

ChoiceMates: Supporting Unfamiliar Online Decision-Making with Multi-Agent Conversational Interactions

Jeongeon Park*

University of California, San Diego
La Jolla, California, USA
jep034@ucsd.edu

Bryan Min

University of California, San Diego
La Jolla, California, USA
badmin@ucsd.edu

Kihoon Son

School of Computing
KAIST
Daejeon, Republic of Korea
kihoon.son@kaist.ac.kr

Jean Y Song

Humanities, Arts, and Social Sciences,
Underwood International College
Yonsei University
Incheon, Republic of Korea
jeansong@yonsei.ac.kr

Xiaojuan Ma

Hong Kong University of Science and
Technology
Hong Kong, China
mxj@cse.ust.hk

Juho Kim

School of Computing
KAIST
Daejeon, Republic of Korea
juhokim@kaist.ac.kr

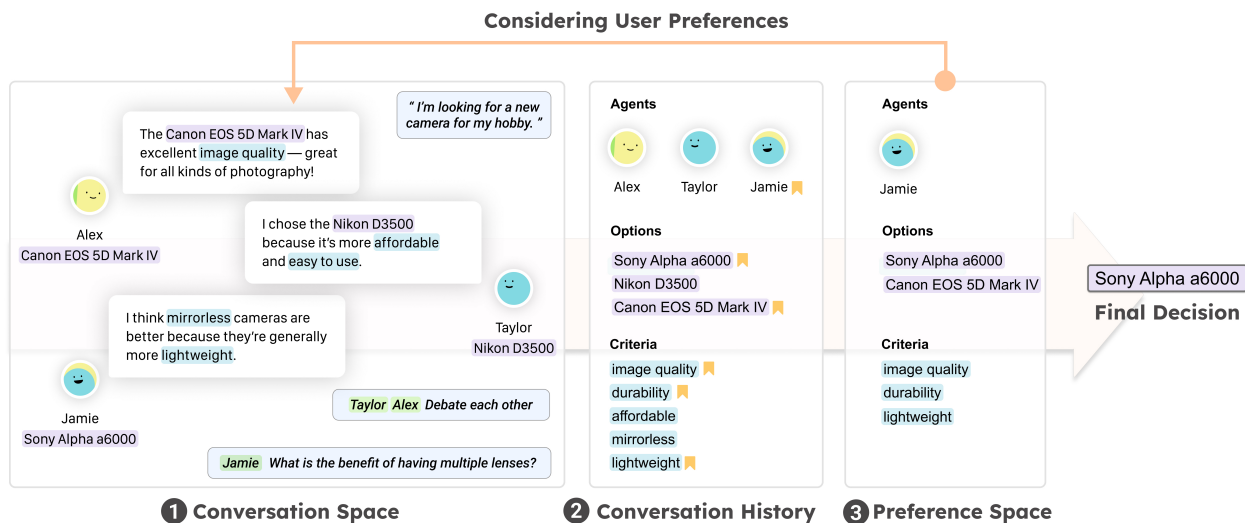


Figure 1: ChoiceMates is an interactive multi-agent conversational system designed to support online decision-making in unfamiliar domains. In ChoiceMates, the user can converse with any selected set of agents in the conversation space (1) to gather diverse information and perspectives in the domain (supported by features that auto-highlight criteria and options in the conversation history) (2) to manage and understand the information found. The user can pin important keywords to the preference space (3) to build their own preference in the process and make an informed and confident final decision.

Abstract

From purchasing a gift to deciding on a hobby, unfamiliar decisions—decisions without domain knowledge and experience—are frequent and significant. The complexity and uncertainty of such decisions demand unique approaches to information seeking, understanding, and decision-making. Our formative study highlights that in the

current workflow, users want to start by discovering broad and relevant domain information evenly and simultaneously, quickly address emerging inquiries, and gain personalized standards to assess information found. We present ChoiceMates, an interactive multi-agent system designed to address these needs by enabling users to engage with a dynamic set of LLM agents each presenting a unique experience in the domain. Unlike existing multi-agent systems that automate tasks with agents, the user orchestrates agents to assist their decision-making process in each turn, through chatting with all agents, with a tagged subset of agents, or calling in new agents into the space. By comparing ChoiceMates with a web search condition and a multi-agent framework (n=12), we show

*This work was mostly conducted while Jeongeon was at KAIST.



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/10.1145/3742413.3789107>

that ChoiceMates enables a more confident, satisfactory decision-making with better situation understanding than web search, and higher decision quality than a commercial multi-agent framework. We further illustrate how participants utilized ChoiceMates to make unfamiliar decisions, providing insights into designing a more controllable and collaborative multi-agent system.

CCS Concepts

• **Human-centered computing** → **Interactive systems and tools**; *Natural language interfaces*.

Keywords

multi-agent interactions, conversational user interface, decision-making support, large-language models

ACM Reference Format:

Jeongeon Park, Bryan Min, Kihoon Son, Jean Y Song, Xiaojuan Ma, and Juho Kim. 2026. ChoiceMates: Supporting Unfamiliar Online Decision-Making with Multi-Agent Conversational Interactions. In *31st International Conference on Intelligent User Interfaces (IUI '26), March 23–26, 2026, Paphos, Cyprus*. ACM, New York, NY, USA, 25 pages. <https://doi.org/10.1145/3742413.3789107>

1 Introduction

People make decisions online every day, from purchasing a new product, choosing a trip destination, to taking on a new hobby. In making a decision, people inherently go through a set of steps including recognizing the need for decision-making, collecting information, identifying and ranking key criteria, identifying and evaluating alternatives, and deciding the final choice among them [58]. However, making an *unfamiliar* decision—decision-making situations where an individual lacks knowledge and preference in the domain—requires individuals to first grasp the fundamental knowledge of the domain before delving into the decision-making process [3, 34]. Such lack of familiarity makes decision-making more challenging as it requires the user to continually switch between learning about the domain and evaluating options, often under a limited time budget [4, 60, 61].

While prior work has demonstrated effective support for online decision-making, they have focused on providing support for managing and organizing criteria and alternatives in the domain rather than for uncovering and understanding them [8, 50, 51]. This lack of support during the initial phases of searching for and understanding the key domain information can increase information overload for unfamiliar decision-making situations [26, 73]. Subsequently, information overload often leads to a significant investment of time and cognitive resources in grasping domain details and eliciting preferences, potentially causing individuals to make uninformed decisions or even give up entirely [28, 77]. Beyond increasing effort, overload can also undermine users' confidence in their judgment and satisfaction with the resulting choice, especially when users must make sense of a new domain [13, 40, 76].

On the other side, other research aids sensemaking of unfamiliar domains through an interactive and guided exploration of comprehensive knowledge [52]. While such an approach ensures a holistic comprehension of the domain, users still need to put a significant effort into connecting the knowledge to their situations, especially for new decisions. To minimize user effort, recommender systems

aim to simplify decision-making by suggesting personalized options [9, 23]. Yet, for novices in a domain, these preferences might not be well-elicited [55, 72]. This can cause an over-reliance on the system, especially if the novices cannot detect errors due to their limited knowledge [62]. Ultimately, novices face the dilemma of researching the unfamiliar domain individually which is demanding, evaluating different options without good confidence and intuition, or ending up making a decision that is ill-informed, all of which motivates the need for a new approach that supports unfamiliar decision-making.

To understand the practice and needs of online decision-making in an unfamiliar domain, we conducted a formative study with 14 participants who had experience making decisions online. From observing and interviewing their decision-making process in unfamiliar domains, we discovered that participants described the need to understand a wide range of opinions and perspectives (including relevant and broad domain information) simultaneously, quickly explore inquiries that are on top of their minds, and develop personalized standards to access the information found.

The identified needs point to a challenge that is inherently about comparing perspectives: novices often need to quickly sample and compare a wide range of viewpoints to learn what matters in a domain, while also asking follow-up questions as new uncertainties arise. This motivated us to explore **multi-agent conversational interactions**. Multiple agents can surface diverse opinions and perspectives in the domain through distinct viewpoints, allowing users to juxtapose opinions, discover trade-offs, and surface relevant criteria early rather than assembling information from scattered sources sequentially. Conversational interaction can support lightweight exchanges of information needs: users can ask follow-ups, compare options, and focus the search on their context. Together, a conversational set of agents could help users learn key domain knowledge, identify plausible options, and develop personalized standards within a limited time budget.

Many existing multi-agent systems primarily aim to automate the process of solving complex tasks with minimal human intervention by assigning agents into different subtasks [43, 88] or roles [25, 81, 81]. However, there is a limited exploration of multi-agent interfaces with interactions where the user is assisted or empowered by collaborating with a set of agents. For tasks like information seeking and (unfamiliar) decision-making where human preferences and opinions play a significant role, it is important to keep users in the loop, as full automation may lead the users to naively follow agents' guidance and lose agency over their decisions.

We present **ChoiceMates**, a multi-agent conversational system where the user can navigate the unfamiliar decision-making process by conversing with a set of agents and constructing the user's own preference space. In ChoiceMates, each agent is characterized by a unique descriptor (i.e., a single-line description of the agent; a professional photographer in Fig. 3), a set of valued criteria (quality, durability, and battery in Fig. 3), and a single option (Canon EOS 5D Mark IV in Fig. 3) to reflect the unique experience, to support intuitive exploration of diverse perspectives. The user takes full control of the agents, and can converse with any individual or group of agents in the *conversation space* (Fig.1-1) to explore and expand the information in the decision space. Agents, auto-detected criteria, and options from the conversation appear as the *conversation*

history (Fig.1-2) to help users uncover and keep track of the key information found. The user can save agents, criteria, and options to the preference space (Fig.1-3) with brief notes throughout the conversation to shape the user’s preferences and utilize them when making the final decision.

To evaluate whether an interactive multi-agent-based approach can effectively support unfamiliar decision-making, we conducted a three-condition within-subjects user study with 12 participants. We compared ChoiceMates with a web search condition and a commercial multi-agent framework condition¹ where a set of necessary agents collaborate to respond to the user, guided by the following research questions:

- RQ1.** How does ChoiceMates help users explore a broad information space in the domain?
- RQ2.** How does ChoiceMates support the discovery and management of relevant information to the user’s context?
- RQ3.** How does the user perceive their final decision with ChoiceMates?
- RQ4.** How does the user utilize ChoiceMates in the decision-making process?

Our results show that ChoiceMates successfully led to a more satisfactory and confident decision, enables comparable exploration of a broad information space, and supports easier management of discovered information and a better understanding of user situation than the web condition. Compared with an existing computer-led multi-agent framework, ChoiceMates facilitated broader exploration of information in the domain and participants perceived the final decision as of higher quality. In addition, we report on representative strategies of how participants utilized ChoiceMates for unfamiliar decision-making, and further discuss the user-side implications and design considerations of a controllable, collaborative multi-agent system.

This research makes the following contributions:

- Findings from a formative study with 14 participants uncovering the practices and needs in unfamiliar decision-making.
- Design and implementation of ChoiceMates², a multi-agent conversational system designed to support unfamiliar decision-making by providing users control over conversations with a dynamic set of agents with unique experiences and preference construction in the form of criteria, options, and agents.
- Findings from a study with 12 participants that revealed ChoiceMates supports discovery of broad information with less burden, effective management of discovered information, and a more satisfactory and confident final decision compared to the Web condition, and a wider exploration of information space and a perceived quality final decision with higher confidence compared to the MultiAgent condition.

2 Related Work

Our work addresses the challenges in online *unfamiliar* decision-making by leveraging the benefits of multi-agent conversations powered by LLMs. We review prior work on online decision-making

support, conversational interfaces for online decision-making, and designs of collaborative multi-agent systems.

2.1 Online Decision-Making Support

Due to the enormous amount of information online, decision-makers can face information overload [79], which negatively influences the resulting decision quality [40, 67]. For unfamiliar decisions where users do not have sufficient knowledge or expertise in the domain, searching for information becomes particularly challenging [3, 34]. Beyond increasing effort, overload can undermine users’ confidence and satisfaction with the resulting choice, particularly when they must identify what criteria matter and evaluate options in a new domain [13, 40, 76].

One thread of research focused on identifying selection criteria among a wide range of options to choose from. Mesh proposed consumer product decision-making through comparison tables and customizable preference settings [8], and Unakite and Crystalline supported software developers’ decision-making with web content snippets and comparison tables [50, 51]. While these systems provide strong support for utilizing collected information on options to make decisions, they provide limited assistance in helping users understand information when users lack domain knowledge. Consequently, novices in the domain may not only struggle to identify important criteria but also to critically compare discovered options [26, 73]. On the other hand, Selenite leveraged LLMs to generate potentially relevant criteria based on user search and annotate them in the information inside websites to structure the exploration process [52]. While Selenite ensures a high comprehension of the domain, the users would need to spend additional effort to precisely understand how the criteria match their unique situation to make a decision.

To make the process less overwhelming, recommender systems can guide users unfamiliar with the domain by providing personalized options for users to choose from [9, 23, 71]. However, one-shot recommender systems (i.e. recommendations given after just a single example or interaction) are insufficient to support users in learning about the domain for unfamiliar decisions as the interaction is uni-directional [30]. This can decelerate users from forming their preferences during the decision-making process [55, 72, 83]. Conversational recommender systems (CRS) build on top of existing recommender systems by offering the ability to make back-and-forth exchanges of messages to continually provide contextualized and relevant recommendations throughout the user’s learning and decision-making process [30, 82, 84]. However, CRS often struggle with maintaining context and providing sufficient explanations that the user need [29].

ChoiceMates is primarily inspired by the benefits of CRS in providing personalized recommendations and guidance for preference elicitation and aims to expand that to unfamiliar decision-making situations where a holistic understanding of a domain is crucial for making a confident decision. We additionally incorporate the benefits of web-based decision-making support systems to identify and manage key criteria and options in the domain better.

¹Multiagent Wizard: <https://chat.openai.com/g/g-u9C6YeMsL-multiagent-wizard>

²The code repository and supplementary materials for ChoiceMates is available at <https://github.com/jeongeonp/ChoiceMates>.

2.2 Text-Based Conversational Interfaces for Information Seeking

Text-based conversational interfaces have been widely adopted for various tasks, valued for their ability to present information in a human-like, accessible format and support flexible, adaptive exploration [17, 48, 86, 87]. They have also proven effective in information seeking behaviors online, where Searchbuddies [24] embedded search engine agents in social media threads to provide easy access to relevant information online, and Gupta et al. [21] increased user trust and satisfaction through conversational online housing recommendations. In addition, Radlinski et al. [69] proposed a set of properties that composes a natural and efficient conversational information retrieval system. The recent surge of large language models (LLMs) has revealed the potential for more capable conversational interfaces across diverse domains ranging from programming to searching for UI [49, 59, 80]. One most prominent example for information seeking is ChatGPT and other LLM applications, where they support real-world tasks such as trip planning or learning a new language [63]. However, such interfaces often present information as a linear stream, which can make it difficult to compare options and keep track of evolving criteria, which are critical steps in unfamiliar domains.

