ARARA: A LLM-based Multi-Agent Development Framework for Conversational Recommender Systems

Fillipe Santos[1,2,3], Juliano Soares[1,2,3], Sildolfo Gomes[1,4], Julio Cesar dos Reis[1,3], and Marcelo S. Reis[1,2,3]

[1] Hub de Inteligência Artificial e Arquiteturas Cognitivas (H.IAAC), Campinas, Brazil
[2] Laboratório de Inteligência Artificial (Recod.ai), Campinas, Brazil
[3] Instituto de Computação, Universidade Estadual de Campinas (UNICAMP), Campinas, Brazil
[4] Instituto de Pesquisas Eldorado, Campinas, Brazil
{fillipesantos, juliano.soares, msreis, jreis}@ic.unicamp.br; sindolfogomes@gmail.com

Abstract. Large Language Models (LLMs) have accelerated the development of Conversational Recommender Systems (CRSs), enabling flexible language understanding and tool-augmented reasoning. However, most LLM-based CRS implementations embed architectural decisions—such as role specialization, memory access, and coordination logic—implicitly within prompts or tightly coupled pipelines. This conflation of reasoning structure with model capacity limits transparency, hinders systematic analysis, and makes architectural extensions ad hoc. This investigation introduces ARARA, a modular, provider-agnostic framework that models CRS as a multi-agent orchestration problem. ARARA defines explicit abstractions for agents, users, modules, orchestration strategies, memory, and reusable skills, establishing governed execution semantics in which conversational behavior emerges from structured role specialization and controlled coordination. To evaluate structural effects, we instantiate ARARA on RecBench+ conversational recommendation benchmarks in the Movie and Book domains, comparing single-LLM baselines with matched multi-agent configurations under identical memory settings. Results show that architectural decomposition and governed routing systematically improve robustness, accuracy, and user satisfaction as reasoning complexity increases, particularly under implicit inference and deceptive inputs. Our findings demonstrate that CRS effectiveness is not solely a function of model capacity, but is critically shaped by architectural design. By elevating coordination, memory, and specialization to first-class architectural primitives, ARARA reframes CRS development as an inspectable and extensible engineering discipline grounded in structured multi-agent orchestration.

Keywords: Conversational Recommender Systems, Multi-Agent Frameworks, Large Language Models.

1 Introduction

Recommender Systems (RSs) play a central role in modern digital platforms, supporting decision-making across domains such as e-commerce and media [13,21]. As user needs become increasingly contextual and time-sensitive, recommendations can no longer be treated as a single prediction step, but as an interactive process shaped by evolving intents and constraints [18].

Conversational Recommender Systems (CRSs) respond to this shift by embedding recommendation within multi-turn dialogue [1,5]. Powered by Large Language Models (LLMs), these systems enable flexible language understanding and tool-augmented reasoning [13,17]. Yet, as conversational complexity increases, reasoning naturally decomposes into multiple interdependent sub-tasks rather than a single generation step.

A practical way to manage this complexity is to distribute responsibilities across specialized *actors*. Instead of relying on a monolithic prompt to implicitly perform all reasoning steps, conversational systems can model retrieval, normalization, verification, ranking, and memory conditioning as distinct components collaborating within a structured interaction [12,3]. Empirical studies indicate that such decomposition improves robustness, supports implicit queries, and mitigates misinformation [19]. These actors can be concretely implemented as *agents* with well-defined responsibilities and interaction boundaries.

Despite this promise, many multi-agent CRS implementations embed architectural decisions implicitly in prompts or tightly coupled pipelines. Questions regarding routing, memory access, tool invocation, and execution flow are often handled ad hoc. General-purpose LLM-agent frameworks provide coordination infrastructure, yet they are not tailored to the structural requirements of conversational recommendation [20,9]. As a result, the system architecture is frequently implicit rather than explicitly modeled and extensible.

To transition from basic recommendation to robust decision support that actively improves decision quality, these underlying structural requirements must first be explicitly defined. This article addresses this gap by establishing a foundational blueprint for CRS architectures, and subsequently introduces *ARARA*, a modular framework for developing multi-agent CRSs. We argue that advancing CRSs requires an explicit architectural foundation that models conversational reasoning as coordinated agent interaction rather than prompt engineering.

The framework *ARARA* is composed of: (**i**) **Agents**: components with clearly defined responsibilities (*e.g.*, retrieval, normalization, verification, ranking), operating as LLM-driven reasoning units; (**ii**) **User**: a specialized agent representing the human participant, interacting through the same message-based interface as other agents. Unlike autonomous agents, the User does not perform LLM-based generation; instead, messages are manually provided through an interaction interface, typically initiating the conversation with a user-issued query; (**iii**) **Modules**: reusable groupings of agents implementing higher-level reasoning patterns; (**iv**) **Orchestrators**: agents responsible for selecting the next component to act. Multiple orchestrators may coexist, and routing strategies may include LLM-based selection, random policies, round-robin strategies, user-driven selection, or custom algorithms; (**v**) **Skills**: plug-and-play capabilities at-

tachable to agents. Examples include *Vision* (image interpretation), *Web Search*, and controlled *Code Execution*.

Agents in ARARA are powered by LLMs. The framework adopts a provider-agnostic client interface, enabling transparent integration of different LLM backends. Current integrations include widely used providers such as OpenAI[5], Anthropic[6], and OpenRouter[7], Groq[8], as well as locally hosted models. New providers can be incorporated by implementing additional client connectors, enabling the system to evolve alongside the rapidly changing LLM ecosystem. ARARA also supports an optional human-in-the-loop mode, distinct from the end user, in which a developer or supervisor may inject input during execution for experimentation, supervision, or debugging.

Memory in ARARA is accessed through skills and currently supports two primary forms: *episodic memory*, which records interaction histories and conversational events, and *shared memory*, which provides a common knowledge space accessible across agents. Both memory structures and their retrieval mechanisms can be customized according to application requirements, allowing developers to define how information is stored, accessed, and updated during execution.

By explicitly modeling users, agents, orchestration, memory, and capabilities, *ARARA* enables conversational behavior to emerge from coordinated interactions among agents. Extension points are exposed at multiple layers—including orchestration strategies, memory mechanisms, skills, and LLM backends—allowing developers to tailor the system to domain-specific CRS deployments.

To evaluate the architectural impact of ARARA, we instantiate it on structured conversational recommendation benchmarks from RecBench+ [7]. We compare single-LLM baselines with ARARA-based multi-agent configurations under matched memory settings, isolating the structural effects of specialization and orchestration across explicit, implicit, and misinformed queries. Rather than claiming state-of-the-art performance, our objective is to quantify the extent to which explicit architectural decomposition and orchestration shape system behavior under varying reasoning conditions.

