



A Tutorial on Recent Advances in Generative Conversational Recommender Systems



Thomas E. Kolb¹



Ahmadou Wagne¹



Ashmi Banerjee²



Fatemeh Nazary³



Julia Neidhardt¹



Yashar Deldjoo³



Tommaso Di Noia³

¹ TU Wien, Austria

² TU Munich, Germany

³ Polytechnic University of Bari, Italy

Agenda

- Introduction
- Core Systems & Components
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- Simulation
- Evaluation
- Open Challenges & Future Directions

Background

The tutorial is based on our upcoming survey paper called:

Advancements in Conversational Recommender Systems Using Generative Models: A Systematic Literature Review

AHMADOU WAGNE, TU Wien, Austria

THOMAS ELMAR KOLB, TU Wien, Austria

ASHMI BANERJEE, Technical University of Munich, Germany

FATEMEH NAZARY, Polytechnic University of Bari, Italy

JULIA NEIDHARDT, TU Wien, Austria

YASHAR DELDJOO, Polytechnic University of Bari, Italy

Link: <https://recsys-lab.at/gen-conv-recsys-tutorial>



Agenda

- **Introduction**
- Core Systems & Components
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- Simulation
- Evaluation
- Open Challenges & Future Directions



Yashar Deldjoo



Fatemeh Nazary

Introduction




Conversational recommendation

powered by Generative AI (Gen-CRS)



Before ...

Need a gift from
Prague for my wife



Information
need

Exploration



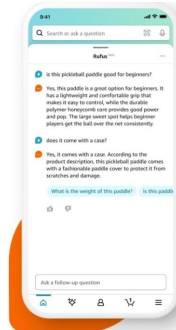

Product Search

Problem

The customer has to go multiple websites and places to find what they want to buy!



Unified Conversation + Recommendation Experience




Information
need

Shopping through Rufus
conversational LLM


What is a Conversational Recommender System (CRS)?

A Conversational Recommender System (CRS) is an interactive AI framework that engages users in **multi-turn dialogue** to **elicit - explicit preferences** and provide **personalized recommendations**.




Advantages

- Assist decision-making and information-seeking
- Support product space exploration e.g., discover unexpected but relevant items
- Elicit users' nuanced or hidden preferences




Search, Interactive Recommendation, RecSys

Traditional paradigms for information-seeking:
Search (pull) or **Recommendation (push)**



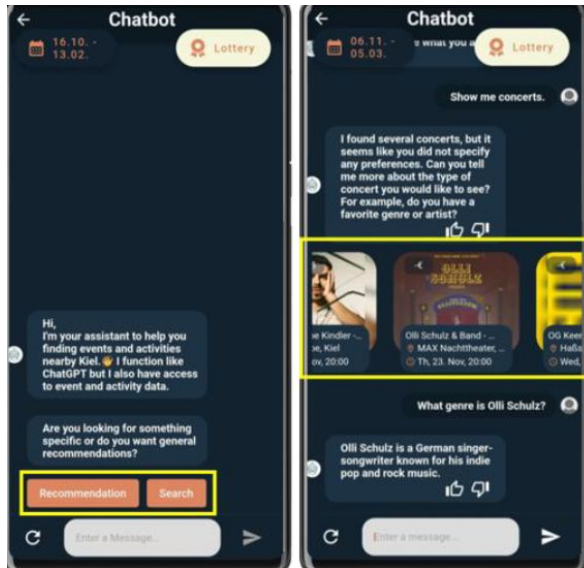
Key Characteristics

- **Multi turn interaction** – conversation persists across rounds.
- **Mixed initiative** – system and user alternate turns and roles.
- **Natural language** – voice or text as the primary interface.
- **Dynamic preference elicitation** – ask, refine and adapt



Conversational AI - High Level Categorization

1. Goal-driven (task-oriented)




Goal-driven (task-oriented): aiming to assist users to complete specific tasks

- Conversational information access: tasks with underlying information need, which can be satisfied through a conversation
- Includes task of **search**, **recommendation**, and QA (boundaries often blurred)