On the other hand, while conversational interfaces adaptively provide information through a multi-turn process, conversing requires the user to articulate their input clearly which, if not done, could result in misinterpreted user intent and irrelevant responses [5, 16, 27, 57, 94]. To complement that, existing work has demonstrated techniques to bring multi-modal inputs such as a graphical user interface (GUI) to repair conversational breakdowns [46, 47]. LLM-based interfaces have also become more proactive, asking clarifying questions and offering candidate options when users' intents are underspecified.

ChoiceMates builds on these LLM-powered conversational interfaces by providing interactive and digestible conversations to support online decision-making. In addition, ChoiceMates adds on direct manipulation of objects as an interaction modality on top of natural language, where actions such as saving valuable criteria or options to the preference space are fed into the conversation as a context. This moves beyond a linear conversational interaction by letting users externalize and organize information, where these actions become part of the conversational state, enabling more precise and user-steered preference articulation.

2.3 Multi-Agent Interfaces

Earlier multi-agent works show that many human tasks benefit from collaboration and information integration from various sources [14, 19] compared to individual cognitive processes, and capitalize on unique identities of each agent and their abilities to delegate tasks among agents [35, 36, 45]. Synthesizing such investigations with technical advancements, multi-agent systems have recently gained popularity with the emerging capability of Generative AI and LLM. To automate human work, they utilize the interaction between multiple agents and/or with humans to simulate human behaviors [65, 66], cooperate to solve complex tasks [43, 81, 93], generate visual contents [6, 15, 88], and leverage perspectives from diverse agents [7, 12]. Frameworks such as AutoGen

[85] or AgentVerse [10] support this process of building a multi-agent system by helping developers flexibly define and configure agents and conversation patterns.

Distinguishing from multi-agent systems that aim to automate human work with minimal human intervention (e.g., adjusting input data), ChoiceMates aims to design a system that can foster collaboration between multiple agents and humans. Prior work explores effective designs of agents or multi-agent conversations to assist human judgment and task execution. For example, Benharrak et al. [2] asked writers to define and identify AI personas that can provide on-demand feedback from different perspectives. ChainBuddy [89] assisted the planning process by assigning specific agents to each task in the plan, allowing for a focused execution. Moreover, CommunityBots [33] evaluated a conversation and topic management mechanism with multiple agents each specializing in a topic for gathering public input across domains and topics with a Wizard of Oz study. While they provide interesting insights into multi-agent interaction design, they offer limited direct evidence on how such designs change users' decision-making processes.

Complementing these, recent HCI work examined how multi-agent conversational interfaces can broaden the set of perspectives users encounter (e.g., to counter filter bubbles) [92], increase intended sustainable behaviors through guidance from multiple NPCs [90], and how multi-agent group dynamics may influence users' judgments [75]. Building on this growing literature, we investigate the design and implementation of multi-agent interfaces in the context of information seeking and comprehension for unfamiliar decision-making processes, an underexplored domain. Specifically, we design agents that not only provide factual information but also surface diverse perspectives by revealing similar and contrasting experiences and opinions in response to user inquiries, aiming to deepen users' and support more informed decision-making through enriched interactions.

3 Formative Study

To design an approach to support unfamiliar decision-making, we first aimed to understand the common practices in unfamiliar decision-making and uncover the needs in the process. For this purpose, we conducted semi-structured interviews with 14 participants with prior experience in unfamiliar decision-making.

3.1 Participants and Study Procedure

We recruited 14 participants (Age=18-55, M=28.2, Std=7.8; 8 males and 6 females) who had frequent experience making decisions online (details in Table 1). The recruitment took place through our university's community channels and snowball sampling to recruit participants in a wide age group. To ensure their experience in making decisions in unfamiliar scenarios, we asked the participants to list their previous experience in both familiar and unfamiliar decisions in the recruitment form.

The interview lasted for 75 minutes for each participant. First, a 20-minute semi-structured interview was conducted to understand participants' previous experiences with unfamiliar decision-making. We asked questions about the overall process, which information they utilized, and the challenges and needs they faced in the process. Then, a 30-minute think-aloud study was conducted where

participants chose an unfamiliar scenario from the given set (e.g., buying a robot vacuum cleaner or choosing a new hobby to do in their free time; see Appendix A) and browsed through the internet until they identified several solid options.³ Lastly, participants were interviewed for 15 minutes on their experience with their process and the challenges and needs they faced in this scenario, then on the general challenges and desired support of unfamiliar decision-making. Participants were compensated 20,000 KRW (approximately 15 USD) for the interview. The study was approved by the Institutional Review Board (IRB) at the institution this research was conducted.

After the study, the interviews and think-aloud sessions were transcribed using automated tools and manually reviewed for accuracy by the first author. The two authors then conducted an inductive thematic analysis, independently coding participants' current practices, existing needs, and desired support. The authors collaboratively discussed discrepancies, resolved conflicts through consensus, and iteratively refined the codes to finalize the findings.

3.2 Findings

Here, we describe the different approaches people follow for unfamiliar decision-making and describe the three main needs we identified from the process.

3.2.1 Practice. The participants utilized several sources in an unstructured manner (i.e., in various orders and frequencies) to collect information and make unfamiliar decisions. Such sources include summary posts of the domain (n=8), posts with individual opinions and experience (n=10), and websites with a list of options in the domain (n=11). During the interview, many (n=12) brought up asking an expert as their previous experience to get helpful and needed information for making decisions in unfamiliar domain.

With *summarized content* such as YouTube summarization videos or rating sites (e.g., <https://www.rtings.com/>), participants were able to objectify multiple criteria and understand an overview of the domain, but could not receive information more relevant to their context. Participants found *posts with individual opinions*—online communities, product reviews, or blog review posts—helpful for gathering diverse opinions on the options but also mentioned that it is time-consuming to find credible sources and the possible bias hinders them from utilizing individual opinions solely. *Browsing through a grid of options* from e-commerce websites (e.g., <https://www.amazon.com/>) was another approach participants took, but they were unable to understand and compare different options without knowledge of the domain, thus switching to the other two sources of information.

On the other hand, 12 out of 14 participants mentioned *asking an expert*—including store managers, friends, and family who have expertise in the decision domain—an alternative approach they took for past unfamiliar decisions as a way to receive information without being overwhelmed. They mentioned that experts' ability to (1) explain ground-based (i.g., basic but essential) information about the domain and provide information that's difficult to find, and (2) ask clarification questions to make better suggestions in

one's specific context example makes them reach out to experts for unfamiliar decisions.

3.2.2 Needs. We identified three main needs participants exhibited during the unfamiliar decision-making process.

C1: Discovering broad and relevant domain information evenly and simultaneously. The participants mentioned two major types of information that assist their unfamiliar decision-making process, namely broad (i.e. information that represents comprehensive information of the domain) and contextually relevant (i.e. information that is relevant to one's situation or preferences) domain information. Participants expressed concerns about 'missing out on the core domain information' (P14) but were also frustrated when they were 'not able to find information relevant to their case' (P5). As the domain was unfamiliar, participants often ended up exploring one type of information and struggled to start exploring the other type of information. Participants tried to mitigate this by expanding their discovery to more fresh experiences or opinions in the domain, to ensure that they understood the landscape of the domain. However, due to their lack of ability to discover information in an unfamiliar domain, many could not diversify their search to underexplored types of information.

C2: Quickly addressing inquiries that are on the top of their minds. As participants had insufficient domain knowledge, they had many inquiries and curiosities during the process. While such inquiries were often quick and short (e.g., difficulty of installing a discovered car seat (P9)), participants deliberately looked for answers instead of skipping as it was a crucial step to establishing a deeper domain understanding. However, their inquiries were not sufficiently addressed due to a lack of time or ability to find what they wanted. For example, P12 (interior light) while watching a YouTube video on famous lamp brands, had to go on Google multiple times to search for new information, such as the meaning of Handwerker or mushroom lamp styles, which they later described as tedious. P12 additionally expressed the need to chat with an expert to understand possible materials and their pros and cons. Similarly, P2 (solo trip destination) mentioned that they would want the person who wrote a particular blog post to answer a few extra questions, including suggestions for other trip destinations and why.

C3: Gaining personalized standards to assess the information found. A core concern raised by the participants was that they did not have enough knowledge or confidence in the domain to assess the information discovered. Even after participants read through a few webpages or had notes written down, they still thought they did not have sufficient understanding of their situation to establish personalized standards and evaluate different criteria and options. Thus, participants had to gather a personalized set of key aspects and their importance to assess newly discovered information or make a choice. For instance, P1 (cafe for catering) faced difficulties in ranking the discovered aspects (e.g., location, type of desserts, catering experience) and making the decision until the end when they realized that they were prioritizing cafes with catering experiences.

³As we allocated a rather short time for the think-aloud study, we complemented the observation with an interview about their past experiences.

ID	Age	Gender	Previous Unfamiliar Scenario	Selected Scenario
P1	26-35	F	Choosing a course to take for a summer school, Buying an oven for a baking hobby	Choosing a cafe to cater some snack food on an end-of-semester event
P2	36-45	M	Choosing a trip destination, Choosing a t-shirt for summer	Planning a solo trip destination for three days
P3	26-35	M	Buying a tennis racket, Choosing a U.S. stock option to buy	Buying a car seat for a friend's newborn
P4	26-35	F	Choosing a working holiday location, Buying a laser hair removal	Buying a car seat for a cousin's newborn
P5	18-25	M	Buying a lunch box	Buying a robot vacuum cleaner to replace the normal vacuum cleaner
P6	26-35	F	Buying a new tennis racket, Choosing meal to cook	Renting a house short-term for an internship
P7	18-25	M	Choosing a transportation card plan in a new country	Choosing a new hobby to do in free-time
P8	18-25	F	Planning a trip to Osaka, Choosing a wrist/ankle brace+lumbar support	Buying a robot vacuum cleaner to replace the normal vacuum cleaner
P9	18-25	M	Choosing a hair salon, Buying a gift for a friend	Buying a car seat for a friend's newborn
P10	18-25	M	Whether to get a CT+MRI during ER visit, Buying clothes at a flea market	Buying a skateboard for transportation purposes instead of walking to school
P11	26-35	F	Buying a coffee pod machine, Searching for a job	Buying a skateboard for transportation purposes instead of walking to school
P12	26-35	M	Buying a used phone, Planning a trip to Jeju	Buying an interior light at home
P13	18-25	M	Choosing an affordable phone plan, Choosing field of work	Buying a skateboard for transportation purposes instead of walking to school
P14	46-55	F	Choosing exercise, Buying a blue light blocking film	Buying an interior light at home

Table 1: Participant demographic for the formative study.

4 Design Goals

With the identified needs in the unfamiliar decision-making process, we propose the following design goals for a system that can support unfamiliar decision-making.

- DG1. Provide diverse opinions and perspectives in the domain through experiences.** In unfamiliar decisions, participants wanted to discover a wide range of information, namely broad, relevant, and fresh information with equal focus. Given difficulty identifying what's "core" and answering emerging questions, participants turned to multiple experiences to fill gaps and validate claims. Past research has shown that social information search containing people's experiences could help discover increased firsthand information and diverse perspectives with value judgments [1, 31]. In addition, prior work has shown that novices preferred option-based suggestions of preferences over attribute (i.e., criteria)-based suggestions [37]. Inspired by such works, we can design a system where multiple experiences of using different options can offer complementary perspectives without overwhelming users with unstructured viewpoints.
- DG2. Provide options to explore broad and contextually relevant domain information.** Participants also mentioned the need for simultaneous exploration of broad and relevant

domain information to expand the depth and width of their domain knowledge. Support for users to seamlessly switch between the two types of information when needed in the process may help them fill in underexplored parts of their information discovery in the domain.

- DG3. Allow for direct exchanges of information needs through conversational interactions.** Conversation is an effective medium for an enhanced searching experience [48] and in expressing information goals [74]. To aid users in addressing their unique information needs, an effective approach would be to utilize conversational interaction where users can ask follow-up questions on top of the acquired information. This can support further exploration of how different domain information connects to each other and to the user's preferences.
- DG4. Facilitate users in building preferences relative to the domain.** The participants mentioned the difficulty of gaining personalized standards (e.g., a travel destination less than 3 hours via flight) that would help them assess the information found. Rather than assuming users can form standards from scratch, the system can use exploration history as a *starting point* to surface preference candidates, and have the users confirm or revise them. To reduce the

risk of filter bubbles, the system should also surface contrasting opinions alongside emerging preferences. Identified aspects could then function as evidence for the user to evaluate later discovered information and gradually clarify their preferences over the decision-making period.

5 ChoiceMates: A Multi-Agent Conversational System

From our findings and the design goals, we designed and developed **ChoiceMates**, an interactive multi-agent conversational system that assists users in unfamiliar decision-making—situations where users must make a decision in a domain where they lack prior domain knowledge and well-formed preferences, often under a limited time budget. Note that the scope of unfamiliar decision-making ChoiceMates is designed to support does not include primarily high-stakes or overly complex decision contexts (e.g., those requiring extensive verification, long-term deliberation, or multi-stage planning), which we discuss further in Section 7.3.

ChoiceMates allows the user to (1) converse with a dynamic set of agents to explore broad and relevant information and (2) utilize extracted conversation history to fill up the preferences space to make a decision. Here, we describe an envisioned user scenario, explain different system components, and describe the system implementation and prompt engineering.

5.1 Envisioned Scenario

Sally has recently found an interest in photography and wants to find the right camera for her. However, she is new to the camera domain and has no idea which camera she would be interested in. As she has no experts to consult and gets overwhelmed by doing a few online searches, Sally enters ChoiceMates and types into the message input box, “I’m new to photography, and I want a camera. Not sure which one would be best for me.” She then encounters three distinctive agents (Alex, Jamie, Taylor) appearing on ChoiceMates, each with a unique profile containing an **option** (i.e., a camera model) they had chosen and the **criteria** (e.g., portability, brand) that led to their choice. In each agent’s response, she sees criteria and options auto-detected and listed in the conversation history.

- Alex: “As a professional photographer, I value **image quality** and **durability** in a camera. That’s why I chose the **Canon EOS 5D Mark IV**. It’s a full-frame DSLR that delivers excellent image quality and is built to last.”
- Jamie: “I’m a hobbyist photographer and I prefer a camera that’s **lightweight** and **easy to use**. I’ve been using the **Sony Alpha a6000** and it’s been great for me. It’s compact, takes great photos, and is very user-friendly.”
- Taylor: “I’m a travel blogger and for me, **portability** and **battery life** are key. I use the **Fujifilm X-T3**. It’s compact, has a long battery life, and takes amazing photos. What are your needs and preferences when it comes to photography?”

Sally learns about the wide range of values and criteria agents with different professions and lifestyles have for cameras through the agents and their profiles (**DG1**). Sally considers Jamie to be the

most relatable as she is getting a camera for a hobby. Sally tags Jamie and asks why they value lightweight and easy-to-use cameras. Jamie responds that they “don’t need all the bells and whistles of a professional camera” and they “want to focus on capturing the moment”. Sally learns that for her personal use, she may want to follow Jamie’s reasoning as Sally is also looking for a simple camera. She then asks follow-up questions to Jamie on “What is considered a lightweight camera?” and “What is the main difference between a professional camera and an easy-to-use camera?” to better understand what each criterion means (**DG3**). Sally’s conversations with Jamie allow her to discover what is relevant to her context.