This study provides the following original contributions:

1. Establish a foundational set of structural requirements for CRS architectures to transition from basic recommendation to robust decision support.
2. Present ARARA as a composable architectural framework for developing multi-agent CRSs, integrating users, agents, orchestrators, memory, skills, and provider-agnostic LLM backends within a unified execution model.
3. Provide a controlled empirical evaluation demonstrating how architectural decomposition and orchestration influence CRS effectiveness under varying reasoning conditions.

The remainder of this article is organized as follows. Section 2 establishes the foundational architectural requirements for CRSs and reviews related work in light of these

[5] https://openai.com/
[6] https://www.anthropic.com/
[7] https://openrouter.ai/
[8] https://groq.com/

requirements; Section 3 details the ARARA framework; Section 4 presents the experimental methodology; Section 5 reports our obtained experimental results; Section 6 discusses our findings and their implications; and Section 7 presents the concluding remarks.

2 Background and Related Work

Conversational recommendation is organized around an iterative recommendation objective rather than general task completion. A CRS must interpret evolving user preferences, connect conversational evidence to a concrete item space, and support the user throughout the recommendation process [8,4].

These requirements serve as architecture-agnostic design principles, supporting both monolithic pipelines and decomposed multi-agent systems. However, making them explicit is important for inspecting, adapting, and systematically comparing CRS architectures [18]. Table 1 summarizes eight CRS-oriented architectural requirements and outlines the design implications necessary to address each one effectively.

Establishing these structural requirements provides a formal baseline for evaluating existing approaches. In the following section, we examine how current LLM-based frameworks address these conversational and architectural demands.

We examine existing LLM-based systems for conversational and agent-oriented applications in light of the architectural requirements established in the previou section 2. Prior work is categorized into single-agent CRSs and general-purpose multi-agent frameworks, evaluating how they address the structural demands of conversational recommendation.

Single-agent systems rely on a single LLM instance augmented with tools, memory, or prompting strategies, without explicit multi-agent coordination. Early CRSs adopt monolithic pipeline architectures combining dialogue management and recommendation backends [5]. Recent LLM-based recommenders integrate language models as rankers or conversational layers over retrieval components [18], improving linguistic flexibility. While these systems target recommendation, they keep reasoning centralized in a single instance, making iterative premise repair, validation, and multi-step elicitation opaque. Tool-augmented frameworks such as Haystack[9] and DSPy[10] support dynamic workflows and external tools, but do not define modular abstractions for coordinated multi-agent orchestration.

Multi-agent systems distribute reasoning across multiple LLM instances or role-based agents. To evaluate these general-purpose frameworks against the specific needs of conversational recommendation, we analyze them across four structural dimensions derived from our foundational requirements:

(i) User and Dialogue Modeling: whether the framework natively models the user as an architectural participant with preference continuity, rather than treating human input as generic task prompts; **(ii) Decision Support Abstractions**: whether the system

[9] https://haystack.deepset.ai/
[10] https://dspy.ai/

Table 1: CRS-oriented architectural requirements and their corresponding design implications.

CRS requirement	Motivation	Architectural Implications
Multi-turn dialogue state and preference continuity	Preferences, constraints, and interaction history accumulate across turns and must inform subsequent recommendations [8,4].	The architecture must explicitly model the user within the interaction loop and provide structured memory abstractions for context storage and injection.
Feedback-driven preference adaptation	User reactions to proposed items, including acceptance, rejection, and critique, should refine subsequent recommendations rather than remain only as dialogue history [10,3].	The system must be capable of parsing feedback signals to dynamically update the criteria used for subsequent retrieval and ranking steps.
Conversational preference elicitation and refinement	Initial requests may be underspecified and require iterative refinement before recommendation [8,10].	The design requires mechanisms to route the conversation toward dedicated elicitation or refinement routines when intents are ambiguous.
Item-grounded retrieval and ranking	Conversational evidence must be connected to items in an external collection rather than relying solely on parametric LLM knowledge [4,18].	The architecture must support explicit tool invocation or skill attachments to execute deterministic queries and external retrieval functions.
Candidate comparison and decision support	When several items are plausible, a CRS should help users assess trade-offs and progress toward a choice rather than return an unstructured list [8,4].	The system must support dedicated reasoning paths that evaluate item trade-offs, focusing on improving overall decision quality rather than merely retrieving candidates.
Recommendation-oriented interaction control	The system must coordinate elicitation, retrieval, ranking, explanation, and continued dialogue as the interaction evolves [4,3].	Execution flow must be governed by an orchestration layer capable of regulating information routing and determining dialogue termination.
Grounded explanation and user control	Users should be able to inspect and refine the basis of a recommendation [4,14].	The architecture must preserve reasoning traces and support the integration of specialized components for explanation or critique generation.
Validation and premise repair	Ambiguous entities and false premises may compromise retrieval and recommendation quality [14,4].	The system must accommodate dedicated verification and error-correction workflows to resolve inaccuracies before items are ranked.

provides modular structures for candidate comparison, feedback adaptation, and decision quality improvement, rather than generic text generation; **(iii) Dynamic Orchestration**: how the execution flow is governed to support the recommendation-oriented interaction control (e.g., static pipelines, graph-based, or dynamic runtime routing); **(iv) Memory and Validation**: whether episodic memory and premise-repair workflows are treated as explicit architectural primitives rather than as auxiliary prompting tricks.

Representative frameworks are summarized in Table 2. **BabyAGI** [16], **CAMEL** [11], and **Multi-Agent Debate** [2] implement role-based or deliberative interactions under largely static coordination schemes. They lack explicit user modeling and provide no abstractions for preference refinement or decision support. **AutoGen** [20] and **LangGraph** [9] provide flexible communication and routing infrastructures. However, as general-purpose orchestration tools, they do not inherently provide the structural decomposition required for CRSs, leaving validation, item-grounding, and user continuity as ad hoc engineering exercises for the developer. **MetaGPT** [6] organizes role-based agents into structured workflows, but is strictly oriented toward software generation.

Table 2: Structural comparison of ARARA with representative multi-agent LLM frameworks based on CRS architectural requirements.

Dimension	**ARARA**	**Baby-AGI**	**CAMEL**	**Multi-Agent Debate**	**AutoGen**	**Lang-Graph**	**MetaGPT**
User & Dialogue Modeling	✓	✗	✗	✗	✗	✗	✗
Decision Support Abstractions	✓	✗	✗	✗	✗	✗	✗
Dynamic Orchestration	dynamic& customizable	static	static	static	flexible	graph-based	static
Memory & Validation	✓	partial	✗	✗	partial	partial	partial

While existing multi-agent approaches demonstrate the benefits of distributed reasoning, they prioritize generic coordination. In contrast, ARARA positions conversational recommendation architecture itself as the primary design object. By building upon foundational multi-agent communication principles [20] and graph-based execution [9,15], ARARA extends these paradigms to fulfill the structural demands of CRSs. It integrates explicit user modeling, decision-support modules, validation structures, and controlled memory into a unified framework, transforming theoretical requirements into concrete, extensible software abstractions.

