properties of an item), *Experience* (the relation of an item to the user’s personal experience); *Exchange* (the relation of an item to other social actors across the user’s various interpersonal spheres); and *Entwinement* (the way an item is situated in relation to other items). We elaborate on each of these dimensions in greater detail, below.

**4.1.1 Essence.** This dimension represents information needs related to the fixed, intrinsic properties of an item. Participants described focusing on four main aspects: *Item Metadata* (e.g., objective identifiers such as genre, creator, or release details), *Item Subject Matter* (e.g., the central topic, theme, or message), *Item Stylistic Expression* (e.g., the tone, aesthetic, or expressive manner), and *Item Format and Form Factor* (e.g., the container or delivery mode). Across examples, participants first referred to basic identifiers like title, creator, or release timing as familiar anchors (“There’s this book by Ursula K. Le Guin called *The Dispossessed*... The title and author caught my eye” [P22]; “I prefer books that were published relatively recently” [P2]). They then sought details about the item’s subject matter and expressive style to understand what it contained and how it might feel to engage with (“Sometimes, dialogue is what makes a book stand out to me” [P23]; “If the writing style is pretty

bad, I am not going to survive” [P12]). Format and form also influenced decision-making, with participants noting factors like length, density, and delivery medium (“Seeing how thick a book is can be a huge deterrent” [P2]; “A good narrator can make the experience better” [P9]). Together, these accounts show that participants wanted to know enough about an item’s intrinsic characteristics to form a grounded sense of what it is, what engaging with it might involve, and whether it warranted further exploration.

**4.1.2 Experience.** This dimension represents information needs related to how an item fits with a person’s own circumstances, capacities, and sense of self. Participants described attending to four main aspects: *Alignment with Time Capacity* (i.e., how the expected duration or pacing of the item fits into one’s available time or attention), *Alignment with Effort or Skill Capacity* (i.e., how approachable or demanding the item seems given one’s current familiarity or energy), *Alignment with Affective or Sensory Capacity* (i.e., how its emotional tone or sensory qualities suit one’s current state or preference), and *Alignment with Values or Identity* (i.e., how it connects with one’s beliefs, priorities, or lived experiences). Participants often evaluated whether they had the time or focus to engage with an item (“How many books in the series? Like how much time I’m going to have to invest in this?” [P18]) and whether it matched their mental bandwidth (“Picking the right one that I can keep reading regardless of how much energy I have would help me keep the habit” [P27]). They also considered whether its mood or tone fit their current emotional state (“I need to know if something is sad or not” [P15]) and valued items that reflected or deepened their own perspectives (“It just felt so reminiscent of real life... It made me think about how I’ve seen that play out in my own life” [P24]). Overall, participants used these considerations to judge whether an unfamiliar item felt feasible, emotionally suitable, and personally worthwhile to explore.

**4.1.3 Exchange.** This dimension represents information needs related to how an item is situated within social relationships and shared cultural contexts. Participants described attending to three main aspects: *Personal and Direct Exchange* (i.e., recommendations or influences from people they personally know), *Community and Shared-Interest Exchange* (i.e., input from groups, forums, or communities built around common interests), and *Public and Cultural Exchange* (i.e., signals from broader media, trends, or cultural discourse). Participants frequently mentioned trusting suggestions from friends, family, or mentors (“I picked it up because a friend of mine told me it was one of those books you had to read at least once in your life” [P25]) and being drawn to items discussed within their social circles or interest communities (“I actually ended up reading it because a couple of people in a book club I used to go to were talking about it” [P22]; “I just saw someone mention it on Reddit... the way people described it made me pick it up” [P30]). They also noticed items highlighted through broader social channels such as online trends, reviews, or cultural lists (“I have a TikTok folder called BookTok... Sometimes they’re like, ‘If you liked this, you’ll love this’” [P18]). Across these accounts, participants drew on others’ experiences and attention at varying levels of interpersonal proximity—from personal networks to wider communities—to interpret unfamiliar items and judge their relevance.

**4.1.4 Entwinement.** This dimension represents information needs related to how an item connects to other items within or beyond its media-type. Participants described attending to two main aspects: *Relationships with Similar Items* (i.e., connections to other items of the same media-type, such as those sharing genre, structure, or tone) and *Relationships with Different Items* (i.e., connections across other media-types, such as links between books, films, music, or other media). Participants often compared unfamiliar items to ones they had previously enjoyed, looking for continuity in mood, style, or theme (“If I really liked something, I want to find something that feels similar... not necessarily the same plot, but a similar writing style or a similar mood” [P11]; “I tend to read in clusters... I’ll look up things similar to it because I want to stay in that headspace” [P14]). They also noticed cross-media references that extended their engagement, such as adaptations, playlists, or discussions in other formats (“You’ll see a movie that’s based on a book and then think, ‘Well, maybe I should read that book’” [P7]; “Sometimes there’s a playlist associated with it, and I’ll listen to that while reading to get into the mood” [P18]). Across these accounts, participants mapped unfamiliar items against known references—both within the same medium and across different ones—using these relationships to position new material within a broader structure of meaning.