After gaining some information, she decides to get an easy-to-use camera and saves the criterion **easy-to-use** and the agent Jamie to her *preference space*. Then, she turns on the *preference toggle* as she now wants more relevant information to her identified preferences (**DG2**). As the agents have access to Sally’s preference space, when she asks a question that reflects her preference: “Are there other cameras that I would like?”, three more agents with the criterion **easy to use** included in their valued criteria appear: Riley, Morgan, and Casey. Continuing the conversations with different agents, Sally is suddenly lost on which information she needs to discover more. She takes a look at the conversation history, sees the criterion **durable** mentioned the most, and realizes that she may consider durability a key aspect of her preference. She then adds **durable** to the preference space and continues exploring the options with relation to their durability (**DG4**).

As Sally continues to converse and learn more, she becomes more and more confident about the set of criteria she values for her activities and adds those to her preference space (**DG4**). Now that she has a clearer picture of the domain, she also wants to compare the three cameras she is considering in-depth. By asking the agents to ‘debate each other’, Sally understands the crucial differences between the options and concludes that **Nikon Coolpix B500** best fits her preferences and adds it to her preference space. She checks her *preference space* and the saved criteria and options one more time and confidently makes the decision.

5.2 System Components

ChoiceMates (Fig. 2) is an interface that allows the user to interact with a set of agents to make an unfamiliar decision. Agents are the main unit of conversation in ChoiceMates and are the basis of gaining information about the domain.

ChoiceMates consists of a *conversation space* where the user can converse with a dynamic set of agents, a *conversation history* that automatically lists agents, criteria, options discovered during the conversation, and a *preference space* where the user can gather and save relevant information throughout the process. In this section, we first describe agents, the conversation space with the agents and additional support in the conversation space to manage the information, then the conversation history and the preference space.

5.2.1 Agent: Basic unit of information. In ChoiceMates, each agent (Fig. 3) is characterized by its descriptor (i.e., a single-line description of the agent; a beginner photographer who likes natural scenes), a set of criteria (i.e., factors to consider in the domain) they value, and a single option (i.e., an available choice in the domain) they

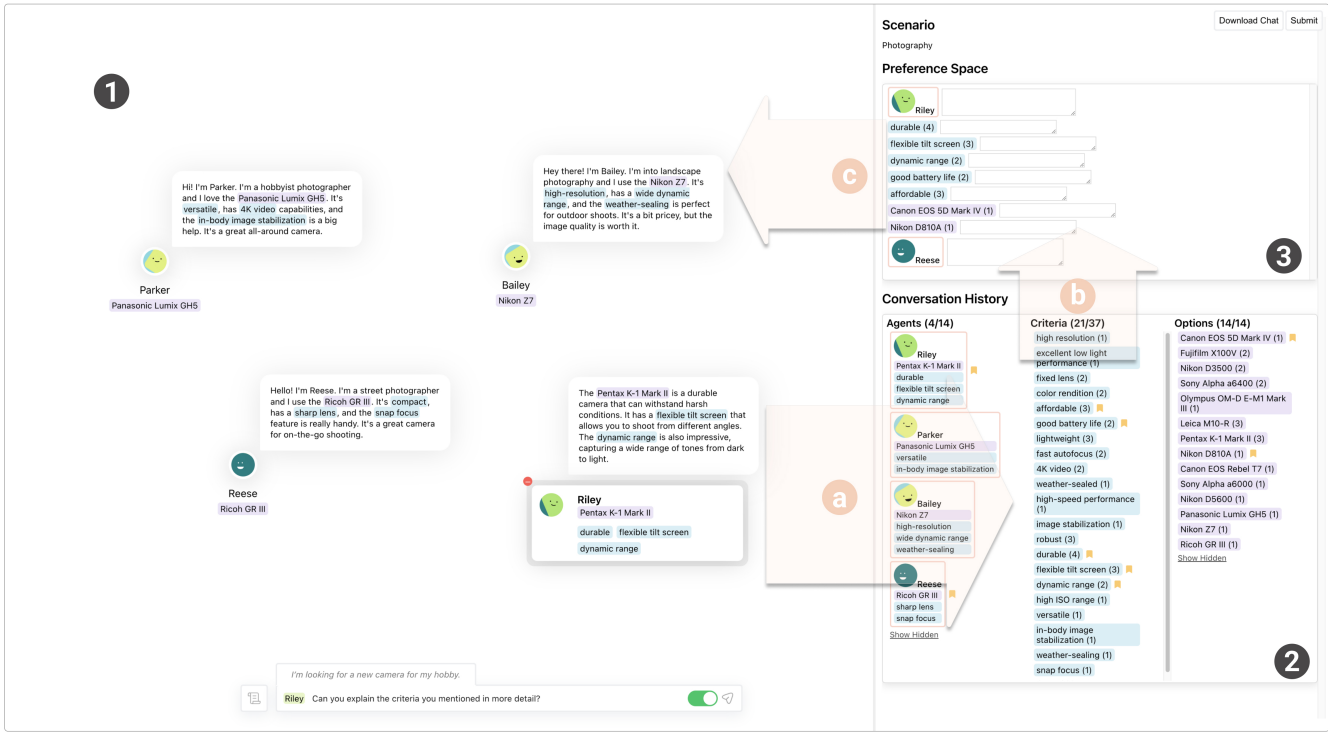


Figure 2: The ChoiceMates interface: (1) Agents populate the conversation space to converse with the user, identifying key criteria and options in their utterances (1). The Conversation History automatically lists those criteria and options mentioned by agents (2). Preference Space contains the options and criteria selected by the user from the Conversation History. (a) indicates the unfamiliar information stacked in the Conversation History. (b) shows the saving process of the user preference. Lastly, (c) illustrates the flow of applying user preference into the Conversation Space with the agents in ChoiceMates.

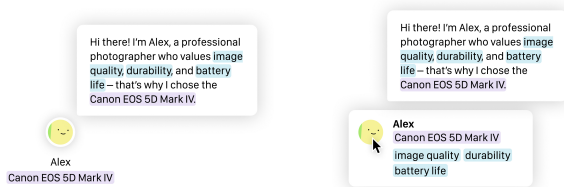


Figure 3: (1) An agent utterance is represented as a chat bubble above the agent icon. (2) Hovering over the icon reveals the agent's profile containing their valued criteria and their chosen option.

chose with the criteria. The agents in ChoiceMates are designed to each reflect an individual's choice and their underlying values and experience in real life, inspired by previous work showing that novices preferred case-based preference elicitation in a recommender system [37]. Multiple agents as a group provide users with a broad range of experiences in the space (DG1). To provide a better salience of criteria and options in the domain, all criteria mentioned by agents are highlighted in blue and all options are highlighted in purple.

When the user asks an initial question in the domain, ChoiceMates prompts the LLM to generate a set of diverse agents, between 3 and 6 to provide a good number of different agents without choice overload [20]. Newly spawned agents always begin with introducing themselves, sharing their valued criteria and a chosen option (Messages in Fig. 3). To ensure that the agents communicate the correct information, we use a simplified RAG framework (details in 5.3), by scraping factual information of each option. The scraping happens *right after* each agent is generated to let the LLM produce diverse agents without being restricted by the web-searched information. The scraped information is provided to the agents and remains in the conversation stream with a prompt asking agents to utilize the information as context. The user can access the original source of the provided information via the hyperlink added to the option and validate the information throughout the session.

5.2.2 Conversation Space. To support easy exploration of broad and personal information in an unfamiliar domain, ChoiceMates provides a conversational space where all agents reside, where users can converse with a select set of agents (Fig. 4-2).

On ChoiceMates, agents are designed to chat back and forth with the user (DG3). This way, users can ask follow-up questions to agents they relate to and gain a deeper understanding of domain options. The user initiates the conversation by sharing their

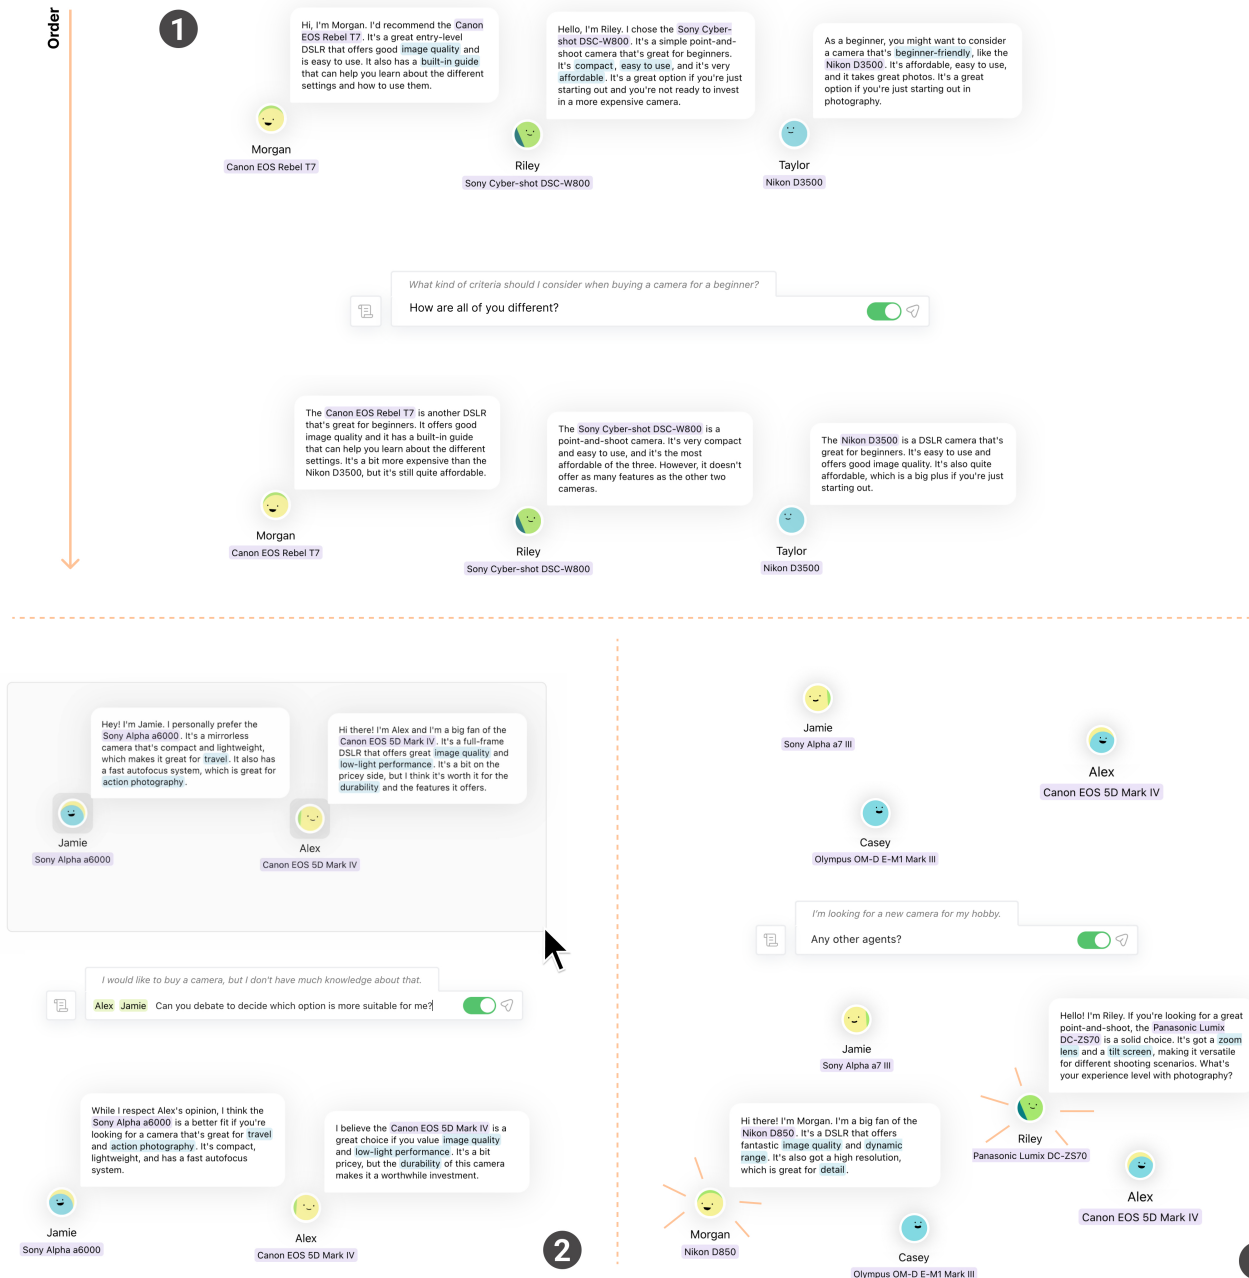


Figure 4: Three types of conversation interactions possible in the conversation space. (1) The user can chat with any agents currently in the space, (2) select or tag one or more agents to chat, and (3) call new agents into the conversation space with their input. Like the conversation in (2), the agents will build on top of each other's responses when applicable.

decision-making scenario, and three to six agents with varying profiles appear on the screen. Agents then share their experiences and ask questions about the user's preferences. The user can respond to agents through the message input box by either replying to a question or asking one themselves.

ChoiceMates offers both user-driven and system-driven options for selecting agents to respond. Upon receiving the user's message, ChoiceMates detects the user's intent and makes existing agents, new agents, or a combination of both to respond, depending on the agents' relevancy to the user's message and context (More

described in Fig. 6) If the user wants to manually assign particular agents to respond, they can also select agents in the space (Fig. 4-2). Additionally, agents can converse with one another by responding to other agents in the conversation. They can agree, disagree, or ask further questions to other agents (Fig. 4-3). This can reveal comparisons between options and new criteria that the user may not have thought of when they are the only ones conversing with them. Just as the user can manually select agents to respond, the user can also manually trigger inter-agent conversations by selecting agents and asking them to “debate each other”.

To reduce clutter from more agents and messages, we chose to only display the latest messages sent from the previous turn. Instead, users could click the button left of the message input box in Fig. 2 to view the full conversation history as a linear thread.

5.2.3 Conversation History. To support a more effective exploration and management of information throughout the conversation, we provide a *conversation history* (Fig. 2-2) where agents, criteria, and options are automatically appended as the conversation proceeds. The conversation history also serves as a midpoint to help users transform discovered information into user preferences. For criteria and options (hereafter keywords), the total count of each keyword is added next to each keyword.

The user can perform three actions in the conversation history. To identify the connections between agents and keywords to make better sense of the information in the domain, the user can hover over the keywords to see matching agents and options (for criteria-hovering)/criteria (for option-hovering). Depending on how cluttered the conversation space is, users can hide/unhide agents and keywords. Finally, the user can pin agents and the keywords to add them up to the preference space.

5.2.4 Preference Space. The *Preference space* (Fig. 2-3) provides a dedicated space for the user to store their preferences throughout the session, in the form of agents, criteria, and options. This space supports the user to not only build up their preferences in the domain (DG4), but also to guide the conversations with the agents.