3 The ARARA Framework

ARARA is a modular engineering framework for developing multi-agent CRSs powered by LLMs. Rather than prescribing a fixed conversational pipeline, ARARA models conversational recommendation as the coordinated interaction of explicitly defined architectural components. Conversational behavior emerges from structured role specialization, explicit orchestration, modular decomposition, and extensible capabilities operating under a unified execution model. Although ARARA exposes reusable multi-agent abstractions, it is designed to make the architectural structure of conversational recommender systems explicit, *e.g.*, modular decomposition, user modeling within the interaction loop, and controlled memory usage—whereas general-purpose agent frameworks provide coordination mechanisms but leave such CRS-specific design decisions implicit, often embedded in prompts or loosely defined pipelines.

Figure 1 presents the architectural flow. Execution begins when a user submits a query. The orchestrator selects the next executable component—either an agent or a module. The selected component processes the message, may invoke external tools or execute code, may access memory through skills, and produces structured output. The output updates the shared state and is routed back to the orchestrator. This message-driven loop continues until a termination condition is reached.

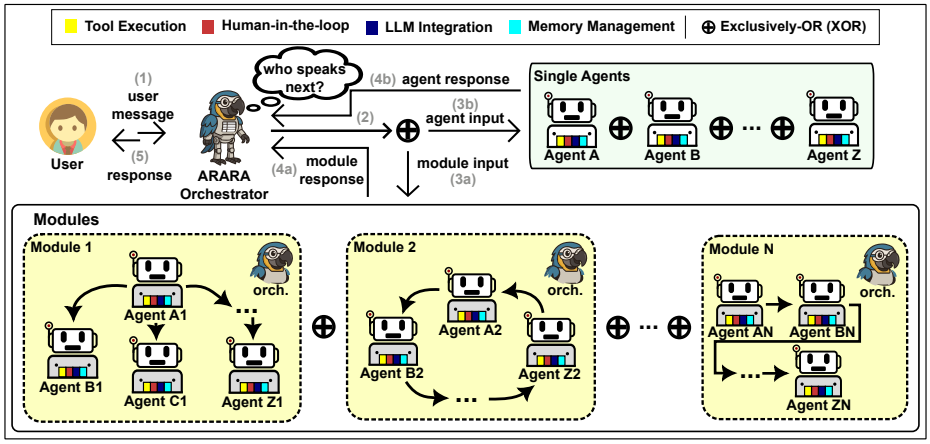

Fig. 1: Overview of the ARARA architecture. A user submits a query, which the orchestrator routes to agents or modules. Agents operate as LLM-driven reasoning units and may access memory, invoke tools, or execute code through skills. The message-driven loop continues under explicit orchestration control until termination.

Agents. Agents are the fundamental reasoning units in ARARA. Each agent embodies a well-defined role—such as retrieval, normalization, verification, ranking, explanation, or evaluation—and transforms structured inputs into structured outputs consumed by other agents or modules. Agents operate under explicit orchestration control and serve as the primary units of reasoning within the system.

An agent's behavior is defined through two complementary mechanisms. First, its role is specified via a system-level instruction that constrains its reasoning scope and behavioral objectives. Second, the agent may invoke external functions or computational routines during execution. For instance, an agent may invoke a classical machine learning model (e.g., a random forest or regression model), execute a deep learning inference routine, or invoke a domain-specific function with structured parameters. The result of such computation can then be reformulated or contextualized by the agent through its LLM-based reasoning process. This design ensures that agents may combine prompt-based reasoning with deterministic or statistical computation within a unified abstraction.

ARARA adopts a provider-agnostic LLM client interface. Agents can be configured with different backend providers—including OpenAI, Anthropic, OpenRouter, Groq, or even locally hosted models—allowing heterogeneous model configurations within the same system. The client layer supports response caching and cost tracking, enabling reproducible experimentation and operational monitoring. New providers can be integrated by implementing additional connectors without modifying the architectural core.

Agents communicate exclusively through message passing. Each message preserves the sender and recipient identities, role, and content, ensuring traceability while maintaining compatibility with chat-based LLM interfaces.

User. Within ARARA, the `User` is modeled as a first-class agent that interacts through the same message-based interface as other components. Execution begins when the user formulates a query. This message is routed through the orchestrator and integrated into the same conversation state as agent-generated responses. Unlike autonomous agents, the `User` does not perform LLM-based generation. Modeling the user architecturally ensures uniform execution semantics and avoids special-case routing or memory handling.

Orchestrator. The orchestrator is modeled as a specialized agent that governs the execution flow. It selects which component—an individual agent or a module—acts next, regulates information routing, and determines termination conditions. Routing policies may rely on deterministic rules, stochastic strategies, LLM-guided decisions, or external decision models. Because the orchestrator follows the same abstraction as other agents, alternative coordination mechanisms can be implemented without altering other architectural layers. This enables experimentation with different dialogue planning strategies while preserving structural consistency.

Modules. Modules provide structural organization by grouping agents into problem-specific units. Each module encapsulates one or more agents designed to address a specific subtask of the conversational recommendation workflow. For example, a *Recommendation Module* may contain agents responsible for candidate retrieval, contextual filtering, and ranking, while a separate *Explanation Module* may include agents dedicated to justification generation, critique, or fidelity analysis. When a module is activated, only the agents within that container participate in the reasoning sequence. This decomposition allows developers to divide complex CRS workflows into well-defined structural units aligned with domain reasoning stages. Modules are flexible design constructs: agents may be reused across modules, and a single query may involve one or multiple modules, depending on the desired decomposition and coordination strategy.

ARARA does not impose a fixed execution structure, leaving these design decisions to the developer. Modules are activated by the system-level orchestrator, which treats each module as a single executable unit. Once activated, control is delegated to the module-level orchestrator, which coordinates the agents contained within that module.

Skills. Skills encapsulate reusable capabilities that can be attached to agents without altering their reasoning roles. This abstraction separates infrastructural functionality from decision logic. Examples include **(i)** Vision, enabling multimodal image interpretation; **(ii)** Web Search, enabling external information retrieval; and **(iii)** Code Execution, enabling controlled execution of generated programs. Skills are plug-and-play and extensible, allowing new capabilities to be implemented and attached to one or multiple agents without structural refactoring.

Memory. Memory is implemented as a specialized skill. The framework supports **(i)** episodic memory, which records chronological interaction histories associated with individual agents; and **(ii)** shared memory, a global knowledge space accessible across agents with ownership-aware write semantics. Both memory types share a common interface that defines storage, retrieval, and context injection. Retrieval strategies and update mechanisms are customizable. Additional memory structures may be introduced without modifying agent logic or orchestration policies.