## 4.2 Study 2: Evaluating Structured Explanations

To empirically assess the effectiveness and generalizability of the four information dimensions identified in Study 1, we compared how explanations structured according to these dimensions performed relative to unstructured explanations in supporting engagement-related outcomes. Specifically, this study examined whether structured explanations would help participants perceive a persona’s engagement with recommended items across four complementary aspects: information discovery, meaningful understanding, sustained interest, and contextual breadth. These aspects reflect how the persona would come to learn about, make sense of, remain interested in, and situate the recommended content within a broader context of meaning.

Across all four engagement aspects, participants consistently perceived the structured explanations as more supportive of the persona’s engagement than the unstructured baselines. Mean ratings ranged between 2.54 and 2.64 ( $SD \approx 1.4\text{--}1.5$ ) on the seven-point comparison scale (1 = *A* much more, 4 = about the same, 7 = *B* much more), indicating a clear directional preference toward the structured condition (values below 4). To test the reliability of these differences, one-sample t-tests were conducted for each engagement aspect against the neutral midpoint (4 = about the same). Results showed that all mean differences were statistically significant ( $p < .001$ ), confirming that participants perceived structured explanations as more engaging across all aspects (see Table 2).

The aggregated mean of 2.57 also differed significantly from the midpoint ( $t(1759) = -41.56, p < .001$ ), demonstrating a robust overall tendency to favour structured explanations. This pattern was consistent across all content domains—books, movies, music, and podcasts—suggesting that the perceived advantage of structured explanations was not restricted to a particular domain.

**Table 2: Comparison of structured (Condition A) and unstructured (Condition B) explanations across engagement questions. One-sample *t*-tests were conducted against the neutral midpoint (4 = “about the same”). Lower means indicate stronger preference for the structured condition.**

Engagement Question	M	SD	<i>t</i> (439)	<i>p</i>	95% CI [LL, UL]
Which explanation helps the persona discover more information about the item?	2.58	1.42	-21.05	< .001	[2.44, 2.72]
Which explanation helps the persona understand the item in a way that feels more thoughtful or meaningful to them?	2.54	1.47	-20.86	< .001	[2.40, 2.68]
Which explanation would make the persona more likely to stay interested or keep looking for related things?	2.54	1.42	-21.62	< .001	[2.41, 2.67]
Which explanation gives the persona a broader sense of what the item relates to or represents?	2.64	1.46	-19.59	< .001	[2.50, 2.78]
<b>Overall (all items combined)<sup>a</sup></b>	2.57	1.44	-41.56	< .001	[2.51, 2.63]

Note. One-sample *t*-tests compared mean ratings to midpoint (4 = “about the same”).

<sup>a</sup>Overall test based on all ratings combined ( $n = 1760$ ).

These findings suggest that organizing explanatory information according to the four identified dimensions can enhance engagement-related experiences. While these results are based on persona-based perceptions rather than direct behavioural measures, they provide converging evidence that the information-dimension framework identified in Study 1 can support the design of explanations that better align with exploratory processes across diverse leisure content domains.

### 4.3 Summary of Study 1 & 2

The two studies collectively contribute to understanding what kinds of information people attend to when exploring unfamiliar content and how this information can be structured to support engagement. Study 1 (addressing **RQ1**) examined users' exploratory information-seeking practices, identifying four recurring information dimensions—*Essence*, *Experience*, *Exchange*, and *Entwinement*—that capture the informational considerations shaping how approach unfamiliar content. Study 2 (addressing **RQ2**) then tested how explanations structured around these dimensions were perceived relative to unstructured ones, finding that participants generally viewed the structured explanations as more supportive of engagement across multiple aspects. Together, these findings indicate that the identified information dimensions provide a useful lens for characterizing exploratory information needs and informing the design of more intelligible, user-aligned explanations in recommendation contexts.

## 5 Discussion

This work introduces an information structure composed of four dimensions—*Essence*, *Experience*, *Exchange*, and *Entwinement*—that reflect the spectrum of users' informational needs for the explanatory goal of engagement. The structure offers a principled way to organize explanatory content in generative recommender systems, foregrounding what kinds of information users require to make sense of unfamiliar recommendations. By emphasizing explanatory

composition rather than presentation style—which has been the predominant focus of previous literature looking at design choices for explanation personalization (e.g., [2, 31])—it contributes to a growing body of research that re-conceptualizes explanations in recommender systems as instruments of *sensemaking* rather than as mere *justifications* of algorithmic reasoning [29, 50].

Our findings suggest that explanatory usefulness, specifically for the goal of engagement, can be operationalized through informational composition, which we have represented at four dimensions of information. Each dimension represents a distinct cognitive or social orientation through which users interpret new items: *Essence* establishes what an item is in and of itself; *Experience* situates it relative to personal capacity, emotion, or identity; *Exchange* places it within social and cultural discourse; and *Entwinement* links it to other items, expanding relational context. These dimensions formalize a kind of informational balance that align, functionally, with what Jafari and Vassileva [16] describe as “balanced explanations”, i.e., those that invite users to deliberate rather than to simply being persuaded. They also help explain why human-generated explanations in prior work (e.g., [25]) often outperform system-generated ones: people naturally weave together multiple kinds of information when explaining recommendations, producing richer cognitive scaffolds for reasoning and exploration for the recipients of those explanations.