Conversational AI - High Level Categorization

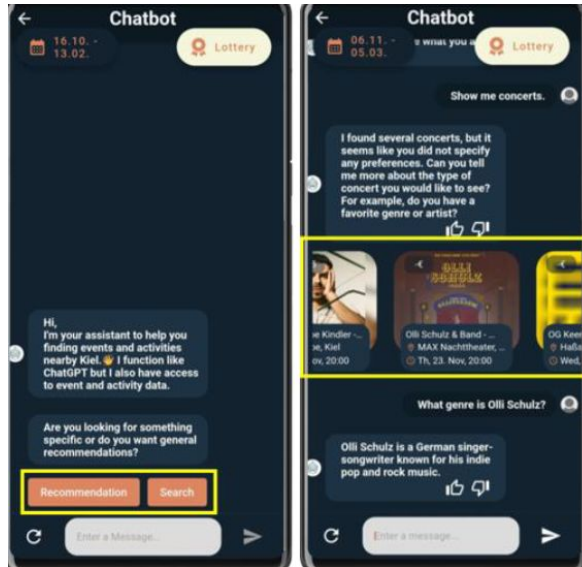
2. Non Goal-driven (chatbots)

Non Goal-driven (chatbots): aiming to carry out an extended conversation (**“chit-chat”**) usually with the purpose of **entertainment**.




Conversational AI - High Level Categorization

Goal-driven (task-oriented)



Non Goal-driven (chatbots)



Our focus

Overview of Conversational AI



CRS (Conversational RecSys)




CQA (Conversational Q&A)



CS (Conversational Search)



Social Chatbot



Gen-CRS

CRS vs Conversational Search

- **Common:** rank or surface relevant options in multi-turn interaction
- **Key difference:** CRS builds a user model & personalised; search focuses on query understanding & retrieval

CRS vs Conversational QA

- **Common:** dialogue to resolve information needs
- **Key difference:** QA answers factual questions; CRS elicits preferences and recommends subjective items

CRS vs Social Chatbot

- **Common:** free-form conversational exchange
- **Key difference:** CRS is goal-oriented (task success); social chat aims for open-ended chit-chat & engagement

CRS vs Gen-CRS (LLM-powered)

- **Common:** goal-oriented recommendation via dialogue
- **Key difference:** Traditional CRS uses rules/templates & separate rankers; Gen-CRS uses LLMs for free-form NLG, mixed-initiative, and tool use (with grounding to reduce hallucination)

Evolution of Conversational RS

2017-2018

Neural CRS Emergence

First neural conversational recommenders appear; separate dialogue and rec modules.

2018-2021

Evolution

Research integrates knowledge graphs & critique-based recommendation.

2022-2023

LLM Revolution

GPT-3 & ChatGPT enable generative dialogue & zero-shot recommendation.

2023-2025

GenCRS Explosion

Unified, modular & agentic architectures flourish with LLMs & tools.

Some traditional approaches...

Traditional approaches rarely involved “conversation” as we might normally think of it:

1. **Thompson et al., 2004** (query refinement):

Elicits users’ preferences and constraints with regard to **item attributes**;

Example of a user model




User Name		Homer					
Attributes	w_i	Values and probabilities					
Cuisine	0.4	Italian	French	Turkish	Chinese	German	English
		0.35	0.2	0.25	0.1	0.1	0.0
Price Range	0.2	one	two	three	four	five	
		0.2	0.3	0.3	0.1	0.1	
...	...	...					
Parking	0.1	Valet		Street		Lot	
		0.5		0.4		0.1	
Item Nbr.		0815	5372	7638	...		6399
Accept/Present		23 / 25	10 / 19	33 / 36	...		12 / 23

Some traditional approaches...

Traditional approaches rarely involved “conversation” as we might normally think of it:

2. Mahmood and Ricci, 2009 (reinforcement learning): **Queries** users about **recommendation attributes** during each round; learns a **policy** to choose queries to efficiently yield a desirable recommendation



Some traditional approaches...

Traditional approaches rarely involved “conversation” as we might normally think of it:

3. Christakopoulou et al., 2016 (iterative recommendation): Collects feedback about recommended items in order to **iteratively learn user preferences**; explores various query strategies to elicit preferences quickly




Table 3: Question selection strategies evaluated.

Greedy: $j^* = \arg \max_j y_{ij}$

A trivial *exploit*-only strategy: Select the item with highest estimated affinity mean.

Random: $j^* = \text{random}(1, N)$

A trivial *explore*-only strategy.

Maximum Variance (MV): $j^* = \arg \max_j \epsilon_{ij}$

A *explore*-only strategy, variance reduction strategy: Select the item with the highest noisy affinity variance.