Runtime Execution and Extensibility. Execution begins when a user submits a query. The orchestrator selects the next agent or module to act. The selected component processes the message, optionally invoking functions, executing code, or accessing memory through skills, and produces structured output. This output updates the shared state and is returned to the orchestrator. The conversational loop terminates under two conditions: **(i)** when an agent signals completion of its reasoning through a configurable termination marker (e.g., "finish"), which must then be confirmed by the user; or **(ii)** when the user explicitly ends the session by issuing the command "exit".

ARARA exposes extensibility at multiple architectural layers. Developers may **(i)** define new agent roles or integrate additional LLM providers; **(ii)** implement alternative orchestration strategies; **(iii)** introduce custom memory structures and retrieval mechanisms; and **(iv)** design new skills encapsulating domain-specific capabilities. Because these extension points are isolated, systems can evolve without structural redesign. ARARA is implemented in Python and released as open source under the MIT License[11].

4 Experimental Methodology

This section presents the procedures of an evaluation to assess how modeling conversational recommendation as a *multi-agent orchestration problem* affects behavior, robustness, and reasoning outcomes. Our goal is not to propose a new recommendation algorithm but to quantify the structural effects of agent decomposition, coordination policies, and controlled memory usage. We instantiate the ARARA architecture on a challenging CRS benchmark and compare it to single-LLM baselines under matched conditions, focusing on queries that require multi-hop inference and correction of incorrect premises.

[11] `https://github.com/fsant0s/arara/`

Subsection 4.1 describes the evaluation datasets and query types. Subsection 4.2 defines the adopted metrics. Subsection 4.3 introduces the single-agent baselines and memory configurations. Subsection 4.4 details the ARARA agents, modules, and orchestration setup used in the experiments.

4.1 Datasets

We adopt the *Movie* and *Book* subsets of RecBench+ [7] as structured and challenging use cases for conversational recommendation. These datasets are designed to stress reasoning, factual consistency, and contextual understanding—properties that are difficult to address with single-agent or purely prompt-based approaches. Each subset extends a traditional collaborative dataset—MovieLens-1M for movies and Amazon-Book for books—by integrating structured knowledge from Wikipedia and Amazon metadata. The resulting Item Knowledge Graph (KG) links items to multiple attributes such as directors, actors, genres, and authors, enabling multi-relational and multi-hop reasoning. RecBench+ defines three condition-based query types that capture distinct reasoning challenges. *Explicit* queries directly specify entities or attributes (e.g., "movies directed by Christopher Nolan"), *Implicit* queries refer to examples that require relational inference (e.g., "movies directed by the director of *The Abyss*"), and *Misinformed* queries intentionally contain factual inaccuracies to simulate user misconceptions (e.g., "books by George Orwell who wrote *Brave New World*"). Each query type is associated with a ground-truth item set extracted from the Item KG, enabling reproducible and semantically grounded evaluation. Following [7], we focus exclusively on these three query types and exclude profile-based queries to maintain controlled reasoning conditions. To balance computational feasibility and representativeness, we sampled 200 queries per reasoning condition across domains. This results in 600 queries per domain (Explicit, Implicit, and Misinformed) and 1,200 queries in total across both domains. These queries serve as a testbed to analyze how different reasoning structures—single-agent versus multi-agent—behave as reasoning complexity increases.

4.2 Evaluation Metrics

Five complementary metrics are adopted. Let R_{rec} denote the set of items recommended within the top-k positions, and R_{true} the ground-truth relevant item set extracted from the Item KG. Recall (R) measures the proportion of relevant items successfully retrieved, capturing how well the system covers the ground-truth set. It is defined as

$$R@k = \frac{|R_{\text{rec}} \cap R_{\text{true}}|}{|R_{\text{true}}|}. \tag{1}$$

Precision (P) measures the concentration of relevance within the recommended list, reflecting how much of what is returned is actually correct. It is defined as

$$P@k = \frac{|R_{\text{rec}} \cap R_{\text{true}}|}{|R_{\text{rec}}|}. \tag{2}$$

Normalized Discounted Cumulative Gain (NDCG) evaluates ranking quality by assigning higher importance to relevant items appearing at top positions, thus capturing ordering sensitivity. It is computed as

$$NDCG@k = \frac{DCG@k}{IDCG@k}, \tag{3}$$

where $DCG@k$ denotes the discounted cumulative gain up to position k, and $IDCG@k$ corresponds to the ideal (maximum) DCG obtained when all relevant items are perfectly ranked at the top positions. Satisfied Ratio (SR) reflects user-level success by verifying whether at least one relevant item appears within the top-k results for a given query. It is defined as

$$SR = \frac{1}{N} \sum_{u=1}^{N} \mathbb{I}(|R^u_{\text{rec}} \cap R^u_{\text{true}}| > 0), \tag{4}$$

where N denotes the total number of evaluated queries. Finally, Fail-To-Recommend (FTR) measures robustness by quantifying the proportion of queries for which the system fails to return any valid recommendation. It is defined as

$$FTR = \frac{N_{\text{fail}}}{N}. \tag{5}$$

A failure is defined as a query for which the system returns no valid item in the top-k results. All ranking-based metrics are computed at a query-dependent cutoff k, where k corresponds to the number of ground-truth relevant items associated with the input query (bookCount/movieCount). In our benchmark, k varies between 2 and 5 across queries. Higher values of R, P, NDCG, and SR indicate stronger recommendation effectiveness and user-facing success, whereas lower FTR indicates greater robustness under increasing reasoning difficulty. All reported metric values correspond to macro-averages computed across the evaluated queries. All experiments were executed once per query without stochastic repetition.

4.3 Baselines

We adopt three representative LLMs from RecBench+ [7] as structural baselines: **GPT-4o**, **DeepSeek-V3**, and **Claude 3.5 Sonnet**. These models are not treated as competing algorithms, but as reference single-agent configurations against which ARARA-based multi-agent instantiations are compared. Each model is evaluated under two settings: *with memory* (*WM*) and *without memory* (*W/O M*). In the memory-enabled setting, the model receives the user's ten most recent interactions (`history_line`) as context. In the memory-free setting, reasoning is based solely on the current query. This design allows us to isolate the effect of *agent orchestration and task decomposition* from the underlying language model capabilities.

4.4 Agents, Modules and Orchestration Setup

We use the *ARARA* architecture to evaluate how modeling conversational recommendation as a *multi-agent orchestration problem* affects behavior under challenging reasoning scenarios. For each underlying LLM, we compare: (i) a **single-agent baseline**

(LLM + RecBench+ prompt) and (ii) a **multi-agent ARARA instantiation** embedding the same LLM within ARARA's agent abstraction, where role-specific system configurations define bounded reasoning responsibilities under orchestration control. Both systems are evaluated under matched memory settings (*WM* / *W/O M*), so differences isolate reasoning structure and coordination rather than model capacity or additional supervision.