The findings from our evaluation study further show that explicitly embedding this structure into LLM prompts can shape users' perceptions of explanatory quality. Across domains—particularly those of books, movies, music, and podcasts—participants consistently rated structured explanations as more conducive to engagement across the four aspects we measured: information discovery, meaningful understanding, sustained interest, and contextual breadth. These outcomes are perceptual rather than behavioural, but their consistency suggests that even lightweight structural scaffolding, using the identified dimensions, can make generative outputs more coherent and intelligible to the users. This aligns

with recent evidence that prompt-level structure and logic scaffolding, e.g., [36, 52], can stabilize the reasoning expressed by Large Language Models. The practical takeaway from this work is that small design interventions—such as explicitly organizing system reasoning for a recommendation across a few stable user-centred information dimensions—may help mitigate the user experience risks associated with generative variability [47] while preserving the advantages of conversational flexibility that LLMs offer [51]; allowing explanations to unfold in response to users' curiosities rather than remain fixed outputs.

## 5.1 Limitations & Future Work

Several limitations qualify these findings but also point toward fruitful directions for future research. The evaluation (Study 2) relied on fictional personas rather than participants' own contexts, which allowed experimental control but limited ecological realism. Participants assessed how well explanations appeared to support a persona's engagement, rather than how they would support their own exploration. This distancing may reduce the role of personal relevance or affective investment in participants' judgments. In addition, fictional personas present a relatively stable and internally consistent user profile. In practice, real user data is often incomplete, evolving, or inconsistent, and explanations grounded in such data may surface different tensions or trade-offs than those observed here. For these reasons, our findings should be understood as evidence of how informational structure can shape perceived engagement under controlled conditions, rather than as a direct account of engagement in situated, long-term use. Future work should examine how these information dimensions perform when instantiated with real user data and in interactive settings.

Relatedly, 'engagement' in this study was measured through comparative judgments rather than behavioural data. While this approach aligns with our focus on perceived support for exploration, it does not capture whether structured explanations lead to enacted exploratory behaviour, such as sustained browsing or broader content consumption over time. Our focus on low-stakes leisure domains further limits generalizability. Engagement dynamics in contexts involving higher perceived risk or commitment—such as health or finance—may differ substantially, as considerations of reliability, accountability, and consequences may outweigh exploratory curiosity. Future work should therefore examine how informational composition affects behavioural engagement, for example by measuring whether structured explanations encourage users to explore more broadly or persist longer. Beyond low-stakes settings, future studies could examine domains that vary in both risk and investment. Tourism, for instance, represents a high-investment yet non-critical domain, where decisions typically unfold over extended periods of exploration and require substantial upfront commitment relative to other product domains, making it a promising context for studying how explanatory structure supports exploratory decision-making [5, 18]. Interactive prototypes could also be used to test adaptive mechanisms that adjust the emphasis of different information dimensions as users' needs evolve, extending principles of selective and layered explanation [24].

The evaluation design itself also suggests avenues for extension. While this study compared a single structured explanation framework against unstructured explanations, future work could examine multiple structured variants to better understand how different ways of organizing explanatory information shape user perceptions. Such comparisons would help disentangle the effects of structure more generally from those of any single framework. Alternative experimental designs, including between-subjects approaches in which participants evaluate a single explanation using agreement-based measures, could further support comparisons across multiple explanation variants while reducing participant cognitive load.

Future work can also explore how the proposed information dimensions are realized across different interaction modalities. In particular, presenting all four dimensions within a single explanation may be inappropriate for voice-based interfaces, where verbosity and cognitive load are key design concerns [10, 35]. Investigating how explanatory information structured along the proposed framework is delivered and distributed over the course of an interaction could help clarify how informational structure interacts with attention, timing, and modality in conversational settings.

Finally, individual differences are likely to play an important role in how explanations are interpreted and valued. Prior work suggests that personality traits and cognitive styles influence explanatory preferences [4, 20, 33, 45]. Because our study relied on standardized fictional personas, we were not able to examine these factors directly. Future research could investigate how individual characteristics relate to preferences for different informational dimensions, particularly how the *Experience* dimension might adapt to personal differences while maintaining broader interpretive balance through *Exchange* and *Entwinement*.

## 6 Conclusion

In this paper, we presented the results of an exploratory study that addresses a core design challenge that exists especially within LLM-based conversational recommender systems—determining what information needs to be presented in explanations and how it should be structured to support exploration. Using book selection as a representative case of exploratory search, we first identified four information dimensions—*Essence*, *Experience*, *Exchange*, and *Entwinement*—that characterize users' informational needs when encountering unfamiliar items. Through a follow-up evaluative study, we showed that explanations organized around these dimensions enhanced perceived engagement across multiple leisure domains (i.e. books, movies, music, podcasts). Together, these findings offer a principled framework for composing explanations that not only justify recommendations but also scaffold user sensemaking and curiosity, ultimately in service of users' epistemic self-actualization goals.