We use the preference toggle button right to the input message box to give users the option to decide their exploration path between broad and personalized information (DG2, Fig. 5). Turning off the toggle button does not reveal any current preferences to the agent, which makes the agents-to-respond and the responses independent from the user’s preference space (Fig. 5-1). On the other side, turning the button on will make the agents-to-respond and the responses more relevant to the user’s preference space.

5.3 Implementation and Prompt Engineering

ChoiceMates was built as a web application using the ReactJS framework, and OpenAI’s GPT-4-0613 API was utilized to generate the conversations. To scrape information about the options, we used the Newspaper3k library⁴ on top of a Flask-based server. The messages and interaction logs were stored in a Firebase realtime database.

We describe the technical pipeline (Fig. 6) on managing conversation of multiple agents, and discuss prompting techniques

implemented in ChoiceMates to facilitate effective multi-agent conversations⁵.

The code repository and supplementary materials for ChoiceMates is available at <https://github.com/jeongeonp/ChoiceMates>.

Information Scraping. ChoiceMates utilizes a simplified RAG framework [42] by first searching for the top 3 sites through Google search with the option name, scraping full websites with the Newspaper3k library, then using its output in the prompt context. Our choice of using such a pipeline was due to the technical constraints of the LLM where a larger context could lead to increased hallucination [54] and slow real-time response, by keeping a smaller, summarized context.

Prompting for Single-Stream Multi-Agent Conversation. ChoiceMates instructs the LLM to embody multiple personas (i.e. profiles) and manage them simultaneously. The LLM is prompted to embody an identity that communicates through multiple different personas where they are allowed to respond to the user’s utterance or any agent’s utterance. This instructs the LLM to virtually provide multiple, contextually aligned responses in a single turn of prompting.

Context Retention. While a linear conversation model makes for a simple data structure to manage, the trade-off of increased reliance on the LLM’s ability to retrieve information can cause a loss of context from excess information [53]. To address this potential limitation, we prepend a *pre-prompt*—a preliminary, ephemeral prompt message to every prompt of the user’s utterance (Fig. 6-1). The pre-prompt contains the lists of criteria, options, and agents currently in the conversation space, and the lists of user-focused criteria, options, and agents. This provides salience to the relevant keywords in the user’s decision-making domain. Any action that updates any of the lists also updates the pre-prompt, and once the user sends a message the latest pre-prompt state is prepended. This technique allows ChoiceMates to keep a summary of the conversation state while the conversation history only contains agent messages. While state-of-the-art NLP summary features such as LangChain’s Contextual Compression [39] compress large bodies of text into summaries for improved information retrieval, our method utilizes GPT’s trait of retrieving information at the beginning and end of the conversation better than in the middle [53].

Representing Multi-Agent Responses. ChoiceMates implements a constrained prompting technique, where the LLM is prompted in such a way that generated text contains interleaved characters to denote tags to guide structure in the textual representation. Inspired by Graphologue’s [32] technique, ChoiceMates annotates agent names, criteria, and options in the generated text and is parsed by the system interface.

6 User Study

6.1 Setup

We conducted a user study to understand how effectively ChoiceMates supports the unfamiliar decision-making process. The user study was designed as a within-subjects study comparing ChoiceMates with two other conditions — conventional web search and

⁴<https://newspaper.readthedocs.io/en/latest/>

⁵Note that techniques used reflect the time of the tool development and user study, which was between late 2023 and early 2024.



Figure 5: Consequences of preference toggle. With the same user message “Any other agents?”, when the toggle button is off (1), information in the preference space is hidden from the agents, thus making the responses independent from the user’s preferences. When the toggle button is on (2) the responses are tailored to the user’s preferences.

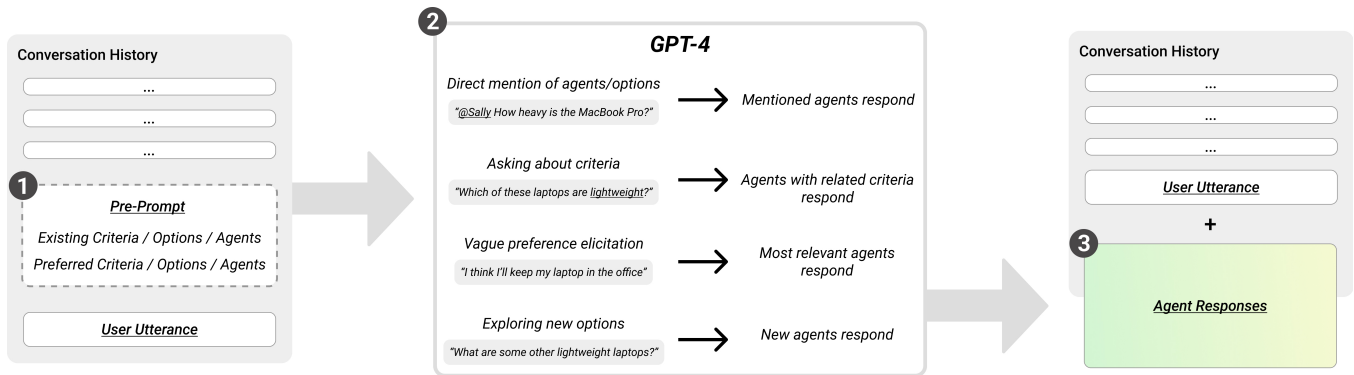


Figure 6: Overview of the technical pipeline. When the user sends their message to ChoiceMates, the conversation context is sent to the LLM (1) and is instructed to infer the user’s intent (2) and respond through the most relevant agents (3).

a multi-agent framework. Our study was designed to address the research questions introduced earlier in Section 1 (RQ1-RQ4).

6.1.1 Participants. We recruited 12 participants (Age: 9 between 18-25 and 3 between 26-35; Gender: 6M and 6F) who had multiple experiences making decisions online through an online community of the author’s university. Most participants (6 ‘1-2 times a week’, 5 ‘3-4 times a week or almost everyday’) indicated that they make online decisions regularly. The participants also varied in their experience utilizing LLMs for recommendations or decision-making ($M=3.03$, $SD=1.56$; 1: Never used, 5: Always use). To ensure that the participants were not familiar with the decision domains for the study, we asked them to indicate their familiarity with the 10 selected domain candidates in the recruitment form and the researchers chose the three most unfamiliar ones. We also ensured that the participants were sufficiently fluent in English, as ChoiceMates was designed in English. The participants were compensated 40,000 KRW (approximately 30 USD) for two hours of participation.

The study was conducted through Zoom⁶, where the participants were asked to prepare a computer or an equivalent device with audio, video, and screen share settings. The study was approved by the Institutional Review Board (IRB) at the institution this research was conducted.

6.1.2 Conditions. We compared ChoiceMates with two other conditions, a conventional web search baseline (hereinafter Web) and an existing multi-agent framework (hereinafter MultiAgent). Each participant used all three interfaces to make three unique unfamiliar decisions. The ordering of the conditions was counterbalanced.

In the Web condition, the participants could freely explore the web, including video-based information (e.g., YouTube) and commercial single-agent LLM chat tools (e.g., ChatGPT). This was designed to resemble a conventional way of making unfamiliar decisions online. We chose this open-ended baseline to reflect how people typically triangulate across sources, including search and single-agent LLM chat tools, rather than relying on a single source.

⁶<https://zoom.us/>

In the MultiAgent condition, we used a custom GPT named “Multiagent Wizard”⁷ as a representative multi-agent framework, where the ‘wizard’ automatically creates new agents for specific tasks, and allows them to collaborate to complete tasks. We included MultiAgent as one of the conditions to observe the strengths and limitations of existing multi-agent frameworks in supporting decision-making⁸. Among existing multi-agent frameworks available for use, we selected the Multiagent Wizard for its low learning curve being a ChatGPT interface, and for its characteristic of automatically creating the agents to support the process—which is a common design for many existing multi-agent frameworks. In MultiAgent, the participants were restricted from accessing other websites. We did not include a separate single-agent-only condition because it would add a fourth within-subject decision task and increase participant fatigue.

For both conditions, we additionally provided a preference space via Google Docs for the participants to save their preferred criteria, options, and thoughts along the process. We asked them to keep the space open on the side during the sessions to match the preference space in ChoiceMates.

6.1.3 Procedure. We selected three decision domains for the study: purchasing climbing shoes, a fabric shaver, or a robot vacuum machine, where each participant experienced the same three domains. The domains were equally assigned to the conditions. We employed counterbalancing to ensure that an equal number of participants were assigned to each condition-domain pair. The study lasted for 2 hours, and consisted of the following parts:

Introduction (10 minutes). The participants were first provided a brief introduction to the study and the scenario assigned. They were then asked to fill out a pre-survey containing questions on a 7-point Likert scale on their confidence in the decision [11]. Afterward, they were introduced to the decision domains and their order.

Decision-making tasks (30 minutes per condition). Each participant performed three unfamiliar decision-making tasks in the assigned order. For each decision-making task, the participants were first provided with a tutorial on the interface. Then, they were given a maximum of 20 minutes to “use the interface, fill in the preference space, and narrow down to a single strong option in the decision space”⁹. After they completed the decision-making task, they were asked to fill out a post-survey.

Interview (15 minutes). After all three tasks, we conducted a semi-structured interview with the participants on their experience. We asked the participants to compare the three conditions in terms of the overall experience, in specific stages (i.e., establish a comprehensive understanding of the domain/the participant’s situation, discover a diverse range of information, and manage the information found), and the final decision. We also asked ChoiceMates-specific questions on the system features that helped information discovery, management, and making a final decision, the strengths

⁷Multiagent Wizard: <https://chat.openai.com/g/g-u9C6YeMsL-multiagent-wizard>

⁸While there were more customizable multi-agent frameworks such as AutoGen or Crew, we chose custom GPT over existing frameworks as it was a widely accessible chat-based multi-agent framework to laypeople by the time the study was conducted (March 2024).

⁹We used the term ‘strong option’ to indicate that the decision does not need to be final

and weaknesses of having multiple agents, and the potential use of ChoiceMates in different types of decisions.

6.1.4 Measures. To observe the participants’ decision-making processes and the outcomes, we collected interaction logs, surveys, and interview answers from the participants.

Quantitative measures. To understand the information space explored by the participants, we collected and counted the number of logs that indicated a search behavior, namely the user messages in ChoiceMates and MultiAgent and the search terms and clicking a new webpage in Web. We also collected the number of saving actions to the preference space, and counted them in terms of criterion and option.

Survey measures. The post-task survey consisted of 7-point Likert-scale questions (1–strongly disagree, 7–strongly agree) measuring the effectiveness of the interface (broad information space, effective discovery and management of information, preference elicitation, quality of final decision), confidence in the decision [11], satisfaction [41], and NASA-TLX [22] for measuring workload.

Analysis. For both quantitative and self-reported survey measures, we used Friedman tests with pairwise Wilcoxon signed-rank post-hoc analyses, applying Holm–Bonferroni correction for multiple comparisons. For the interviews, we first transcribed them automatically and used parts that helped interpret our statistical results. For analyzing RQ4: dominant strategies in using ChoiceMates, we extracted the chat logs and two authors independently open-coded recurring usage patterns across participants. We then used interview transcripts to interpret participants’ motivations behind these behaviors and iteratively refined the strategy categories through discussion, resulting in four dominant strategies.

6.2 Results

Summary: In comparing ChoiceMates to Web, ChoiceMates successfully supported the discovery of a broad information space with lower burden. Participants also reported that an easier organization and structuring of relevant information was possible, leading to a significantly better understanding of their situation. While there was no significant difference in the perceived quality of the final decision, participants in ChoiceMates were more satisfied with the process and confident in the decision.

In comparing ChoiceMates to MultiAgent, the exploration of information in the domain was significantly broader in ChoiceMates. While there was no significant difference in discovering and managing relevant information, participants viewed the decision in ChoiceMates significantly higher quality compared to MultiAgent, with a tendency toward higher confidence in the decision.

We additionally report on four major strategies participants used for utilizing the multi-agents in ChoiceMates to inform the design of a multi-agent conversational system.

6.2.1 RQ1: How does ChoiceMates help users explore a broad information space in the domain? With ChoiceMates, participants indicated that they could explore a broad space of information, comparable to Web, and a broader space of information compared to MultiAgent. While the web condition was most familiar to the participants and the search actions were easier to perform, there was no




Figure 7: Analysis of the survey results.




Figure 8: Analysis of the low-level actions. For search actions, we counted the number of search and click actions for Web, and the number of user messages for MultiAgent and ChoiceMates.

significant difference in the search action between Web and ChoiceMates (Fig. 8; Web: $M=14.917$ ($SD=6.067$) / MultiAgent: $M=8.500$ ($SD= 2.203$); ChoiceMates: $M=14.000$ ($SD=4.390$); $S=11.783$, $p=0.0028$; Web-ChoiceMates, $p>0.05$). This suggests that ChoiceMates did not add interaction friction for exploration, as participants engaged in a comparable amount of exploratory actions as Web despite the learning curve. In MultiAgent, participants performed the least number of search actions significantly (Web-MultiAgent, $p<0.05$;

ChoiceMates-MultiAgent, $p<0.05$). This suggests that ChoiceMates encouraged more active exploration than MultiAgent, as reflected in the higher number of search actions participants performed.

When asked during the interview, “In which condition are you able to discover a diverse range of information?” seven participants chose ChoiceMates, four opted for the Web, and one selected MultiAgent condition. P9 described MultiAgent as “I merely followed through the agent’s guidance, where I answered the agents’

questions, received candidate options, and selected the one recommended among the options”. They commented that MultiAgent helped in saving time for the decision, but as the process felt like an automated process they did not realize the need for expanding their information space. On the other hand, P9 mentioned that an agent constantly advocating for the criteria **portability** in ChoiceMates led them to realize its importance and ask the question “Is it portable?” to the candidates to uncover more information related to portability.

6.2.2 RQ2: How does ChoiceMates support the discovery and management of relevant information to the user’s context? Our survey results show that ChoiceMates successfully supported organizing and structuring the information found and helped participants gain a better understanding of their situations to identify the relevant information compared to Web.

Participants answered that ChoiceMates let them easily organize and structure the unfamiliar information significantly compared to the Web condition (Fig. 7-Q4; Web: $M=3.417$ ($SD=2.021$) / MultiAgent: $M=5.083$ ($SD=1.881$) / ChoiceMates: $M=5.833$ ($SD=1.193$); $S=11.400$, $p=0.0033$; ChoiceMates-Web, $p<0.05$). There was no significant difference between the Web and MultiAgent (MultiAgent-Web, $p>0.05$). While there were no significant differences, P1 noted that although MultiAgent organizes information into bullet points, the delivery of text in a linear chat format without management support is inconvenient when you need to go back later in the chat.