Agents and Memory. All ARARA agents implement the framework's unified agent abstraction, are instantiated through role-specific system configurations that define their reasoning behavior, and share the same LLM client within each run. For each model (**GPT-4o**, **DeepSeek-V3**, **Claude 3.5 Sonnet**), we instantiate two ARARA variants: *ARARA-*$_{WM}$ and *ARARA-*$_{W/O\,M}$. To isolate memory effects at the decision stage, only *Recommender* agents access the user's ten most recent interactions (`history_line`) through the episodic memory skill; all other agents operate statelessly. Agents are explicitly specialized into bounded reasoning functions (e.g., retrieval, normalization, validation, correction, ranking), enabling complex queries to be decomposed into interpretable steps.

Modules. Agents are grouped into three modules aligned with RecBench+ query types: *Explicit*, *Implicit*, and *Misinformed*. Exactly one module is activated per query under a mutually exclusive (XOR) selection policy. **Explicit module** consists of a `RetrieverAgent` for deterministic candidate retrieval and a `RecommenderAgent` for ranking (optionally conditioned on `history_line`). **Implicit module** decomposes multi-hop inference into `TitleNormalizer`, `CommonAttributeExtractor`, `RelationTypeDetector`, `MultiHopRetriever`, and `RecommenderAgent`. **Misinformed module** introduces explicit validation and correction via `EntityResolverAgent`, `FactCheckerAgent`, `FactCorrectionAgent`, `RetrieverAgent`, and `RecommenderAgent`, preventing incorrect premises from propagating into retrieval and ranking.

Orchestration. We use an **LLM-guided routing** strategy in which a prompt-based decision step selects one module (Explicit/Implicit/Misinformed) per query based on the module descriptions. Upon completion of the selected module, the orchestrator returns the final ranked recommendations to the user (Figure 2).

This setup is designed to attribute observed gains to *structural properties*—role specialization, governed coordination, and controlled memory—rather than algorithmic novelty. In the present experimental setup, agent-level termination signals were not enabled. Execution proceeds until the user explicitly issues the command `exit`, which terminates the conversational loop and concludes the experimental run. All LLM APIs were executed with a temperature of 0.0 to eliminate stochastic variability, and identical decoding configurations were used across all experimental conditions. All agent definitions, system prompts, orchestration configurations, and experimental settings are publicly available on GitHub for full reproducibility[12].

[12] `https://github.com/fsant0s/arara_experiments`

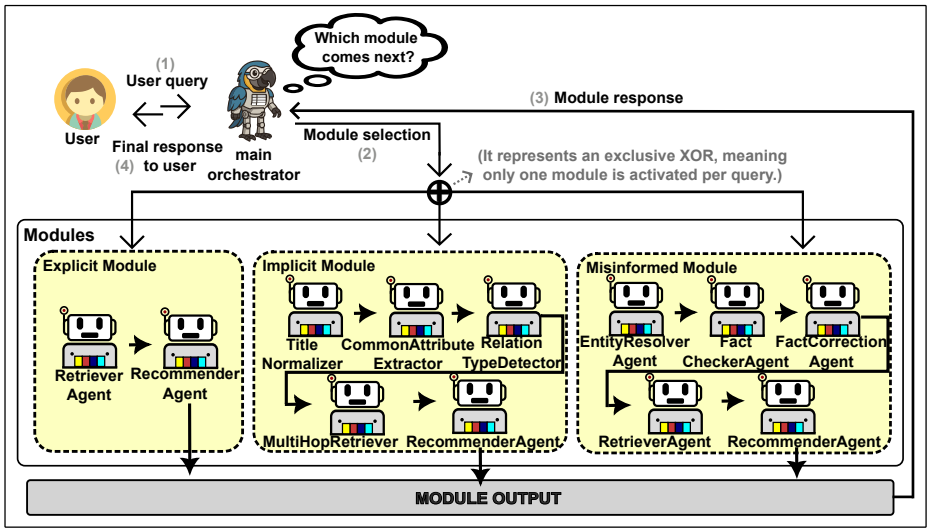

Fig. 2: Overview of the instantiated ARARA-based architecture. A prompt-based, LLM-guided orchestrator routes each query to exactly one module (Explicit, Implicit, or Misinformed), which coordinates specialized agents for structured reasoning.

5 Experimental Results

This section evaluated ARARA-based multi-agent instantiations from a structural perspective, focusing on how reasoning decomposition and orchestration affected conversational recommendation behavior. Rather than targeting absolute performance gains, we examine how multi-agent coordination affects retrieval quality, ranking quality, robustness to misinformation, and user-facing satisfaction as reasoning complexity increases.

The evaluation was organized into three analyses. First, Subsection 5.1 reports results for the *Movie* domain. Second, Subsection 5.2 presents results for the *Book* domain using the same protocol. Together, these two subsections compared single-agent LLM baselines against their ARARA-based multi-agent counterparts across three reasoning conditions (*Explicit*, *Implicit*, and *Misinformed*). Third, Subsection 5.3 analyzed orchestration behavior itself, examining routing accuracy and module selection patterns across both domains. Recall that, for each underlying LLM baseline, two configurations were evaluated: *with memory* (*WM*) and *without memory* (*W/O M*). ARARA instantiations embedded the same LLM backends within the previously described modular architecture (*Explicit*, *Implicit*, and *Misinformed* modules coordinated by an orchestrator), preserving identical memory settings. Consequently, observed differences were interpreted as structural effects induced by agent specialization, governed coordination, and controlled memory usage, rather than by changes in model capacity.

5.1 Movie Domain Results

Table 3 reports the results for the *Movie* domain across three reasoning conditions: *Explicit* (Easy), *Implicit* (Medium), and *Misinformed* (Hard). Results were shown for base LLM configurations and their ARARA-based counterparts, under both memory settings.

In the *Explicit* condition, base models and ARARA configurations exhibited comparable behavior, with modest gains in recall and user satisfaction. For example, GPT-4o$_{(WM)}$ obtained R = 0.36 and SR = 0.72, while ARARA GPT-4o$_{(WM)}$ reached R = 0.40 and SR = 0.82. Comparable trends were observed across DeepSeek-V3 and Claude-3.5-Sonnet families, indicating limited separation between single-agent and multi-agent settings for well-specified queries. In the *Implicit* condition, ARARA configurations consistently yielded higher user satisfaction and precision. Under GPT-4o$_{(WM)}$, SR increased from 0.33 to 0.72, while precision increased from 0.16 to 0.28. For DeepSeek-V3$_{(WM)}$, ARARA maintained comparable recall while improving SR from 0.65 to 0.69. Similar patterns were observed across models and memory settings. The *Misinformed* condition exhibited the largest differences. Single-agent baselines showed higher failure rates, whereas ARARA configurations reduced FTR and increased SR. For instance, GPT-4o$_{(W/O\ M)}$ recorded FTR = 0.27, compared to 0.16 for ARARA GPT-4o$_{(W/O\ M)}$, while SR increased from 0.10 to 0.35. For Claude-3.5-Sonnet$_{(W/O\ M)}$, ARARA increased SR from 0.12 to 0.28 while maintaining comparable robustness levels.