## References

- [1] Juan Ahmad, Jonas Hellgren, and Alan Said. 2025. Tell Me the Good Stuff: User Preferences in Movie Recommendation Explanations. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*. 103–108.
- [2] Qurat Ul Ain, Mohamed Amine Chatti, Mouadh Guesmi, and Shoeb Joarder. 2022. A multi-dimensional conceptualization framework for personalized explanations in recommender systems. In *Companion Proceedings of the 27th International Conference on Intelligent User Interfaces*. 11–23.

- [3] Ibrahim Al-Hazwani, Nimra Ahmed, Mennatallah El-Assady, and Jürgen Bernard. 2024. Towards Personal Explanations for Recommender Systems: A Study on the Impact of Familiarity and Urgency. In *Adjunct Proceedings of the 2024 Nordic Conference on Human-Computer Interaction*. 1–8.
- [4] Gabrielle Alves, Dietmar Jannach, Luan Soares de Souza, and Marcelo Garcia Manzano. 2025. Towards Personality-Aware Explanations for Music Recommendations Using Generative AI. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems*. 684–689.
- [5] Hans H Bauer, Nicola E Sauer, and Christine Becker. 2006. Investigating the relationship between product involvement and consumer decision-making styles. *Journal of Consumer Behaviour* 5, 4 (2006), 342–354.
- [6] Brett Binst, Lien Michiels, and Annelien Smets. 2025. What Is Serendipity? An Interview Study to Conceptualize Experienced Serendipity in Recommender Systems. In *Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*. 243–252.
- [7] Michelle Brachman, Arielle Goldberg, Andrew Anderson, Stephanie Houde, Michael Muller, and Justin D Weisz. 2025. Towards Personalized and Contextualized Code Explanations. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*. 120–125.
- [8] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [9] Mohamed Amine Chatti, Mouadh Guesmi, Laura Vorgerd, Thao Ngo, Shoeb Joarder, Qurat Ul Ain, and Arham Muslim. 2022. Is more always better? The effects of personal characteristics and level of detail on the perception of explanations in a recommender system. In *Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization*. 254–264.
- [10] Michael Harris Cohen, James P Giangola, and Jennifer Balogh. 2004. *Voice user interface design*. Addison-Wesley Professional.
- [11] Keren Dali. 2014. From book appeal to reading appeal: Redefining the concept of appeal in readers' advisory. *The Library Quarterly* 84, 1 (2014), 22–48.
- [12] Katy Ilonka Gero, Chelse Swoopes, Ziwei Gu, Jonathan K Kummerfeld, and Elena L Glassman. 2024. Supporting sensemaking of large language model outputs at scale. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [13] Mouadh Guesmi, Mohamed Amine Chatti, Laura Vorgerd, Shoeb Joarder, Shadi Zumor, Yiqi Sun, Fangzheng Ji, and Arham Muslim. 2021. On-demand personalized explanation for transparent recommendation. In *Adjunct Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*. 246–252.
- [14] Rully Agus Hendrawan, Peter Brusilovsky, Arun Balajee Lekshmi Narayanan, and Jordan Barria-Pineda. 2024. Explanations in open user models for personalized information exploration. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*. 256–263.
- [15] Kasper Hornbæk, Per Ola Kristensson, and Antti Oulasvirta. 2025. *Introduction to Human-Computer Interaction*. Oxford University Press.
- [16] Elahesh Jafari. 2024. Balanced Explanations in Recommender Systems. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*. 25–29.
- [17] Elahesh Jafari and Julita Vassileva. 2023. Ethical issues in explanations of personalized recommender systems. In *Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization*. 215–219.
- [18] Monika Jain. 2019. A study on consumer behavior-decision making under high and low involvement situations. *IJRAR-International Journal of Research and Analytical Reviews* 6, 1 (2019).
- [19] Anthony Jameson, Martijn C Willemsen, Alexander Felfernig, Marco De Gemmis, Pasquale Lops, Giovanni Semeraro, and Li Chen. 2015. Human decision making and recommender systems. In *Recommender systems handbook*. Springer, 611–648.
- [20] Yucheng Jin, Li Chen, Wanling Cai, and Xianglin Zhao. 2024. CRS-Que: A user-centric evaluation framework for conversational recommender systems. *ACM Transactions on Recommender Systems* 2, 1 (2024), 1–34.
- [21] Rutuja Joshi, Yu-Jou Lee, and Klaus Bengler. 2025. User preferences in conversational ai for healthcare: Insights from an interview study. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces*. 1–13.
- [22] Manveer Kalirai and Anastasia Kuzminykh. 2024. You Today, Better Tomorrow: Envisioning the Role of Conversation in Recommender Systems of the Future. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces*. 1–5.
- [23] Bart P Knijnenburg, Saadhika Sivakumar, and Daricia Wilkinson. 2016. Recommender systems for self-actualization. In *Proceedings of the 10th ACM conference on recommender systems*. 11–14.
- [24] Vivian Lai, Yiming Zhang, Chacha Chen, Q Vera Liao, and Chenhao Tan. 2023. Selective explanations: Leveraging human input to align explainable ai. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–35.