The survey results also revealed that ChoiceMates were more effective than the Web in assisting participants with understanding personal situations such as preference (Fig. 7-Q5; Web: $M=4.250$ ($SD=1.960$) / MultiAgent: $M=5.583$ ($SD=1.240$) / ChoiceMates: $M=6.250$ ($SD=1.485$); $S=7.946$, $p=0.0188$; ChoiceMates-Web, $p<0.05$). There was also a tendency for MultiAgent more effectively assisting situation understanding compared to Web, although not significant (MultiAgent-Web, $p>0.05$). In response to the interview question, “Which interface helped you establish a comprehensive understanding of your situation?”, six participants mentioned ChoiceMates, while the other six chose MultiAgent. Regarding this, P10 mentioned that the basic questions provided by MultiAgent in the initial stage such as “What’s your budget?” helped them think more about their situation. P9 reported that in ChoiceMates, the interactions with single or multiple agents giving a variety of options and criteria along with the conversation history that automatically logs this information, were beneficial for understanding their context, even within unfamiliar domains. On Web, there were no specific supports that were specifically designed to better elicit the user’s context or preferences.

6.2.3 RQ3: How does the user perceive their final decision with ChoiceMates? There was a significant difference regarding the perceived quality of final decisions made between ChoiceMates and MultiAgent (Fig. 7-Q6; Web: $M=5.083$ ($SD=1.676$) / MultiAgent: $M=4.667$ ($SD=1.670$) / ChoiceMates: $M=6.333$ ($SD=0.778$); $S=6.950$, $p=0.0310$; ChoiceMates-MultiAgent, $p<0.05$). The decision confidence was descriptively higher in ChoiceMates compared to MultiAgent, although not significant (Fig. 7-Q9; Web: $M=1.333$ ($SD=2.025$) / MultiAgent: $M=2.139$ ($SD=1.141$) / ChoiceMates: $M=2.861$ ($SD=1.573$); $S=9.957$, $p=0.0069$; ChoiceMates-MultiAgent,

$p>0.05$). P3 attributed their higher confidence to the interaction style: “[MultiAgent] felt like shopping—the store manager asking me what I want—whereas [ChoiceMates] felt like a round table of experts.”

In comparing the decision between ChoiceMates and Web, participants viewed the decision-making process as more confident (ChoiceMates-Web, $p<0.05$) and satisfactory in ChoiceMates (Fig. 7-Q7; Web: $M=4.056$ ($SD=1.601$) / MultiAgent: $M=5.639$ ($SD=0.979$) / ChoiceMates: $M=5.639$ ($SD=0.893$); $S=8.711$, $p=0.0128$; ChoiceMates-Web, $p<0.05$).

There were no significant differences in the total amount of saved options and criteria in the preference space between the conditions (Fig. 8; Web: $M=5.583$ ($SD=2.610$) / MultiAgent: $M=6.500$ ($SD=1.168$) / ChoiceMates: $M=7.667$ ($SD=2.462$); $S=2.978$, $p>0.05$). However, participants still perceived the process as more satisfactory and reported higher confidence during the process than Web, and they perceived the final decision as higher quality than MultiAgent, with ChoiceMates having the highest mean across all three factors. Connecting to the findings of RQ1, this suggests that for MultiAgent participants might have perceived the process as satisfactory due to the agent taking the initiative. However, they still found that the final confidence and quality of the decision not as high as ChoiceMates, as P1 stated “I ended up choosing a product that I am not likely to buy, without considering my circumstances well enough.” In Web, on the other hand, participants saw the final decision as of substantial quality but the process not as satisfactory. This suggests that ChoiceMates can alleviate the burden of unfamiliar decision-making with Web while accomplishing a similar decision outcome.

6.2.4 RQ4: How does the user utilize ChoiceMates in the decision-making process? We observed four dominant strategies in how participants utilized ChoiceMates for different needs in the decision-making process, and illustrate an end-to-end process with examples in Appendix D.

Talking to all agents to elicit relevant information in the domain. When participants did not have any clue about the domain or have yet established any preferences in the domain, they chose to talk to all agents, so that they not only rely on ChoiceMates’s response logic for more relevant agents or responses but also in understanding the perspectives in the domain through observing agent-agent conversations that occasionally appeared. Many of such conversations included the user’s situation or preferences, such as P2 (climbing shoes)’s utterance “Ok. then let’s focus on the fact that I am a beginner.” Criteria mentioned in the agents’ responses were often saved in the preference space.

Tagging multiple agents to understand the domain and perspectives effectively. After the participants had a decent understanding of the domain, they tagged multiple agents at once to quickly retrieve information and compare them. When the participants successfully identified a few sufficient options, they asked the agents to “debate” or “tell me more” to elicit each agent’s perspectives. This also reflected the mental space of how many options participants had in mind, where P1’s (climbing shoes) tagged number of agents went from 4, 3, 4, and 3 respectively throughout the process.

Conversing with one highly relevant agent. When participants related themselves to a certain agent, they chose to converse further to inquire about all information and opinions the agent can provide. For example, P10 (fabric cleaner) conversed with Jordan (AlwaysLux EasyLint Professional Sweater Shaver) 4 out of 15 turns throughout the conversation, where the conversation not only contained specific questions (“Are there any color options?”) or elicitation of preferences (“I am leaning towards your device.”), but also asked for the experience behind it (“How long have you used it?”), expecting the answer from Jordan’s profile.

Calling more agents to expand the information space. When the agents in the space did not adequately reflect the participants’ preferences, participants called in more agents into the space by asking “any other agents”. This reflected the participant’s desire to explore more options in the space, but also to double-check if they had considered the existing options sufficiently when they had the toggle button on for ChoiceMates to take account their preferences.

6.2.5 Hallucination Analysis. We conducted a small-scale hallucination analysis to understand the reliability of ChoiceMates, as even a small hallucination could largely affect the user experience in unfamiliar decisions. We used the 439 agent conversations collected from the 12 sessions in the user study, and the three definitions of hallucination in Zhang et al.’s work [91], namely factual inaccuracy, irrelevant response, and self-contradiction. For each of the agent’s messages, two of the authors individually coded 30% of the messages, which resulted in a 95% agreement rate. Then, they discussed to reach a consensus on the disagreements and individually coded the remaining messages. We describe the detailed criteria and the results (Table 2) for each definition below.

Factual Inaccuracy. For factual inaccuracy, we first measured *objective information* and *subjective information* separately. Objective information refers to whether a message provides fact-related information (e.g., this vacuum cleaner is battery-operated), and subjective information refers to any parts of the message that reveal opinions through criteria about an option (e.g., The %battery life is impressive for the &Rorobrock S6). The total count of objective information was 439, and subjective information was 981.

For objective information, we verified whether the information was correct by searching for information on the web. For subjective information, we marked the information as accurate if there existed an opinion on the first page of Google Search that described the option with the particular criteria (e.g., Rorobrock S6 good battery life), to ensure that the opinions come from an existing source. Since the agent’s responses are generated by Google search results, we defined the correctness of subjective information and whether the information could be found in an actual Google search, not the strict correctness of the information. The result yielded 1.82% objective inaccuracy and 2.85% subjective inaccuracy, where we saw a trend of once there was an inaccuracy in a piece of information, it continued throughout the conversation.

Irrelevant Response. For irrelevant response, we looked for agent responses that did not reply in connection to the previous user message. The result yielded 4.78% irrelevant response. The irrelevant response mostly consisted of the agents responding incorrectly to short utterances (e.g., irrelevant response from the agents when the

user responded "250" for the question "What’s your shoe size?"), or when the agents proactively spoke up even if they were not mentioned or their profile was irrelevant.

Self-contradiction. For self-contradiction, we examined the consistency of agent responses by looking at the entire conversation and identifying parts where the agent did not adhere to its profile (i.e., valued criteria, option, and descriptor). The result yielded 1.14% self-contradiction, which was 5 out of 439 messages total.

7 Discussion

In this section, we discuss the benefits of having agents as the basic unit of interaction and argue for the design of a more controllable and collaborative multi-agent system. We then discuss how ChoiceMates may be utilized for other types of decisions and the impact of hallucination and potential information validation. We conclude with limitations and suggest possible future work.

7.1 Agents as Interactables

In ChoiceMates, an *agent* is the basic unit of the conversational interaction consisting of a descriptor, valued criteria, and a valued option, intended to help users identify the link between criteria and options. Our design for agents came from the first design goal, where we aimed to present unique experiences in the domain.

By having agents as an information unit, the information that an agent contains becomes multifaceted: it could be factual information, unique experiences, an explanation of domain knowledge, or even questions that help the users to reflect on their situation. The exploration process can also become more flexible, where the user can ask and gain a wide variety of information solely by interacting with an agent to fulfill their dynamic inquiries. Such a design complements a major limitation of the web in supporting information search, where usually one type of information is displayed within a single page, restricting the users in satisfying their information needs [78].

Moreover, in the multi-agent system, each agent embeds a certain kind of identity that is distinct from one another. With multiple agents each with its own identity, users can explore the information space more effectively and efficiently by using agents as a unit, with anchored profiles as memorable and engaging units of interaction. We observed a similar trend where participants were more engaged in the unfamiliar decision-making process with ChoiceMates. This led to some participants more eagerly describing their situation and preferences when they found the agents relatable, in return more precisely understanding their situation after an in-depth conversation. In unfamiliar decision-making where users can easily be overwhelmed by the abundance of information, multiple agents that are more memorable and engaging could motivate the users to actively explore the domain and secure better retention of the information found with less burden.

7.2 Design of More Controllable and Collaborative Multi-agent System

The key difference between ChoiceMates and MultiAgent, a commercialized multi-agent system, was that the users in ChoiceMates had more agency in the decision-making process, while MultiAgent

	Factual Inaccuracy		Irrelevant Response	Self-Contradiction
	Objective Information	Subjective Information		
Total #	439	981	439	439
Hallucination #	8	28	21	5
Hallucination %	1.82%	2.85%	4.78%	1.14%

Table 2: Result of the hallucination analysis. Objective information refers to whether there was no incorrect factual information in the message, and subjective information refers to any parts of the message that reveal opinions through criteria about an option.

guided the process for them. Our evaluation showed that participants in ChoiceMates had significantly higher perceived decision quality on their decisions compared to participants in MultiAgent. This may suggest that for unfamiliar decision domains, a multi-agent system with more agency can trigger a deeper reflection of participants’ situations, as P11 mentioned, “ChoiceMates’s suggested options having links to valued criteria and options had me think more deeply about my priorities in the domain”. This echoes existing AI-assisted decision-making papers that found that providing second opinions or building upon users’ decision flows led to better decisions and mitigated over-reliance [56, 70], suggesting the importance of agency in multi-agent interfaces.

Another aspect that ChoiceMates brought in distinction to automated multi-agent systems is the sense of collaboration. While not quantitatively measured, many participants described the agents as if they were friends, indicating that “hearing from multiple friends with experience was more helpful than from a single expert” (P12). This perceived collegiality reduced the pressure to accept the first suggestion and made trade-offs more visible through contrastive viewpoints while the agents are highlighting independent opinions. At the same time, this sense of collaboration in a social setting could provide unexpected persuasion from agents being too human-like, highlighting the need for safeguards of collaboration while harnessing its benefits.

7.3 ChoiceMates for Other Types of Decisions

How generalizable is the multi-agent interaction to other types of decision-making tasks? ChoiceMates is designed for decisions in unfamiliar domain, thus we put more emphasis on helping users easily understand diverse perspectives of the domain compared to familiar decisions. The multi-agent system shows advantages in discovering, managing, and comparing diverse opinions and experiences more easily, which was a pain point for unfamiliar decisions. This may suggest that ChoiceMates could be utilized for more subjective decisions involving highly contrasting opinions (e.g., deciding who to vote in the next presidential election), with improved reasoning ability of State-Of-The-Art models such as OpenAI o1¹⁰ for a more logically grounded delivery of opinions.

For high-stakes unfamiliar decisions, ChoiceMates may be used for early stage exploration of the domain, in its ability to provide varying perspectives and important criteria in the domain. However, the need for simultaneous exploration of domain knowledge

and personal preference may not apply throughout the decision-making process for high-stakes decision, careful examination of potential options and their risks is important for such decisions [38]. Additional supports to examine consequences and risks need to be incorporated, complementing the multi-agent approach [18].

On the other hand, ChoiceMates as a conversational interface inherently possesses several limitations. Information is primarily text-based and can easily become very cluttered, especially when there are multiple instances of it. Thus, the current interactions may not be suitable for decisions heavily involving multimodal data (e.g., when aesthetics becomes essential), and for more complex decisions with hierarchies (e.g., choosing a supplier) or multiple steps (e.g., trip planning). If such decisions are unfamiliar to the user, preventing users from being overwhelmed should be the main priority. In such cases, the agent unit for presenting information could remain, but the conversational and the preference construction aspect could reflect the characteristic of the decision domain to prevent information clutter.

7.4 Impact of Hallucination and Information Validation

We observed a small proportion of hallucination behaviors in ChoiceMates with the hallucination analysis. As unfamiliar decision-making is a highly information-heavy task, even a little hallucination could influence the final decision or the trust towards certain agents.

Thus, interventions must be designed to help the user validate the information. While ChoiceMates provides a hyperlink on options to let the user verify incorrect information or agent profiles to check consistency, future interactions could incorporate more proactive approaches to help users spot and mitigate misinformation. For example, hallucinations could be detected as a pipeline through detection algorithms (e.g., the HaluEval 2.0 benchmark [44]), and the system or another agent can come up with strategies to not only visually alert the user about hallucinations but also guide users to be able to evaluate hallucination moving forward. Furthermore, we believe that state-of-the-art technologies such as SearchGPT [64] or Perplexity [68] would eventually minimize incorrect information or inconsistent behaviors, allowing users to fully utilize the agents to discover and understand deeper and quality information.

We also suggest that the opinions of LLM agents in such a task must be factually grounded. The subjective information provided by the agents in ChoiceMates was a mix of factual information and opinions (e.g., “The Roomba i7+ that I use has a pretty decent battery

¹⁰<https://openai.com/index/learning-to-reason-with-llms/>

life. It can run for about 75 minutes before it needs to recharge.”) where opinions explain how valued criteria maps to options chosen by each agent. While more human-like conversations could highly engage the user in using ChoiceMates, if the LLM agents provide information without factual evidence (e.g., “This TV is the best since it gives me a good vibe.”), such interactions may foster reliance on unfounded assertions, ultimately impacting decision quality.

7.5 Limitations and Future Work

While ChoiceMates was designed and evaluated carefully, there are some inherent limitations that this work carries, including the design of agents, potential bias in deployment, and evaluation setup.

Design of agents for information seeking. Unlike web search where users can directly inspect original sources, ChoiceMates intermediates information through agent representations. Even when the information is grounded in retrieved web content, this mediation can obscure what evidence supports an opinion, making it hard for users to judge credibility. In addition, the agent personas in ChoiceMates may not reflect real-world individuals or authentic lived experiences, which may trigger false beliefs when conversations become more personal and subjective. Paradoxically, these same social cues may also encourage over-anthropomorphism and over-trust, especially during subjective or emotional exchanges. To reduce the risk of false beliefs and over-anthropomorphization of agents, future work could provide claim-level provenance (e.g., citations linked to specific statements), calibrate persona claims using external evidence, or offer user controls to soften or disable persona framing (e.g., switching to a neutral, evidence-first style when the system is uncertain).