Across several configurations, ARARA exhibited trade-offs between recall and satisfaction. For example, ARARA DeepSeek-V3$_{(W/O\ M)}$ achieved lower recall (R = 0.34 vs. 0.40) but higher satisfaction (SR = 0.87 vs. 0.62).

5.2 Book Domain Results

Table 4 presents results for the *Book* domain using the same evaluation protocol and reasoning conditions. In the *Explicit* condition, ARARA and base models exhibited different recall profiles. While base models achieved lower recall, ARARA configurations substantially increased recall, often at the cost of higher failure rates. For example, GPT-4o$_{(WM)}$ achieved R = 0.14, while ARARA GPT-4o$_{(WM)}$ reached R = 0.62, with a higher FTR. Similar recall–failure trade-offs were observed across other models. In the *Implicit* condition, ARARA configurations consistently improved ranking and satisfaction metrics. ARARA DeepSeek-V3$_{(WM)}$ increased NDCG from 0.20 to 0.50 and SR from 0.04 to 0.09. Comparable improvements were observed across models and memory settings. The *Misinformed* condition again showed larger separations. ARARA Claude-3.5-Sonnet$_{(W/O\ M)}$ achieved R = 0.49 and SR = 0.15, compared to R = 0.05 and SR = 0.09 for the corresponding baseline. In several cases, ARARA configurations traded higher recall and user satisfaction for increased failure rates, particularly under misinformed queries.

5.3 Orchestration Analysis Results

To better interpret the differences observed in the results from the previous sections, we carried out an orchestration analysis to isolate the effect of the orchestrator's routing

Table 3: Quantitative results for the *Movie* domain across the three reasoning conditions — *Explicit* (Easy), *Implicit* (Medium), and *Misinformed* (Hard). Each block reports five metrics: FTR ↓, *R* ↑, *P* ↑, NDCG ↑, and SR ↑. Results compare base LLMs and their ARARA-enhanced counterparts, under both memory configurations: $_{(\text{WM})}$ for *with memory* and $_{(\text{W/O M})}$ for *without memory*. Bold values indicate the best result per condition.

	Explicit (Easy)					**Implicit (Medium)**					**Misinformed (Hard)**				
Model	FTR↓	R↑	P↑	NDCG↑	SR↑	FTR↓	R↑	P↑	NDCG↑	SR↑	FTR↓	R↑	P↑	NDCG↑	SR↑
GPT-4o$_{WM}$	**.03**	.36	.33	.36	.72	**.01**	.23	.16	.21	.33	.26	**.05**	**.03**	**.03**	.22
+ ARARA	.09	**.40**	**.43**	**.43**	**.82**	.15	**.29**	**.28**	**.28**	**.72**	**.18**	.04	**.03**	**.03**	**.24**
GPT-4o$_{W/O}$	**.01**	**.40**	**.30**	**.39**	.71	**.02**	.22	.14	.20	.30	.27	.03	.01	.02	.10
+ ARARA	.07	.26	.28	.27	**.85**	.16	**.28**	**.28**	**.28**	**.71**	**.16**	**.06**	**.06**	**.06**	**.35**
DeepSeek-V3$_{WM}$	**.00**	.44	.22	.39	.65	**.00**	**.26**	.09	.20	.65	**.00**	.06	.03	.05	.25
+ ARARA	.02	**.45**	**.46**	**.48**	**.86**	.12	**.26**	**.26**	**.28**	**.69**	.15	**.09**	**.08**	**.09**	**.26**
DeepSeek-V3$_{W/O}$	**.00**	**.40**	.18	**.35**	.62	**.00**	.26	.09	.20	.21	**.01**	**.07**	.02	.05	.10
+ ARARA	.02	.34	**.35**	**.35**	**.87**	.15	**.28**	**.28**	**.28**	**.69**	.07	**.07**	**.07**	**.07**	**.27**
Claude-3.5$_{WM}$	**.01**	.36	.20	.31	.65	**.01**	.16	.10	.12	.09	**.02**	.04	.06	.03	.21
+ ARARA	.06	**.42**	**.42**	**.45**	**.87**	.13	**.23**	**.23**	**.24**	**.59**	.05	**.08**	**.08**	**.09**	**.23**
Claude-3.5$_{W/O}$	**.00**	.31	.26	**.44**	.65	**.00**	**.21**	.12	**.23**	.24	.08	.07	.03	.05	.12
+ ARARA	.02	**.38**	**.38**	.39	**.79**	.09	.17	**.17**	.18	**.61**	**.04**	**.08**	**.08**	**.07**	**.28**

Table 4: Quantitative results for the *Book* domain across the three reasoning conditions — *Explicit* (Easy), *Implicit* (Medium), and *Misinformed* (Hard). Each block reports five metrics: FTR ↓, R ↑, P ↑, NDCG ↑, and SR ↑. Results compare base LLMs and their ARARA-enhanced counterparts, under both memory configurations: $_{(\text{WM})}$ for *with memory* and $_{(\text{W/O M})}$ for *without memory*. Bold values indicate the best result per condition.

	Explicit (Easy)					**Implicit (Medium)**					**Misinformed (Hard)**				
Model	FTR↓	R↑	P↑	NDCG↑	SR↑	FTR↓	R↑	P↑	NDCG↑	SR↑	FTR↓	R↑	P↑	NDCG↑	SR↑
GPT-4o$_{WM}$	**.01**	.14	.07	.13	.06	**.01**	.12	.07	.11	.01	**.26**	.05	.03	.05	.00
+ ARARA	.15	**.62**	**.55**	**.62**	**.10**	.27	**.45**	**.42**	**.44**	**.06**	.28	**.42**	**.37**	**.41**	**.02**
GPT-4o$_{W/O}$	**.02**	.22	.09	.19	.05	**.01**	.14	.07	.13	.04	.39	.03	.02	.03	**.00**
+ ARARA	.12	**.56**	**.50**	**.54**	**.20**	.21	**.43**	**.40**	**.42**	**.15**	**.22**	**.34**	**.31**	**.32**	**.00**
DeepSeek-V3$_{WM}$	**.01**	.17	.06	.14	.05	**.02**	.18	.07	.20	.04	.10	.01	.01	.01	.00
+ ARARA	**.01**	**.57**	**.42**	**.63**	**.13**	.03	**.46**	**.32**	**.50**	**.09**	**.03**	**.28**	**.17**	**.30**	**.06**
DeepSeek-V3$_{W/O}$	**.00**	.26	.08	.21	.00	.04	.21	.07	.16	.00	**.08**	.00	.00	.00	.00
+ ARARA	.01	**.56**	**.43**	**.60**	**.27**	**.01**	**.37**	**.25**	**.39**	**.15**	.12	**.30**	**.21**	**.33**	**.07**
Claude-3.5$_{WM}$	**.00**	.16	.05	.15	.03	**.01**	.22	.05	.18	.02	**.04**	.09	.03	.06	.04
+ ARARA	.21	**.60**	**.52**	**.61**	**.16**	.03	**.41**	**.31**	**.43**	**.08**	.12	**.42**	**.33**	**.43**	**.12**
Claude-3.5$_{W/O}$	**.01**	.18	.06	.16	**.02**	**.01**	.18	.05	.16	.00	**.01**	.05	.02	.04	.09
+ ARARA	.07	**.34**	**.29**	**.33**	.00	.08	**.40**	**.33**	**.41**	**.08**	.22	**.49**	**.42**	**.48**	**.15**