- [25] Hongyu Lu, Weizhi Ma, Yifan Wang, Min Zhang, Xiang Wang, Yiqun Liu, Tat-Seng Chua, and Shaoping Ma. 2023. User perception of recommendation explanation: Are your explanations what users need? *ACM Transactions on Information Systems* 41, 2 (2023), 1–31.
- [26] Sebastian Lubos, Thi Ngoc Trang Tran, Alexander Felfernig, Seda Polat Erdeniz, and Viet-Man Le. 2024. LLM-generated explanations for recommender systems. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*. 276–285.
- [27] Marshall McLuhan. 1994. *Understanding media: The extensions of man*. MIT press.
- [28] Ashlee Milton, Juan F Maestre, Abhishek Roy, Rebecca Umbach, and Stevie Chancellor. 2024. Seeking in cycles: how users leverage personal information ecosystems to find mental health information. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [29] Cataldo Musto, Fedelucio Narducci, Pasquale Lops, Marco De Gemmis, and Giovanni Semeraro. 2016. Explod: A framework for explaining recommendations based on the linked open data cloud. In *Proceedings of the 10th ACM Conference on Recommender Systems*. 151–154.
- [30] Cataldo Musto, Marco Polignano, Amon Rapp, Giovanni Semeraro, and Jürgen Ziegler. 2025. 7th Workshop on Explainable User Models and Personalised Systems (ExUM 2025). In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*. 174–176.
- [31] Mohammad Naiseh, Nan Jiang, Jianbing Ma, and Raian Ali. 2020. Personalising explainable recommendations: literature and conceptualisation. In *World conference on information systems and technologies*. Springer, 518–533.
- [32] Victor Nell. 1988. The psychology of reading for pleasure: Needs and gratifications. *Reading research quarterly* (1988), 6–50.
- [33] Ayano Okoso, Keisuke Otaki, Satoshi Koide, and Yukino Baba. 2025. Impact of tone-aware explanations in recommender systems. *ACM Transactions on Recommender Systems* 3, 4 (2025), 1–34.
- [34] Ege Otene. 2024. Towards Designing for Multimodal Remembering: Findings from an Interview Study. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–6.
- [35] Cathy Pearl. 2016. *Designing voice user interfaces: Principles of conversational experiences*. O'Reilly Media, Inc.
- [36] Behnam Rahdari, Hao Ding, Ziwei Fan, Yifei Ma, Zhuotong Chen, Anoop Deoras, and Branislav Kveton. 2024. Logic-scaffolding: Personalized aspect-instructed recommendation explanation generation using llms. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 1078–1081.
- [37] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2021. Recommender systems: Techniques, applications, and challenges. *Recommender systems handbook* (2021), 1–35.
- [38] Catherine Sheldrick Ross. 2021. Making choices: What readers say about choosing books to read for pleasure. In *Readers, reading, and librarians*. Routledge, 5–21.
- [39] Catherine Sheldrick Ross, Paulette M Rothbauer, and Lynne EF McKechnie. 2018. *Reading Still Matters*. (2018).
- [40] Ítallo Silva, Leandro Marinho, Alan Said, and Martijn C Willemsen. 2024. Leveraging chatgpt for automated human-centered explanations in recommender systems. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*. 597–608.
- [41] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling multilevel exploration and sensemaking with large language models. In *Proceedings of the 36th annual ACM symposium on user interface software and technology*. 1–18.
- [42] Emily Sullivan, Dimitrios Bountouridis, Jaron Harambam, Shabnam Najafian, Felicia Loecherbach, Mykola Makhortyk, Domokos Kelen, Daricia Wilkinson, David Graus, and Nava Tintarev. 2019. Reading news with a purpose: Explaining user profiles for self-actualization. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*. 241–245.
- [43] Nava Tintarev and Judith Masthoff. 2010. Designing and evaluating explanations for recommender systems. In *Recommender systems handbook*. Springer, 479–510.
- [44] Nava Tintarev and Judith Masthoff. 2012. Beyond explaining single item recommendations. In *Recommender Systems Handbook*. Springer, 711–756.
- [45] Marko Tkalčić and Li Chen. 2012. Personality and recommender systems. In *Recommender systems handbook*. Springer, 757–787.
- [46] Thi Ngoc Trang Tran, Alexander Felfernig, Viet Man Le, Thi Minh Ngoc Chau, and Thu Giang Mai. 2023. User needs for explanations of recommendations: In-depth analyses of the role of item domain and personal characteristics. In *Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization*. 54–65.
- [47] Justin D Weisz, Jessica He, Michael Muller, Gabriela Hofer, Rachel Miles, and Werner Geyer. 2024. Design principles for generative AI applications. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [48] Ryan W White and Chirag Shah. 2025. *Information Access in the Era of Generative AI*. Springer.
- [49] An Yan, Zhankui He, Jiacheng Li, Tianyang Zhang, and Julian McAuley. 2023. Personalized showcases: Generating multi-modal explanations for recommendations. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2251–2255.
- [50] Sojeong Yun and Youn-kyung Lim. 2025. User Experience with LLM-powered Conversational Recommendation Systems: A Case of Music Recommendation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*.

1–15.