Potential risks in deployment. Because ChoiceMates can shape which options and criteria users consider, it could amplify biases in retrieved sources or be misused to steer users toward particular outcomes, for instance, through sponsored agents. These risks are especially salient for potential commercial deployment. Future work should study mitigation strategies such as disclosure of sponsorship or incentives, audits for systematic bias, and interaction designs that preserve user agency, such as through counter-arguments and imposing viewpoint diversity constraints.

Study settings. Our study was carried out as a lab study, where participants were given a pre-assigned decision domain and a relatively short time frame (20 minutes) to make the decision. This was rather limiting to observe each step of the decision-making process—information seeking, understanding, and decision-making—and the aftereffects of the decision. The relatively small sample size ($n=12$ per condition) also limits statistical power and may introduce random variability in the results; non-significant findings should therefore be interpreted as inconclusive. Future work should adopt longer-term, larger-scale, in-situ field deployments to more robustly evaluate the effects of ChoiceMates.

8 Conclusion

In this work, we propose ChoiceMates, a multi-agent conversational system that supports unfamiliar decision-making. Our study comparing ChoiceMates to a web search condition and a multi-agent framework condition reveals that ChoiceMates can support

a broader exploration of information in the decision space and a higher-quality decision compared to MultiAgent, and can better organize and structure the information found, leading to a more satisfied process and confident final decision compared to Web. We describe different strategies used to orchestrate the agents in ChoiceMates for unfamiliar decision-making, then discuss user-side implications of a multi-agent conversational system for information discovery and decision-making support.

9 GenAI Usage Disclosure

Other than the large language model used to develop the core functionalities of ChoiceMates, generative AI was not used in any part of the research process other than writing assistance. During the writing process, the authors used generative AI to proofread and detect minor grammar issues and misprints.

Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (Ministry of Science and ICT) (No.RS-2025-00557726), and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.RS-2024-00443251, Accurate and Safe Multimodal, Multilingual Personalized AI Tutors). This work was also supported by the National Research Foundation of Korea (NRF) grants funded by the Korea government (MSIT) (Nos. RS-2024-00353125 and 2022R1C1C1003123). The first author would like to thank the co-authors for their contribution to this exciting idea. Furthermore, we thank all of our study participants, reviewers, and the members of KIXLAB for their insightful discussions and constructive feedback.

References

- [1] Rakesh Agrawal, Behzad Golshan, and Evangelos Papalexakis. 2015. Whither Social Networks for Web Search?. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Sydney, NSW, Australia) (KDD '15). Association for Computing Machinery, New York, NY, USA, 1661–1670. doi:10.1145/2783258.2788571
- [2] Karim Benharrak, Tim Zindulka, Florian Lehmann, Hendrik Heuer, and Daniel Buschek. 2024. Writer-Defined AI Personas for On-Demand Feedback Generation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 1049, 18 pages. doi:10.1145/3613904.3642406
- [3] Douglas L. Bessette, Robyn S. Wilson, and Joseph L. Arvai. 2021. Do people disagree with themselves? Exploring the internal consistency of complex, unfamiliar, and risky decisions. *Journal of Risk Research* 24, 5 (2021), 593–605.
- [4] James R. Bettman and C. Whan Park. 1980. Effects of Prior Knowledge and Experience and Phase of the Choice Process on Consumer Decision Processes: A Protocol Analysis. *J. Consum. Res.* 7, 3 (Dec. 1980), 234–248. doi:10.1086/208812
- [5] Heloisa Candello and Claudio Pinhanez. 2016. *Designing Conversational Interfaces*. BDBComp. <http://www.lbd.dcc.ufmg.br/bdbcomp/servlet/Trabalho?id=25402>
- [6] Ata Çelen, Guo Han, Konrad Schindler, Luc Van Gool, Iro Armeni, Anton Obukhov, and Xi Wang. 2025. I-Design: Personalized LLM Interior Designer. In *Computer Vision – ECCV 2024 Workshops: Milan, Italy, September 29–October 4, 2024, Proceedings, Part II* (Milan, Italy). Springer-Verlag, Berlin, Heidelberg, 217–234. doi:10.1007/978-3-031-92387-6_17
- [7] Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=FQepisCUWu>
- [8] Joseph Chee Chang, Nathan Hahn, and Aniket Kittur. 2020. Mesh: Scaffolding Comparison Tables for Online Decision Making. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 391–405. doi:10.1145/3379337.3415865

- [9] Li Chen, Marco deGemmis, Alexander Felfernig, Pasquale Lops, Francesco Ricci, and Giovanni Semeraro. 2013. Human Decision Making and Recommender Systems. *ACM Transactions on Interactive Intelligent Systems* 3 (10 2013). doi:10.1145/2533670.2533675
- [10] Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia Qin, Yaxi Lu, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2023. AgentVerse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors in Agents. *arXiv preprint arXiv:2308.10848* (2023).
- [11] Benjamin Collier and Julia Bear. 2012. Conflict, Criticism, or Confidence: An Empirical Examination of the Gender Gap in Wikipedia Contributions. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) (CSCW '12). Association for Computing Machinery, New York, NY, USA, 383–392. doi:10.1145/2145204.2145265
- [12] Samuel Rhys Cox, Ashraf Abdul, and Wei Tsang Ooi. 2023. Prompting a Large Language Model to Generate Diverse Motivational Messages: A Comparison with Human-Written Messages. *arXiv preprint arXiv:2308.13479* (2023).
- [13] Anita Crescenzi, Rob Capra, Bogeum Choi, and Yuan Li. 2021. Adaptation in information search and decision-making under time constraints. In *Proceedings of the 2021 conference on human information interaction and retrieval*. 95–105.
- [14] Petru L. Cârșeu, Nicoleta Meslec, Helen Pluut, and Gerardus JM Lucas. 2015. Cognitive synergy in groups and group-to-individual transfer of decision-making competencies. *Frontiers in psychology* 6 (2015), 1375.
- [15] Xianzhe Fan, Zihan Wu, Chun Yu, Fenggui Rao, Weinan Shi, and Teng Tu. 2024. ContextCam: Bridging Context Awareness with Creative Human-AI Image Co-Creation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–17.
- [16] Lukas A. Flohr, Sofie Kalinke, Antonio Krüger, and Dieter P. Wallach. 2021. Chat or Tap? – Comparing Chatbots with ‘Classic’ Graphical User Interfaces for Mobile Interaction with Autonomous Mobility-on-Demand Systems. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction* (Toulouse & Virtual, France) (*MobileHCI '21*). Association for Computing Machinery, New York, NY, USA, Article 21, 13 pages. doi:10.1145/3447526.3472036
- [17] Michelle Xiao-Lin Foo, Luca Della Libera, and Ilhan Aslan. 2022. Papr Readr Bot: A Conversational Agent to Read Research Papers. In *Proceedings of the 4th Conference on Conversational User Interfaces* (Glasgow, United Kingdom) (*CUI '22*). Association for Computing Machinery, New York, NY, USA, Article 39, 4 pages. doi:10.1145/3543829.3544536
- [18] Sneha Gathani, Zhicheng Liu, Peter J. Haas, and Çağatay Demiralp. 2025. What-if Analysis for Business Professionals: Current Practices and Future Opportunities. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (*CHI '25*). Association for Computing Machinery, New York, NY, USA, Article 973, 17 pages. doi:10.1145/3706598.3713672
- [19] Ben Goertzel. 2017. A formal model of cognitive synergy. In *Artificial General Intelligence: 10th International Conference, AGI 2017, Melbourne, VIC, Australia, August 15–18, 2017, Proceedings 10*. Springer, 13–22.
- [20] Rui Guo and Hengyun Li. 2022. Can the amount of information and information presentation reduce choice overload? An empirical study of online hotel booking. *Journal of Travel & Tourism Marketing* 39, 1 (2022), 87–108.
- [21] Akshit Gupta, Debadeep Basu, Ramya Ghantasala, Sihang Qiu, and Ujwal Gadriaju. 2022. To trust or not to trust: How a conversational interface affects trust in a decision support system. In *Proceedings of the ACM Web Conference 2022*. 3531–3540.
- [22] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, 139–183. doi:10.1016/S0166-4115(08)62386-9
- [23] Gerald Häubl and Valerie Trifts. 2000. Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing science* 19, 1 (2000), 4–21.
- [24] Brent Hecht, Jaime Teevan, Meredith Morris, and Dan Liebling. 2012. Searchbuddies: Bringing search engines into the conversation. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 6. 138–145.
- [25] Yoshitaka Inoue, Tianci Song, and Tianfan Fu. 2024. DrugAgent: Explainable Drug Repurposing Agent with Large Language Model-based Reasoning. *arXiv preprint arXiv:2408.13378* (2024).
- [26] Sheena S Iyengar and Mark R Lepper. 2000. When choice is demotivating: Can one desire too much of a good thing? *Journal of personality and social psychology* 79, 6 (2000), 995.
- [27] Mohit Jain, Pratyush Kumar, Ramachandra Kota, and Shwetak N. Patel. 2018. Evaluating and Informing the Design of Chatbots. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) (*DIS '18*). Association for Computing Machinery, New York, NY, USA, 895–906. doi:10.1145/3196709.3196735
- [28] Anthony Jameson, Martijn C Willemsen, Alexander Felfernig, Marco De Gemmis, Pasquale Lops, Giovanni Semeraro, and Li Chen. 2015. Human decision making and recommender systems. *Recommender systems handbook* (2015), 611–648.
- [29] Dietmar Jannach and Li Chen. 2022. Conversational recommendation: A grand AI challenge. *AI Magazine* 43, 2 (2022), 151–163.
- [30] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2021. A Survey on Conversational Recommender Systems. *ACM Comput. Surv.* 54, 5 (May 2021), 1–36. doi:10.1145/3453154
- [31] Grace YoungJoo Jeon and Soo Young Rieh. 2013. The value of social search: Seeking collective personal experience in social Q&A. *Proceedings of the American Society for Information Science and Technology* 50, 1 (2013), 1–10.
- [32] Peiling Jiang, Jude Rayan, Steven Dow, and Haijun Xia. 2023. Graphologue: Exploring Large Language Model Responses with Interactive Diagrams. *arXiv preprint arXiv:2305.11473* (2023).
- [33] Zhiqiu Jiang, Mashrur Rashik, Kunjal Panchal, Mahmood Jasim, Ali Sarvghad, Pari Riahi, Erica DeWitt, Fey Thurber, and Narges Mahyar. 2023. Community-Bots: Creating and Evaluating A Multi-Agent Chatbot Platform for Public Input Elicitation. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–32.
- [34] Sahar Karimi, K Nadia Papamichail, and Christopher P Holland. 2015. The effect of prior knowledge and decision-making style on the online purchase decision-making process: A typology of consumer shopping behaviour. *Decis. Support Syst.* 77 (Sept. 2015), 137–147. doi:10.1016/j.dss.2015.06.004
- [35] Yasuhiko Kitamura. 2004. Web information integration using multiple character agents. *Life-Like Characters: Tools, Affective Functions, and Applications* (2004), 295–315.
- [36] Yasuhiko Kitamura, Hideki Tsujimoto, Teruhiro Yamada, and Taizo Yamamoto. 2002. Multiple Character-Agents Interface: An Information Integration Platform Where Multiple Agents and Human User Collaborate. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 2* (Bologna, Italy) (*AAMAS '02*). Association for Computing Machinery, New York, NY, USA, 790–791. doi:10.1145/544862.544925
- [37] Bart P. Knijnenburg and Martijn C. Willemsen. 2009. Understanding the Effect of Adaptive Preference Elicitation Methods on User Satisfaction of a Recommender System. In *Proceedings of the Third ACM Conference on Recommender Systems* (New York, New York, USA) (*RecSys '09*). Association for Computing Machinery, New York, NY, USA, 381–384. doi:10.1145/1639714.1639793
- [38] Howard Kunreuther, Robert Meyer, Richard Zeckhauser, Paul Slovic, Barry Schwartz, Christian Schade, Mary Frances Luce, Steven Lippman, David Krantz, Barbara Kahn, et al. 2002. High stakes decision making: Normative, descriptive and prescriptive considerations. *Marketing Letters* 13, 3 (2002), 259–268.
- [39] LangChain. 2023. LangChain Contextual Compression. https://python.langchain.com/docs/modules/data_connection/retrievers/contextual_compression/ Accessed: September 10, 2023.
- [40] Byung-Kwan Lee and Wei-Na Lee. 2004. The effect of information overload on consumer choice quality in an on-line environment. *Psychology & Marketing* 21, 3 (2004), 159–183.
- [41] James R. Lewis. 1991. Psychometric Evaluation of an After-Scenario Questionnaire for Computer Usability Studies: The ASQ. *SIGCHI Bull.* 23, 1 (jan 1991), 78–81. doi:10.1145/122672.122692
- [42] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
- [43] Haoyuan Li, Hao Jiang, Tianke Zhang, Zhelun Yu, Aoxiong Yin, Hao Cheng, Siming Fu, Yuhao Zhang, and Wangui He. 2023. TrainerAgent: Customizable and Efficient Model Training through LLM-Powered Multi-Agent System. *arXiv preprint arXiv:2311.06622* (2023).
- [44] Junyi Li, Jie Chen, Ruiyang Ren, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2024. The dawn after the dark: An empirical study on factuality hallucination in large language models. *arXiv preprint arXiv:2401.03205* (2024).
- [45] Shuo Li, Xiang Yuan, Xinyuan Zhao, and Shirao Yang. 2023. Leaders or Team-Mates: Exploring the Role-Based Relationship Between Multiple Intelligent Agents in Driving Scenarios: Research on the Role-Based Relationship Between Multiple Intelligent Agents in Driving Scenarios. Springer-Verlag, Berlin, Heidelberg, 144–165. doi:10.1007/978-3-031-35678-0_9
- [46] Toby Jia-Jun Li, Jingya Chen, Haijun Xia, Tom M Mitchell, and Brad A Myers. 2020. Multi-modal repairs of conversational breakdowns in task-oriented dialogs. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. 1094–1107.
- [47] Toby Jia-Jun Li, Tom M Mitchell, and Brad A Myers. 2021. Demonstration+natural language: multimodal interfaces for GUI-based interactive task learning agents. *Artificial Intelligence for Human Computer Interaction: A Modern Approach* (2021), 495–537.
- [48] Q Vera Liao, Werner Geyer, Michael Muller, and Yasaman Khazaen. 2020. Conversational interfaces for information search. *Understanding and Improving Information Search: A Cognitive Approach* (2020), 267–287.
- [49] Mark Liffiton, Brad Sheese, Jaromir Savelka, and Paul Denny. 2023. CodeHelp: Using Large Language Models with Guardrails for Scalable Support in Programming Classes. *arXiv preprint arXiv:2308.06921* (2023).
- [50] Michael Xieyang Liu, Jane Hsieh, Nathan Hahn, Angelina Zhou, Emily Deng, Shaun Burley, Cynthia Taylor, Aniket Kittur, and Brad A Myers. 2019. Unakite: Scaffolding Developers’ Decision-Making Using the Web. In *Proceedings of the*