decisions. In ARARA, downstream module specialization could only improve performance if the query was first assigned to the appropriate reasoning module. Therefore, evaluating routing accuracy allowed us to disentangle whether observed gains or failures stem from module specialization itself or from errors in module selection. Table 5 reports routing accuracy for representative ARARA configurations across the *Movie* and *Book* domains. One representative per model family was analyzed, selected based on robustness under *Misinformed* queries, factual reliability (low FTR), and user satisfaction (high SR), to avoid redundant analysis across highly similar configurations. In the *Movie* domain (Table 5a), the orchestrator achieved consistently high accuracy for *Explicit* and *Implicit* queries across all representatives. Most errors were concentrated in the *Misinformed* condition, where deceptive inputs were occasionally misclassified as simpler reasoning types, which suggested increased ambiguity in deceptive queries. In the *Book* domain (Table 5b), routing remained comparatively reliable for *Explicit* queries but showed increased overlap between *Implicit* and *Misinformed* reasoning. This reflected higher linguistic ambiguity and weaker intent signals in the domain, which led to a tendency to overestimate reasoning complexity in challenging cases.

Table 5: Orchestration accuracy across Movie (a) and Book (b) domains, reporting inferred reasoning conditions (Explicit, Implicit, Misinformed), with 200 queries per condition and model.

(a) Movie domain

Model		Condition		
		Exp.	**Imp.**	**Mis.**
ARARA GPT-4o$_{(W/O\ M)}$	**Exp.**	197	0	4
	Imp.	2	200	15
	Mis.	1	0	181
ARARA DeepSeek-V3$_{(W/O\ M)}$	**Exp.**	200	0	126
	Imp.	0	200	21
	Mis.	0	0	53
ARARA Claude-3.5-Sonnet$_{(W/O\ M)}$	**Exp.**	199	0	29
	Imp.	0	200	1
	Mis.	1	0	170

(b) Book domain

Model		Condition		
		Exp.	**Imp.**	**Mis.**
ARARA GPT-4o$_{(W/O\ M)}$	**Exp.**	198	30	1
	Imp.	0	117	0
	Mis.	2	57	199
ARARA DeepSeek-V3$_{(WM)}$	**Exp.**	200	130	73
	Imp.	0	69	0
	Mis.	0	1	127
ARARA Claude-3.5-Sonnet$_{(WM)}$	**Exp.**	199	14	10
	Imp.	0	180	0
	Mis.	1	6	190

6 Discussion

Domain-level effectiveness. Across domains, ARARA reshapes CRS behavior by distributing reasoning across coordinated validation and refinement stages. In the *Movie* domain, it consistently improves recall, precision, and ranking quality through structured multi-agent collaboration, while in the *Book* domain, configurations like **ARARA Claude-3.5-Sonnet(W/O M)** maintain high recall and ranking quality, demonstrating that modular orchestration preserves semantic coherence even in denser relational contexts. These observations suggest that ARARA's effectiveness stems less from model

capacity and more from its structured architectural design, which comprises traceable, interdependent stages that support transparency, reliability, and adaptive behavior across domains.

Reasoning across difficulty levels. For *Explicit* queries, baseline models and ARARA configurations behaved similarly due to the deterministic nature of retrieval. However, under *Implicit* and *Misinformed* conditions, where ambiguity, relational inference, and correction are required, single-agent reasoning exhibits increased error propagation under complex conditions. ARARA mitigates this effect by distributing reasoning across specialized modules that explicitly verify, correct, and recontextualize information. As a result, higher task difficulty is transformed into structured collaboration rather than compounded error, demonstrating that coordinated reasoning scales more reliably than isolated generation as reasoning complexity increases.

Memory analysis. Across both domains, memory-enhanced configurations improved contextual recall and user satisfaction under *Implicit* and *Misinformed* conditions by enabling the reuse of verified context during ranking. At the same time, memory-free configurations retained comparable stability, indicating that orchestration and modular validation alone already sustain reliable reasoning. These findings suggest that memory primarily supports continuity—reducing redundant reasoning and strengthening referential grounding—while its marginal gains diminish when coordination mechanisms are well calibrated. This reinforces ARARA's design goal of treating memory as an optional, controlled component rather than a mandatory source of performance gains.

Orchestrator effectiveness. In the *Movie* domain, routing accuracy was generally high, with minor confusion between *Misinformed* and *Explicit* queries, while in the *Book* domain, ambiguity was concentrated in *Implicit* cases due to dense semantic relations. Overall, ARARA effectively governs reasoning through explicit routing and coordination, though prompt-based routing remains sensitive to linguistic overlap. Rather than signaling structural flaws, these behaviors highlight natural extension points for the orchestration layer, such as adaptive routing thresholds, without compromising interpretability.

Overall analysis. Taken together, the results demonstrated that ARARA reshapes CRS by structuring how reasoning unfolds rather than modifying the underlying language models. While maintaining comparable behavior in simpler queries, it improves stability and robustness as reasoning becomes more implicit or misinformed. These effects highlight orchestration—through role specialization, explicit routing, and controlled memory—as the framework's core contribution.

Limitations and outlook. Routing remains sensitive to linguistic ambiguity, and memory usage is currently static. Moreover, ARARA does not aim to optimize absolute performance, but to expose the structural effects of architectural design. As an extensible framework, ARARA provides a foundation for advancing multi-agent reasoning, adaptive orchestration, and memory strategies in CRS.

7 Conclusion

As conversational recommendation increasingly relies on LLMs, architectural decisions are often embedded implicitly within prompts or monolithic pipelines, limiting trans-

parency and systematic analysis. This study introduced **ARARA**, a modular framework that models conversational recommendation as an explicit multi-agent orchestration problem, treating agents, modules, memory, and routing as composable architectural components. Rather than proposing a new recommendation algorithm, ARARA provides a structured foundation for decomposing reasoning into specialized roles coordinated under governed execution. Empirical results across two domains showed that such architectural decomposition improves stability and robustness as reasoning complexity increases, particularly under *Implicit* and *Misinformed* queries. While ARARA demonstrates clear benefits in structuring conversational reasoning, it also introduces additional architectural complexity and design overhead, highlighting the need for future research on adaptive orchestration mechanisms and more efficient memory management strategies to ensure scalability and broader applicability. In light of these trade-offs, architecture—rather than model capacity—emerges as a central design dimension in CRSs. Future work includes comparisons with alternative multi-agent frameworks to evaluate ARARA along additional dimensions such as adaptability and extensibility.