- [51] Gangyi Zhang. 2023. User-centric conversational recommendation: Adapting the need of user with large language models. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1349–1354.
- [52] Haobo Zhang, Qiannan Zhu, and Zhicheng Dou. 2025. A unified prompt-aware framework for personalized search and explanation generation. *ACM Transactions on Information Systems* 43, 3 (2025), 1–26.

## A Generating Persona Data Templates

### Prompt to ChatGPT 5:

Create a comprehensive blank YAML template for simulating a ChatGPT user profile. The goal is to design a fictional but realistic schema that mirrors the kinds of structured data a conversational AI system might store about a user, without revealing or referencing any private or internal OpenAI data formats.

### System Output:

Listing 1 presents the simulated YAML template generated in response to the above prompt.

**Listing 1: Simulated User Persona YAML Template**

```

user_id: ""
tenant_id: ""
created_at: ""
updated_at: ""
status: ""
labels: []

account:
  email: ""
  email_verified: false
  name:
    given_name: ""
    family_name: ""
    display_name: ""
  username: ""
  plan: ""
  plan_started_at: ""
  plan_expires_at: ""
  locale: ""
  timezone: ""
  marketing_opt_in: false
  early_access_opt_in: false
  referral:
    code: ""
    source: ""

profile:
  demographics:
    age: null
    birth_year: null
    gender: ""
    pronouns: ""
    country: ""
    region: ""
    city: ""
    language_primary: ""
    languages_other: []
  roles_personas: []
  biography_short: ""
  biography_long: ""
  education:
    highest_level: ""
    fields_of_study: []
    institutions: []
  employment:
    current_title: ""
    current_org: ""
    industry: ""
    seniority: ""
    history: []
  preferences:

```

```

writing_style: []
tone: []
formatting:
  bullet_density: ""
  markdown_ok: false
  headings_ok: false
  citations_style: ""
  spelling_variant: ""
depth_defaults:
  explanation_depth: ""
  examples_preferred: false
  code_included_by_default: false
content_filters_user:
  avoid_topics: []
  avoid_phrases: []
accessibility:
  reading_ease: ""
  visual_support: []
  screen_reader_mode: false
session_behavior:
  ask_followups_automatically: false
  summarize_long_responses: false
  include_tl_dr: false
exploration_preferences:
  novelty_tolerance: ""
  serendipity_preference: ""
  risk_aversion: ""
  explanation_alignment_schema: []
interests:
  primary: []
  secondary: []
  disinterests: []
goals:
  long_term: []
  short_term: []
constraints:
  time_per_session_minutes: null
  budget_monthly: null
favorites:
  topics: []
  authors_sources: []
  tools: []

privacy:
  data_collection_consent: false
  analytics_consent: false
  personalized_models_consent: false
  research_consent: false
  share_with_partners_consent: false
  content_retention_days: null
  export_requested_at: ""
  delete_requested_at: ""
  legal_basis: []

security:
  mfa_enabled: false
  last_password_change_at: ""
  recovery_emails: []
  sessions_active_count: 0
  last_suspicious_activity_at: ""
  risk_score: 0.0
  blocked_ips: []

devices: []

network:
  preferred_region: ""
  last_ip: ""
  vpn_detected: false

model_access:
  allowed_models: []
  default_model: ""
  feature_flags: {}
  rate_limits:
    requests_per_minute: null
    tokens_per_minute: null
  safety_overrides_user:

```

```

disallow_sensitive_topics: []
max_output_length_tokens: null

usage:
  account_age_days: 0
  sessions_total: 0
  messages_total: 0
  avg_messages_per_session: 0.0
  avg_response_length_tokens: 0.0
  active_days_last_30: 0
  active_weeks_last_12: 0
  most_active_local_hours: []
  topic_counts: {}
  tool_usage_counts: {}
  satisfaction_signals:
    thumbs_up: 0
    thumbs_down: 0
    comment_feedbacks: 0
  last_active_at: ""

integrations:
  files:
    count: 0
    last_uploaded_at: ""
  connected_apps: []

automations:
  count_active: 0
  items: []

conversations: []

prompts: []

bookmarks: []

feedback:
  sessions_reviewed: 0
  items: []

exploration_profile:
  trust_calibration:
    baseline_trust: ""
    explanation_importance: ""
  preference_dimensions: []
  exploration_goals: []
  diversity_target: ""
  familiarity_target: ""
  longlist_size_preference: null
  shortlist_size_preference: null
  feedback_style: {}
  history_signals: {}

moderation:
  user_mutes_topics: []
  assistant_refusals_count: 0
  interventions: []

billing:
  currency: ""
  default_payment_method: ""
  invoices_count: 0
  last_invoice_at: ""
  credits_balance: 0
  promo_codes_applied: []

experiments:
  assignments: []
  metrics_snapshot: []

web_usage:
  allowed: false
  last_used_at: ""
  domains_frequently_cited: []
  citations_total: 0
  pdf_screenshots_taken: 0

files:

```

```

total_count: 0
storage_bytes: 0
items: []

memories:
  enabled: false
  items: []

notifications:
  channels: {}
  digests:
    frequency: ""
    last_sent_at: ""
  rules: []

support:
  tickets_count: 0
  last_ticket_at: ""
  tickets: []

retention:
  delete_after_days: {}
  next_purge_at: ""
  exceptions: []

builder:
  api_keys_count: 0
  apps: []

```

## B Populating Persona Data Templates

### Prompt to ChatGPT 5:

Create 8 realistic, diverse, and fully-populated simulated ChatGPT user profiles in YAML format based on the schema provided below. Do not reference or infer any OpenAI internal data structures — use the schema purely as a fictional simulation template.