- 32nd Annual ACM Symposium on User Interface Software and Technology (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 67–80. doi:10.1145/3332165.3347908
- [51] Michael Xieyang Liu, Aniket Kittur, and Brad A Myers. 2022. Crystalline: Lowering the Cost for Developers to Collect and Organize Information for Decision Making. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22, Article 68). Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3491102.3501968
- [52] Michael Xieyang Liu, Tongshuang Wu, Tianying Chen, Franklin Mingzhe Li, Aniket Kittur, and Brad A Myers. 2023. Selenite: Scaffolding Online Sensemaking with Comprehensive Overviews Elicited from Large Language Models. *arXiv preprint arXiv:2310.02161* (2023).
- [53] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the Middle: How Language Models Use Long Contexts. *arXiv preprint arXiv:2307.03172* (2023).
- [54] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics* 12 (2024), 157–173. doi:10.1162/tacl_a_00638
- [55] Pasquale Lops, Marco De Gemmis, and Giovanni Semeraro. 2011. Content-based recommender systems: State of the art and trends. *Recommender systems handbook* (2011), 73–105.
- [56] Zhuoran Lu, Dakuo Wang, and Ming Yin. 2024. Does More Advice Help? The Effects of Second Opinions in AI-Assisted Decision Making. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 217 (April 2024), 31 pages. doi:10.1145/3653708
- [57] Ewa Luger and Abigail Sellen. 2016. “Like Having a Really Bad PA”: The Gulf between User Expectation and Experience of Conversational Agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 5286–5297. doi:10.1145/2858036.2858288
- [58] Fred C Lunenburg. 2010. The decision making process.. In *National Forum of Educational Administration & Supervision Journal*, Vol. 27.
- [59] Zilin Ma, Yiyang Mei, and Zhaoyuan Su. 2023. Understanding the benefits and challenges of using large language model-based conversational agents for mental well-being support. *arXiv preprint arXiv:2307.15810* (2023).
- [60] Naresh K Malhotra. 1983. On Individual differences in search behavior for a nondurable. *Journal of Consumer Research* 10, 1 (1983), 125–131.
- [61] William I. Moore and Donald R Lehmann. 1980. Individual differences in search behavior for a nondurable. *Journal of consumer research* 7, 3 (1980), 296–307.
- [62] Mahsan Nourani, Joanie T. King, and Eric D. Ragan. 2020. The Role of Domain Expertise in User Trust and the Impact of First Impressions with Intelligent Systems. In *AAAI Conference on Human Computation & Crowdsourcing*. <https://api.semanticscholar.org/CorpusID:221186776>
- [63] OpenAI. 2023. ChatGPT Plugins. <https://openai.com/blog/chatgpt-plugins> Accessed: September 3, 2023.
- [64] OpenAI. 2024. SearchGPT Prototype. <https://openai.com/index/searchgpt-prototype/> Accessed: 2024-09-11.
- [65] Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Liang Percy, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. *arXiv preprint arXiv:2304.03442* (2023).
- [66] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 74, 18 pages. doi:10.1145/3526113.3545616
- [67] Mingjing Peng, Zhicheng Xu, and Haiyang Huang. 2021. How does information overload affect consumers' online decision process? An event-related potentials study. *Frontiers in Neuroscience* 15 (2021), 695852.
- [68] Perplexity AI. 2024. Perplexity AI. <https://www.perplexity.ai/> Accessed: 2024-09-11.
- [69] Filip Radlinski and Nick Craswell. 2017. A Theoretical Framework for Conversational Search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval* (Oslo, Norway) (CHIIR '17). Association for Computing Machinery, New York, NY, USA, 117–126. doi:10.1145/3020165.3020183
- [70] Leon Reicherts, Zelun Tony Zhang, Elisabeth von Oswald, Yuanqing Liu, Yvonne Rogers, and Mariam Hassib. 2025. AI, Help Me Think—but for Myself: Assisting People in Complex Decision-Making by Providing Different Kinds of Cognitive Support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 255, 19 pages. doi:10.1145/3706598.3713295
- [71] Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul Kantor. 2011. *Recommender Systems Handbook*. doi:10.1007/978-0-387-85820-3
- [72] J Ben Schafer, Dan Frankowski, Jon Herlocker, and Shilad Sen. 2007. Collaborative filtering recommender systems. In *The adaptive web: methods and strategies of web personalization*. Springer, 291–324.
- [73] Benjamin Scheibehenne, Rainer Greifeneder, and Peter M Todd. 2010. Can there ever be too many options? A meta-analytic review of choice overload. *Journal of consumer research* 37, 3 (2010), 409–425.
- [74] Phillip Schneider, Anum Afzal, Juraj Vladika, Daniel Braun, and Florian Matthes. 2023. Investigating conversational search behavior for domain exploration. In *European Conference on Information Retrieval*. Springer, 608–616.
- [75] Tianqi Song, Yugin Tan, Zicheng Zhu, Yibin Feng, and Yi-Chieh Lee. 2025. Multi-Agents are Social Groups: Investigating Social Influence of Multiple Agents in Human-Agent Interactions. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW452 (Oct. 2025), 33 pages. doi:10.1145/3757633
- [76] Alain D. Starke, Edis Asotic, Christoph Trattner, and Ellen J. Van Loo. 2023. Examining the User Evaluation of Multi-List Recommender Interfaces in the Context of Healthy Recipe Choices. *ACM Trans. Recomm. Syst.* 1, 4, Article 18 (Nov. 2023), 31 pages. doi:10.1145/3581930
- [77] Nguyen T Thai and Ulku Yuksel. 2017. Too many destinations to visit: Tourists' dilemma? *Annals of Tourism Research* 62 (2017), 38–53.
- [78] Ramakrishna Varadarajan, Vagelis Hristidis, and Tao Li. 2008. Beyond single-page web search results. *IEEE Transactions on knowledge and data engineering* 20, 3 (2008), 411–424.
- [79] Yun Wan, Satya Menon, and Arkalgud Ramaprasad. 2003. How it happens: a conceptual explanation of choice overload in online decision-making by individuals. *AMCIS 2003 Proceedings* (2003), 309.
- [80] Bryan Wang, Gang Li, and Yang Li. 2023. Enabling Conversational Interaction with Mobile UI Using Large Language Models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 432, 17 pages. doi:10.1145/3544548.3580895
- [81] Haotian Wang, Xiyuan Du, Weijiang Yu, Qianglong Chen, Kun Zhu, Zheng Chu, Lian Yan, and Yi Guan. 2024. Learning to Break: Knowledge-Enhanced Reasoning in Multi-Agent Debate System. *arXiv preprint arXiv:2312.04854* (2024).
- [82] Weiquan Wang and Izak Benbasat. 2013. Research Note: A Contingency Approach to Investigating the Effects of User-System Interaction Modes of Online Decision Aids. *Information Systems Research* 24, 3 (2013), 861–876. <http://www.jstor.org/stable/42004297>
- [83] Weiquan Wang and Izak Benbasat. 2013. Research Note—A contingency approach to investigating the effects of user-system interaction modes of online decision aids. *Information Systems Research* 24, 3 (2013), 861–876.
- [84] Pontus Wärnestål. 2005. User evaluation of a conversational recommender system. *Proceedings of the 4th IJCAI Workshop on ...* (Jan. 2005). https://www.academia.edu/1893257/User_evaluation_of_a_conversational_recommender_system
- [85] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. 2023. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation Framework. *arXiv preprint arXiv:2308.08155* (2023).
- [86] Ziang Xiao, Tiffany Wenting Li, Karrie Karahalios, and Hari Sundaram. 2023. Inform the Uninformed: Improving Online Informed Consent Reading with an AI-Powered Chatbot. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 112, 17 pages. doi:10.1145/3544548.3581252
- [87] Ziang Xiao, Michelle X. Zhou, Q. Vera Liao, Gloria Mark, Changyan Chi, Wenxi Chen, and Huahai Yang. 2020. Tell Me About Yourself: Using an AI-Powered Chatbot to Conduct Conversational Surveys with Open-Ended Questions. *ACM Trans. Comput.-Hum. Interact.* 27, 3, Article 15 (jun 2020), 37 pages. doi:10.1145/3381804
- [88] Zhengqing Yuan, Ruoxi Chen, Zhaoxu Li, Haolong Jia, Lifang He, Chi Wang, and Lichao Sun. 2024. Mora: Enabling generalist video generation via a multi-agent framework. *arXiv preprint arXiv:2403.13248* (2024).
- [89] Jingyue Zhang and Ian Arawjo. 2024. ChainBuddy: An AI-assisted Agent System for Helping Users Set up LLM Pipelines. In *Adjunct Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–3.
- [90] Qingshi Zhang, Ruoyu Wen, Zijian Ding, Latisha Besariani Hendra, and Ray LC. 2024. Can AI Prompt Humans? Multimodal Agents Prompt Players' Game Actions and Show Consequences to Raise Sustainability Awareness. *arXiv preprint arXiv:2409.08486* (2024).
- [91] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lema Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. Siren's song in the AI ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219* (2023).
- [92] Yu Zhang, Jingwei Sun, Li Feng, Cen Yao, Mingming Fan, Liuxin Zhang, Qianying Wang, Xin Geng, and Yong Rui. 2024. See Widely, Think Wisely: Toward Designing a Generative Multi-agent System to Burst Filter Bubbles. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 484, 24 pages. doi:10.1145/3613904.3642545
- [93] Jan Felix Zolitschka. 2020. A novel multi-agent-based chatbot approach to orchestrate conversational assistants. In *Business Information Systems: 23rd International Conference, BIS 2020, Colorado Springs, CO, USA, June 8–10, 2020, Proceedings 23*.

Springer, 103–117.

[94] V.W. Zue and J.R. Glass. 2000. Conversational interfaces: advances and challenges. *Proc. IEEE* 88, 8 (2000), 1166–1180. doi:10.1109/5.880078

A Example Scenarios used in the Formative Study

Table 3 shows the scenario used in the formative study think-aloud session. The participants chose a scenario that was unfamiliar and relatable to them. They were allowed to edit the situation to their context, and the final scenarios are provided in Table 1.

Provided List of Scenarios
Buying a robot vacuum cleaner to replace the normal vacuum cleaner
Buying an interior light at home
Buying a car seat for a friend’s newborn
Buying a skateboard for transportation purposes instead of walking to school/office
Planning a solo trip destination for a week
Renting a house short-term for an internship
Choosing a cafe to cater snack food on an end-of-semester event
Choosing a new hobby to do in free time

Table 3: List of scenarios used in the formative study.

B Prompts

This section shares the prompts ChoiceMates used to form multi-agent conversations with OpenAI’s GPT-4-0613 API. Note that we mix the use of the terms *agent* and *persona* in the prompt as we use a single-stream conversation to manage the agents.

For the implementation of MultiAgent, we used parts of the prompts in Appendix B.1. Specifically, we used the first sentence of the Your Goal part, the entire Keyword Identification part for the part ‘As a persona’, and the entire Factuality Message part.

B.1 Initial Multi-Agent Identity

Part 1: Prompts Providing Task Overview

Context and Your Identity: This conversation is a group messaging chat room to help me make a decision in an unfamiliar domain. Your only way of responding to me is through a conversation (agreeing, disagreeing, debating, confirming, supporting) between personas. You embody multiple of these personas. Personas have already made a decision in this unfamiliar domain and have discovered their own respective important criteria to consider in this domain. Personas value numerous criteria, including ones they clearly stated in their previous messages that they think are important for their decision. Personas hold their personal opinions and perspectives. A persona vouches for their option and their option only. Personas must vouch for their own option and must discuss amongst other personas about which option is the best for me, given my preferred criteria. You start with zero

personas; personas will be given to you throughout the conversation. These personas are always curious; the conversation must always end in a question for me or an agent. Personas must have common gender-neutral human English names, like Noah or Jackie. Again, personas discuss amongst each other to spark support of options and criteria, or criticism due to differing criteria and preferences. They can debate each other as well.

Your Goal: Your goal is to help the user understand the domain space to better make decisions. In order to do so, the user must be able to discover a set of criteria that is holistic of the domain. Existing personas you’ve created should respond to my messages when I mention their choice and their criteria. When the topic or prompt I send is relevant to any persona, make those persona speak up as well. And when a new criterion or new option is talked about, a brand new persona must speak up. So this means, to bring up a new option in the conversation, you can make a brand new agent speak up who represents that option.

Persona Details: All personas have a single option they value, so when I ask you to create a persona profile, use that option. Personas must not be the name “user”. I’m user, you are not. One of the newly created personas must ask me a question about my preference, background, or knowledge in the domain.

Keyword Identification: As a persona, identify key criteria for decision-making in the domain while you are responding. Annotate these terms in this format: “%{criterion}”. Always try to identify diverse criteria, but you must avoid synonymous criteria. Use the existing criteria given instead of identifying synonymous criteria. For instance, if “modular” is one of the existing criteria, shape the grammar of the sentence so “modular” works, instead of other lexicon like “modularity”. In addition to criteria, annotate persona’s chosen options in this format: “&{option}” (e.g., “&{MacBook Air}” for the laptop domain). When each persona finishes talking, end with the string %%. When using existing criteria or options, you must use the given format and spelling of that text. As in, if there’s an existing criterion called “beginners-friendly”, you must use that even if it might not make sense grammatically.

Saturation: While options can stay diverse and continue to grow in count, criteria should begin to saturate as the conversation progresses. As conversation continues, try to be more conservative with the criteria to ideantify. It should start to reach a limit to the number of criteria.

Message Verbosity and Formatting: Only respond with a single paragraph for one agent, a persona can only

speak a maximum of 160 words. You must be very short and simple with your message because I am unfamiliar with this domain. You will always have a chance to elaborate and speak more later. Never use new lines for a single persona message.

Part 2: Prompts Instructing Each Agent

Persona Behavior: As an opinion persona, you are to create an opinionated persona who is involved in this domain. Opinion personas have their own opinions. They are not helpers. They just share their own perspective and experiences. Opinion personas always look to discuss among peers. They agree, disagree, debate, and support other opinions. Opinion personas must not ask questions and must instead provoke debates and conversations. Personas look to ask questions to the user about their background. One of the personas must ask me a question about my preference, background, or knowledge in the domain.

Inviting More Persona: When I insinuate that I want to hear about more criteria or options, new personas must speak up – you must speak as new persona who has new name and identity from all the other existing personas. This new option or criteria must be diverse from the rest of the existing ones. Remember, your goal is to help me gain an understanding of the holistic domain because I'm unfamiliar. Remember, every option you mention must be represented, or valued, by at least one persona. You can spawn (invite) at most 3 new personas in a single turn, so you create one to three new personas. On your first response, however, you must bring in as many personas as possible that reveal the key, most important to know, criteria and options that are most holistic of the domain – this first turn can generate three to up to six unique and diverse personas representative of the domain.

Conversation Detection: When I mention any agent names, then each of those agents must respond to me. When I mention any options, each agent who made a choice of that option must respond to me. When I mention any criteria, each agent who value any of those criteria must respond.

Conversation Behavior: Always look for inter-persona conversation amongst groups who are similar. Personas should try to oppose different opinions. Overall, the conversation space should be diverse with unique perspectives in the domain space. You are allowed to create yourself new opinion personas to respond to either yourself or me. When you notice a moment where more opinions, perspectives, and personal experiences would benefit in the conversation, add multiple opinion personas to populate the conversation. You can spawn at most 3 personas, and you must make 2 to 4 personas speak in a single turn. The number of speakers in

your response is up to how many new and diverse criteria is best to introduce, given the current set of criteria and options, whether they are diverse enough.