Acknowledgments. This project was supported by the Ministry of Science, Technology, and Innovation of Brazil, with resources granted by the Federal Law 8.248 of October 23, 1991, under the PPI-Softex. The project was coordinated by Softex and published as Intelligent agents for mobile platforms based on Cognitive Architecture technology [01245.003479/2024-10]. We also thank the partial support by the National Council for Scientific and Technological Development (CNPq), Brazil (grant #301337/2025-0), and São Paulo Research Foundation (FAPESP) (grant #2025/10403-2).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Deng, Y., Liao, L., Lei, W., Yang, G.H., Lam, W., Chua, T.S.: Proactive conversational AI: A comprehensive survey of advancements and opportunities. ACM Trans. Inf. Syst. **43**(3) (Mar 2025), https://doi.org/10.1145/3715097
2. Du, Y., Li, S., Torralba, A., Tenenbaum, J.B., Mordatch, I.: Improving factuality and reasoning in language models through multiagent debate. In: Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., Berkenkamp, F. (eds.) Proceedings of the 41st International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 235, pp. 11733–11763. PMLR (21–27 Jul 2024), https://proceedings.mlr.press/v235/du24e.html
3. Fang, J., Gao, S., Ren, P., Chen, X., Verberne, S., Ren, Z.: A multi-agent conversational recommender system. arXiv preprint arXiv:2402.01135 (2024)
4. Friedman, L., Ahuja, S., Allen, D., Tan, Z., Sidahmed, H., Long, C., Xie, J., Schubiner, G., Patel, A., Lara, H., et al.: Leveraging large language models in conversational recommender systems. arXiv preprint arXiv:2305.07961 (2023)
5. Gao, C., Lei, W., He, X., de Rijke, M., Chua, T.S.: Advances and challenges in conversational recommender systems: A survey. AI Open **2**, 100–126 (2021). https://doi.org/10.1016/j.aiopen.2021.06.002

6. Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Wang, J., Zhang, C., Wang, Z., Yau, S., Lin, Z., Zhou, L., Ran, C., Xiao, L., Wu, C., Schmidhuber, J.: MetaGPT: Meta programming for a multi-agent collaborative framework. In: Kim, B., Yue, Y., Chaudhuri, S., Fragkiadaki, K., Khan, M., Sun, Y. (eds.) International Conference on Learning Representations. vol. 2024, pp. 23247–23275 (2024), https://proceedings.iclr.cc/paper_files/paper/2024/file/6507b115562bb0a305f1958ccc87355a-Paper-Conference.pdf
7. Huang, J., Wang, S., Ning, L., Fan, W., Wang, S., Yin, D., Li, Q.: Towards next-generation recommender systems: A benchmark for personalized recommendation assistant with LLMs. In: Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining. p. 217–226. WSDM '26, Association for Computing Machinery, New York, NY, USA (2026), https://doi.org/10.1145/3773966.3777954
8. Jannach, D., Manzoor, A., Cai, W., Chen, L.: A survey on conversational recommender systems. ACM Comput. Surv. **54**(5) (May 2021). https://doi.org/10.1145/3453154
9. LangChainAI: LangGraph: Compositional graph framework for LLM-orchestrated workflows. https://github.com/langchain-ai/langgraph (2024), version 1.2.8. Accessed: 2026-07-08
10. Lei, W., Zhang, G., He, X., Miao, Y., Wang, X., Chen, L., Chua, T.S.: Interactive path reasoning on graph for conversational recommendation. In: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 2073–2083 (2020). https://doi.org/10.1145/3394486.3403258
11. Li, G., Hammoud, H., Itani, H., Khizbullin, D., Ghanem, B.: CAMEL: Communicative agents for "mind" exploration of large language model society. Advances in neural information processing systems **36**, 51991–52008 (2023)
12. Li, X., Wang, S., Zeng, S., Wu, Y., Yang, Y.: A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges. Vicinagearth **1**(1), 9 (2024). https://doi.org/10.1007/s44336-024-00009-2
13. Lin, J., Dai, X., Xi, Y., Liu, W., Chen, B., Zhang, H., Liu, Y., Wu, C., Li, X., Zhu, C., et al.: How can recommender systems benefit from large language models: A survey. ACM Transactions on Information Systems **43**(2), 1–47 (2025). https://doi.org/10.1145/3678004
14. Moon, S., Shah, P., Kumar, A., Subba, R.: OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs. In: Korhonen, A., Traum, D., Màrquez, L. (eds.) Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 845–854. Association for Computational Linguistics, Florence, Italy (Jul 2019). https://doi.org/10.18653/v1/P19-1081, https://aclanthology.org/P19-1081/
15. Moura, J.: CrewAI: A framework for orchestrating role-playing autonomous AI agents. https://github.com/crewAIInc/crewAI (2024), version 1.15.2. Accessed: 2026-07-08
16. Nakajima, Y.: BabyAGI. https://github.com/yoheinakajima/babyagi (2023), accessed: 2026-02-19
17. OpenAI: GPT-4 Technical Report. arXiv preprint arXiv:2303.08774 (2023)
18. Peng, Q., Liu, H., Huang, H., Yang, J., Yang, Q., Shao, M.: A survey on LLM-powered agents for recommender systems. In: Christodoulopoulos, C., Chakraborty, T., Rose, C., Peng, V. (eds.) Findings of the Association for Computational Linguistics: EMNLP 2025. pp. 11574–11583. Association for Computational Linguistics, Suzhou, China (Nov 2025). https://doi.org/10.18653/v1/2025.findings-emnlp.620, https://aclanthology.org/2025.findings-emnlp.620/
19. Wang, Z., Yu, Y., Zheng, W., Ma, W., Zhang, M.: MACRec: A multi-agent collaboration framework for recommendation. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 2760–2764. SIGIR '24, Association for Computing Machinery, New York, NY, USA (2024). https://doi.org/10.1145/3626772.3657669

20. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A.H., White, R.W., Burger, D., Wang, C.: AutoGen: Enabling next-gen LLM applications via multi-agent conversation. arXiv preprint arXiv:2308.08155 (2023). https://doi.org/10.48550/arXiv.2308.08155
21. Zhang, Y., Chen, X.: Explainable recommendation: A survey and new perspectives. Foundations and Trends in Information Retrieval **14**(1), 1–101 (2020). https://doi.org/10.1561/1500000066