Input: [*Persona YAML template from Listing 1*]

### Example System Output:

Listing 2 shows an example of a fully populated simulated user profile generated using the persona schema defined in Listing 1.

#### Listing 2: Example Simulated User Persona (YAML)

```

user_id: "u001"
tenant_id: "us_west_01"
created_at: "2024-01-10T09:00:00Z"
updated_at: "2025-08-25T12:00:00Z"
status: "active"
labels: ["casual_user"]

account:
  email: "karen.fl@example.com"
  email_verified: true
  name:
    given_name: "Karen"
    family_name: "Flowers"
    display_name: "KarenF"
  username: "karen_65"
  plan: "free"
  plan_started_at: "2024-01-10T09:00:00Z"
  plan_expires_at: ""
  locale: "en-US"
  timezone: "America/New_York"
  marketing_opt_in: false
  early_access_opt_in: false
  referral:
    code: ""
    source: ""

profile:

```

<pre> demographics:   age: 65   birth_year: 1960   gender: "female"   pronouns: ["she/her"]   country: "US"   region: "Florida"   city: "Tampa"   language_primary: "en"   languages_other: [] roles_personas: ["retiree"] biography_short: "Retired teacher, now gardening and writing poetry." biography_long: "Karen taught English for 30+ years. In retirement, she enjoys gardening, reading mystery novels, and writing poems for her community group. She uses ChatGPT for recipes, summaries, and learning new hobbies." education:   highest_level: "Master's"   fields_of_study: ["Education", "Literature"]   institutions: ["University of South Florida"] employment:   current_title: "Retired"   current_org: ""   industry: ""   seniority: ""   history: ["High school teacher"] preferences:   writing_style: ["simple", "clear"]   tone: ["warm", "friendly"]   formatting:     bullet_density: "low"     markdown_ok: false     headings_ok: false     citations_style: ""     spelling_variant: "US"   depth_defaults:     explanation_depth: "low"     examples_preferred: true     code_included_by_default: false   content_filters_user:     avoid_topics: ["politics", "violence"]     avoid_phrases: []   accessibility:     reading_ease: "high"     visual_support: []     screen_reader_mode: false   session_behavior:     ask_followups_automatically: true     summarize_long_responses: true     include_tl_dr: false   exploration_preferences:     novelty_tolerance: "medium"     serendipity_preference: "high"     risk_aversion: "medium"     explanation_alignment_schema: []  interests:   primary: ["gardening", "poetry", "literature"]   secondary: ["recipes", "travel"]   disinterests: ["gaming"]  goals:   long_term: ["publish a small poetry collection"]   short_term: ["learn how to compost effectively"]  constraints:   time_per_session_minutes: 20   budget_monthly: 0  privacy:   data_collection_consent: false   analytics_consent: false   personalized_models_consent: false   research_consent: false   share_with_partners_consent: false   content_retention_days: null   export_requested_at: "" </pre>	<pre> delete_requested_at: "" legal_basis: []  security:   mfa_enabled: false   last_password_change_at: ""   recovery_emails: []   sessions_active_count: 1   last_suspicious_activity_at: ""   risk_score: 0.1   blocked_ips: []  devices: []  network:   preferred_region: "us-east"   last_ip: "98.201.45.33"   vpn_detected: false  model_access:   allowed_models: ["gpt-5"]   default_model: "gpt-5"   feature_flags: {}   rate_limits:     requests_per_minute: 30     tokens_per_minute: 20000   safety_overrides_user:     disallow_sensitive_topics: []     max_output_length_tokens: null  usage:   account_age_days: 600   sessions_total: 120   messages_total: 800   avg_messages_per_session: 6.7   avg_response_length_tokens: 250.0   active_days_last_30: 10   active_weeks_last_12: 7   most_active_local_hours: [9, 20]   topic_counts: {"gardening": 60, "literature": 30}   tool_usage_counts: {}   satisfaction_signals:     thumbs_up: 45     thumbs_down: 2     comment_feedbacks: 1   last_active_at: "2025-08-25T12:00:00Z"  integrations:   files:     count: 2     last_uploaded_at: "2025-07-10T10:00:00Z"   connected_apps: []  automations:   count_active: 0   items: []  conversations: [] prompts: [] bookmarks: []  feedback:   sessions_reviewed: 2   items: []  exploration_profile:   trust_calibration:     baseline_trust: "medium"     explanation_importance: "medium"   preference_dimensions: []   exploration_goals: []   diversity_target: ""   familiarity_target: ""   longlist_size_preference: null   shortlist_size_preference: null   feedback_style: {}   history_signals: {} </pre>
---	---

```

moderation:
  user_mutes_topics: ["politics"]
  assistant_refusals_count: 0
  interventions: []

billing:
  currency: "USD"
  default_payment_method: ""
  invoices_count: 0
  last_invoice_at: ""
  credits_balance: 0
  promo_codes_applied: []

experiments:
  assignments: []
  metrics_snapshot: []

web_usage:
  allowed: true
  last_used_at: "2025-08-25T12:00:00Z"
  domains_frequently_cited: []
  citations_total: 5
  pdf_screenshots_taken: 0

files:
  total_count: 2
  storage_bytes: 10240
  items: []

memories:
  enabled: false
  items: []

notifications:
  channels: {}
  digests:
    frequency: ""
    last_sent_at: ""
  rules: []

support:
  tickets_count: 0
  last_ticket_at: ""
  tickets: []

retention:
  delete_after_days: {}
  next_purge_at: ""
  exceptions: []

builder:
  api_keys_count: 0
  apps: []