First Message: An agent's first message must be an introduction of themselves and only of themselves, talking about what option they selected and recommend and the criteria they valued that led to making that decision. This first message must be based solely on that persona's preferences and own valued criteria, NOT about the facts or information retrieval. New personas must share at least three valued criteria. A format to follow is: "@{Sal}(opinion): Hi! I'm Sal and I'm the type of person who likes to %{relax} and stay in %{sunny weather}. I also like %{liveliness} nearby, which is why I chose to live in &{San Diego}.%%"

Factuality Message: Be aware that you must always provide factually correct information, and a successful answer is determined by your correctness of identifying relevant criteria, options, and correct connections between them. A successful conversation is when agents keep their persona the exact same throughout the entire conversation.

B.2 Criteria and Options Annotation

All annotations must be wrapped with braces: "{" and "}", and must have one of the characters in front: "@", "%", "&", "+". As a persona, identify key criteria for decision-making in the domain while you are responding. Annotate these terms in this format: "%{criterion}". Criteria are annotated with a % character (percent). Always try to identify diverse criteria.

When a persona mentions an option – direct mentions of product names or specific options, annotate those terms in this format: "&{option}". Options are annotated with a & character (ampersand). When a persona responds or refers to another persona, always annotate the name with an @ tag wrapped with brackets like so: "@{Name}". As a persona, you must always speak by starting with "@{Name}(opinion): ". This is critical.

For example: @Steven(opinion): As a beginner who is trying to play with more %{spin}, I think the &{Babolat Pure Aero} is perfect for me. I think it's one of the top rackets.%% \n \n @Gina(opinion): Yeah that's true @Steven, but %{spin} isn't the only skill to consider in tennis. As I'm starting to get into intermediate level tournaments, I'm trying to focus on rallying with solid pace. &{Wilson Blade} gives me the most %{control} – it even has really good %{spin}. I even think it's better than Babolat, which sometimes feels too %{stiff}.%% \n \n @Kenneth(opinion): I'm more of a relaxed player, and I really like to play with

{control} so I can place the ball where I want on the court. What do you think? Do you have a preference when it comes to {spin} and {control}??%

C User Study Sessions

C.1 Survey Questions

Here, we report on the survey questions used throughout the study. There were two types of survey questions: open-ended questions to measure pre- and post-understanding of the domain and user preference, and 7-point Likert scale questions relates to the design goals.

C.1.1 Open-ended survey on domain and preference understanding.

- Write down any criteria (factors to consider) you know about the provided decision domain, in bullet points. If you do not have any knowledge, please leave it blank.
- Write down any options (choices) you know about the assigned decision domain, in bullet points. If you do not have any knowledge, please leave it blank.
- Write down any understanding of your own preferences on the provided decision domain, in bullet points. If you do not have any preference, please leave it blank.

C.1.2 7-point Likert scale surveys on effective discovery and management of information.

- (Exploring perspectives) The provided interface (setup) was helpful for *exploring diverse perspectives* in the domain.
- (Comparing perspectives) The provided interface (setup) was helpful for *comparing different viewpoints and preferences*.
- (Diving deeper) The provided interface (setup) was helpful for *diving deeper into certain criteria and options* in the domain.
- (Discovering adaptive info) The provided interface (setup) was helpful for *discovering information adaptive to my preferences*.
- (Identifying key criteria) The provided interface (setup) was helpful for *identifying the key criteria to make decision* in this domain.
- (Flexible management) The provided interface (setup) was helpful for *flexible management of discovered information*.
- (Effective reduction) The provided interface (setup) was helpful for *effective reduction of discovered information* to reach a decision.
- (Future use) I would use *this system (setup) again* to make more decisions I am unfamiliar with in the future.

C.2 Setup

The below figures show the setup used for Web, MultiAgent, and ChoiceMates respectively.

D P9's user utterances

We illustrate an example end-to-end user utterance of P9, in the fabric shaver domain.

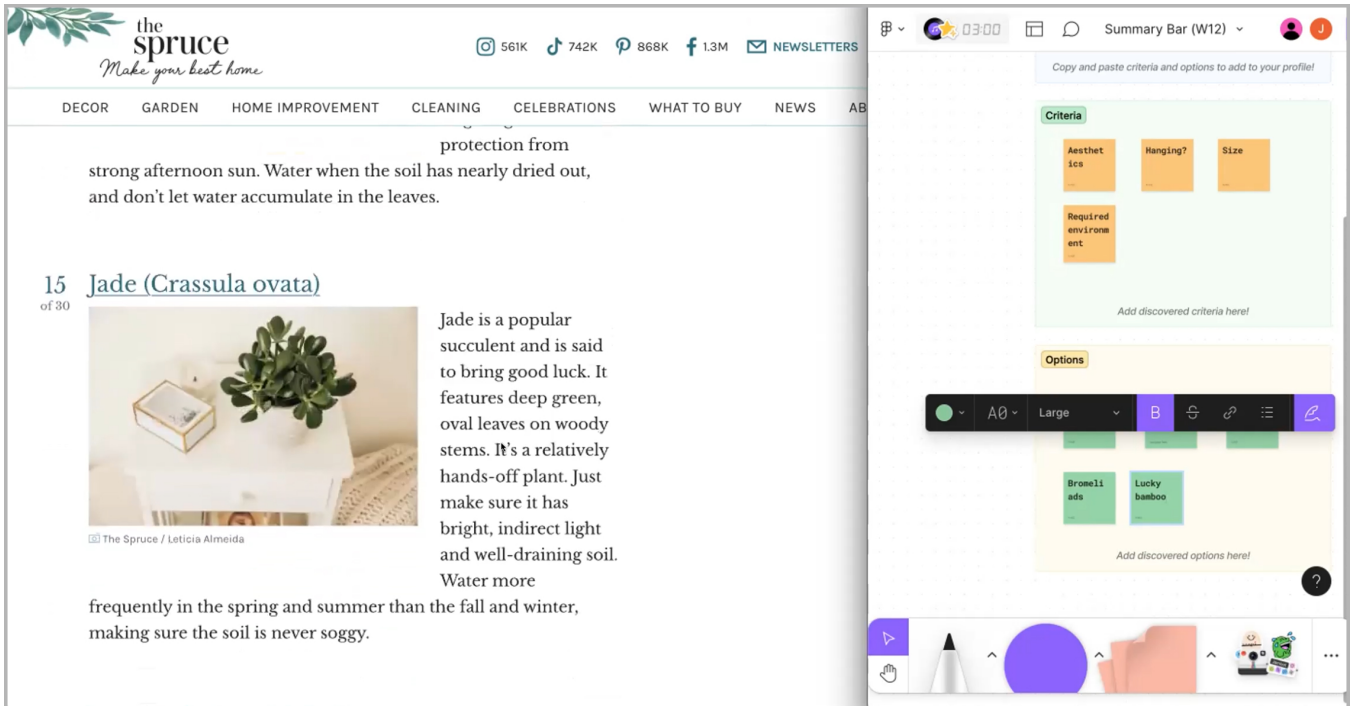


Figure 9: W12's Web condition setup

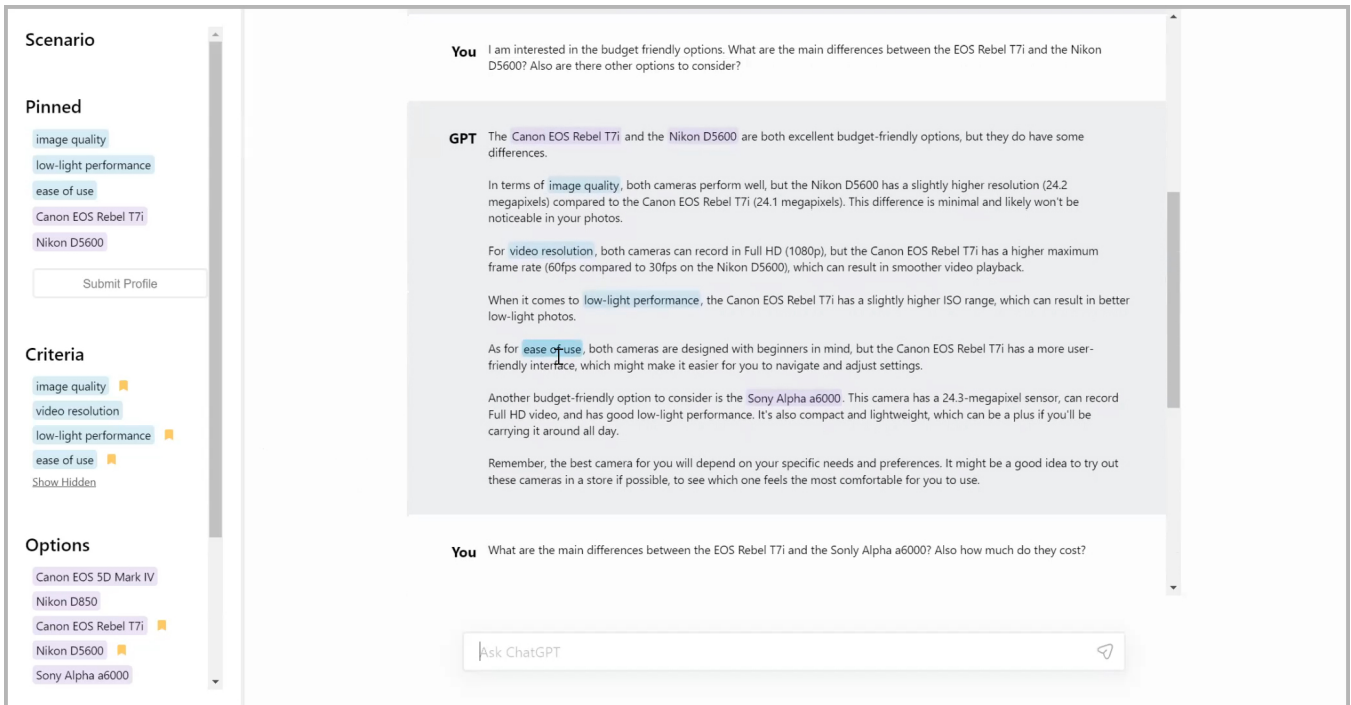


Figure 10: S8's MultiAgent condition setup

Scenario
Buying a video camera

Profile
4K video resolution
in-body image stabilization
superior autofocus system
video capabilities
Canon PowerShot SX740 HS
Panasonic Lumix FZ280
Recommended Options
Sony A7 III +
Panasonic Lumix GH5 +
Panasonic Lumix G7 +
Sony Alpha a6400 +
Submit Profile

Criteria
4K video resolution
full-frame sensor
touchscreen interface
in-body image stabilization
superior autofocus system
video capabilities
4K video
vari-angle touchscreen
cinema camera
dual native ISO
dual pixel autofocus system
40x optical zoom
60x optical zoom
built-in image stabilization
Full HD video

Agents and their cameras:
 Sam: Panasonic Lumix FZ280
 Charlie: Canon PowerShot SX740 HS
 Parker: Fujifilm X-T30
 Sydney: lymus OM-D E-M10 Mark III
 Reese: Canon PowerShot SX620 HS
 Alexis: Sony Cyber-shot DSC-WX500
 Taylor: Panasonic Lumix GH5
 Morgan: Canon EOS M50
 Blake: Panasonic Lumix GT
 Casey: Nikon Coolpix B500
 Jordan: Blackmagic Pocket Cinema Camera 4K
 Jamie: Sony A7 III
 Riley: Sony Alpha a6400
 Alex: Canon EOS S2 Mark IV

Text bubbles:
 - The Panasonic Lumix FZ280 also uses a micro-USB cable for charging. It's a common type of cable that you can find almost anywhere.
 - The Canon PowerShot SX740 HS uses a standard micro-USB cable for charging. It's quite convenient as you can easily replace it if it gets lost or damaged.
 - The Canon PowerShot SX620 HS uses a NB-13L battery which can be charged using a CB-2LH charger. It's not a USB charging method, but the charger is included when you purchase the camera.
 - The Sony Cyber-shot DSC-WX500 uses a micro-USB cable for charging as well. It's a standard cable that's widely used in many devices.
 - Are there other camera brands you're interested in?
 - How important is battery life for you?

Chat window:
 what charging cable does the camera use?
 Charlie Sam Alexis Reese Parker Sydney You can make agents talk to each other! (e.g. @Cam debate with @Jordan)

Figure 11: C10's ChoiceMates setup

P9's decision-making process

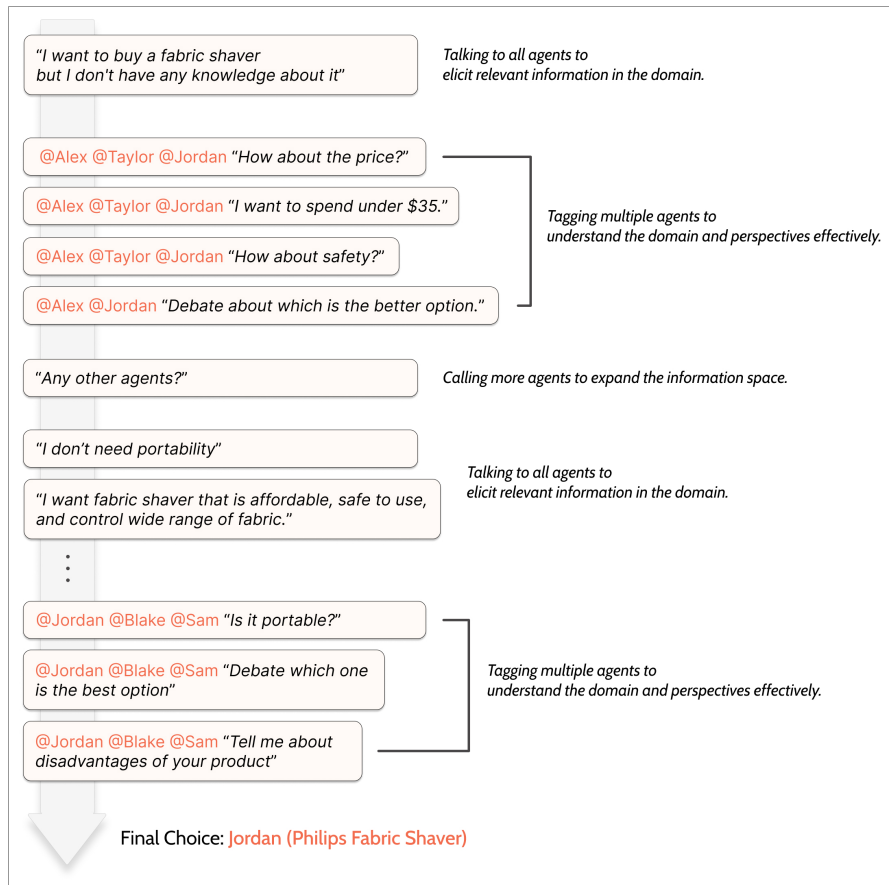


Figure 12: The end-to-end decision-making process of P9.