```

## C Generating Human-Readable Versions of Persona Profile

### Prompt to ChatGPT 5:

Write a 250–300 word natural, human-like character description based on the YAML user profile below. Turn the structured data into a realistic narrative that describes who the person is, what they do, what they care about, and how they use ChatGPT. Keep it warm, believable, and cohesive—no lists or YAML, just flowing prose.

Input: [Example simulated user persona YAML from Listing 2]

### Example System Output:

**KAREN**

Karen Flowers is a 65-year-old retiree living in Tampa, Florida. After teaching English in high school classrooms for more than 30 years, she's enjoying a quieter rhythm of life filled with gardening, reading, and writing poetry. Her days often begin in the garden and end with a mystery novel or a notebook full of verses she shares with her community group. A life-long learner, Karen uses ChatGPT as a practical companion—whether she's looking up new recipes, finding summaries of books, or exploring fresh hobbies like composting.

Her background in education and literature shapes how she approaches the world: she values clarity, warmth, and simplicity in communication, and she avoids topics like politics or violence that don't fit her current focus on gentler, more creative pursuits. She loves poetry that reflects on everyday life and nature, often drawing inspiration from authors like Mary Oliver. While money isn't something she wants to spend on digital tools, she devotes small windows of time—about 20 minutes here and there—to learning, writing, or experimenting with ideas.

Karen's long-term dream is to publish a small poetry collection, something that feels like a culmination of her love for words and her years of guiding others through literature. In the meantime, she sets tangible, everyday goals, like learning to compost more effectively so her garden can thrive. She isn't chasing risk or novelty for its own sake, but she appreciates moments of surprise and delight—whether it's a new flower blooming in her backyard or a line of verse that suddenly feels just right. Karen is warm, thoughtful, and steady, with a deep love for cultivating both her garden and her words.

## D Generating Baseline Condition (i.e., Non-structured) Explanations

### Prompt to ChatGPT 5:

In 100 words explain, in second person, why [PERSONA] should read [RECOMMENDED BOOK]. Here's some information on [PERSONA]: [Example simulated user persona YAML from Listing 2].

### Example System Output:

Karen, you should read *The Book of Delights* by Ross Gay because it feels like sitting with a kind friend who reminds you how much wonder still lives in the ordinary. Each short essay captures a small moment—a blooming flower, a stranger's kindness, a shared laugh—and turns it into something luminous. As someone who loves poetry and gardening, you'll recognize the same gentle noticing that comes from tending both words and plants. Gay's reflections invite you to slow down and savor daily joys, which feels especially comforting and grounding in this stage of life. It's a quietly nourishing read.

## E Generating Experimental Condition (i.e., Structured) Explanations

### Prompt to ChatGPT 5:

In 100 words explain, in second person, why [PERSONA] should read [RECOMMENDED BOOK].

Write as one continuous paragraph that integrates these ideas naturally, without listing or labeling them.

#### Essence – what the item is

Item Metadata (genre/creator/release), Item Subject Matter (core theme/topic), Item Stylistic Expression (tone/aesthetic/voice), Item Format & Form Factor (format/medium/length).

#### Experience – how it fits the user now

Alignment with Time Capacity (time/pacing fit), Alignment with Effort/Skill Capacity (effort/skill fit), Alignment with Affective/Sensory Capacity (mood/sensory fit), Alignment with Values/Identity (values/self fit).

#### Exchange – how it is situated socially and culturally

Personal & Direct Recommendations (e.g., named peers/mentors), Community & Shared-Interest Signals (e.g., named groups/audiences), Public & Cultural Signals (e.g., awards/lists/media).

### Entwinement – how it connects further

Relationships to Similar Items of the Same Kind (back/forward within medium), Connections to Different Kinds of Items (cross-media resonance).

Here's some information on [PERSONA]: [Example simulated user persona YAML from Listing 2].

### Example System Output:

Karen, you should read *The Book of Delights* by Ross Gay because it gathers 102 of his brief daily essays that uncover joy in small, grounded moments—like planting flowers, savoring fruit, or exchanging kindness with strangers. His tone is warm, funny, and deeply observant, echoing the mindfulness you bring to your garden and poetry. The short pieces fit easily into quiet Florida mornings or restful evenings, offering gentle inspiration without hurry. Many readers and writers admire how Gay turns attention into gratitude. When you finish, you might enjoy his public readings on PBS or NPR, where his voice makes those delights sing.