Maximum Item Trait (MaxT): $j^* = \arg \max_j \|\mathbf{v}_j\|_2$

Select the item whose trait vector $\mathbf{v}_j$ contains the most information, namely has highest L2 norm $\|\mathbf{v}_j\|_2 = \sqrt{v_{j1}^2 + v_{j2}^2 + \dots + v_{jd}^2}$.

Minimum Item Trait (MinT): $j^* = \arg \min_j \|\mathbf{v}_j\|_2$

Select the item with trait vector with least information.

Upper Confidence (UCB): $j^* = \arg \max_j y_{ij} + \epsilon_{ij}$

Based on UCB1 [3]: Pick the item with the highest upper confidence bound, namely mean plus variance (95% CI)

Thompson Sampling (TS) [5]: $j^* = \arg \max_j \hat{y}_{ij}$

For each item, sample the noisy affinity from the posterior. Select item with the maximum sampled value.

Actual conversation...

Li et al. (2018) sought approaches more closely matching “free-form” conversation. Roughly:

- Dialogs (around 10k) are constructed by crowd workers, who assume roles of a *recommender* or *seeker*;
- Conversations between the recommender and the seeker are tagged in terms of the movies mentioned, as well as explicit feedback (has the seeker seen the movies mentioned and did they like them);
- Train a dialog generation model that can fulfil the role of the recommender;
- Preferences can then be estimated and the output controlled to reference specific movies

Table 5: Sample conversation from validation set

SEEKER:	hi !
HUMAN:	hi ! how is your morning going ?
HRED:	hi , what kind of movies do you like
OURS:	hello ! what kind of movie are you looking for ?
SEEKER:	great , thanks ! i was thinking of starting the day off with a great comedy , any suggestions ?
HUMAN:	i love the hangover it always makes me laugh . and i bet lots of people are hungover this morning ! airplane ! is another classic of mine .
HRED:	that 's a good one ! have you seen the lego movie ? it 's a great movie !
OURS:	i have n't seen that one ! what kind of movies do you like ? i can help you with that .
SEEKER:	i 've seen the hangover it 's really good ! i thought about watching knocked up have you seen it ?
HUMAN:	i enjoyed that movie . i enjoyed spy it 's a well done parody .
HRED:	it 's about a man who gets revenge on a cyborg and a hockey team . it 's really good . it 's really good . i also like the girl on the train
OURS:	i have n't seen that one . i will have to check it out . i hope i was able to help you find a good one ! have a great day !
SEEKER:	thanks ! i 'll check it out !
HUMAN:	sure ! have a great day !
HRED:	you 're welcome ! have a great day !
OURS:	have you seen revenge of the nerds ? it 's really good . it 's a classic .
SEEKER:	you too .

(from ReDial)

Actual conversation...

Li et al. (2018)'s approach has a number of virtues:

- Actually looks (more or less) like “real” conversation, especially compared to approaches that came before
- Contributes a (now widely used) benchmark dataset for training and evaluation
- Elegant / principled in terms of how the model is trained and the objective it's trained for (i.e., reach a goal movie in the fewest possible number of steps)

Actual conversation...

Though it also has some **limitations**

- Conversations aren't particularly "real": the users aren't actually seeking some item, but play a synthetic game in which they are told which item to seek
- It's unclear to what extent the data collection effort could be applied in other settings, in particular ones not based on "general knowledge" (i.e., for which crowd workers would struggle to engage in synthetic conversations)
- Even within movies, it's hard to tell how closely conversations in ReDial (or similar efforts) represent "organic" conversations

Actual conversation...

Li et al. (2018) sought approaches more closely matching “free-form” conversation. Roughly:

- Dialogs (around 10k) are constructed by crowd workers, who assume roles of a *recommender* or *seeker*;
- Conversations between the recommender and the seeker are tagged in terms of the movies mentioned, as well as explicit feedback (has the seeker seen the movies mentioned and did they like them);
- Train a dialog generation model that can fulfil the role of the recommender;
- Preferences can then be estimated and the output controlled to reference specific movies

Table 5: Sample conversation from validation set

SEEKER:	hi !
HUMAN:	hi ! how is your morning going ?
HRED:	hi , what kind of movies do you like
OURS:	hello ! what kind of movie are you looking for ?
SEEKER:	great , thanks ! i was thinking of starting the day off with a great comedy , any suggestions ?
HUMAN:	i love the hangover it always makes me laugh . and i bet lots of people are hungover this morning ! airplane ! is another classic of mine .
HRED:	that 's a good one ! have you seen the lego movie ? it 's a great movie !
OURS:	i have n't seen that one ! what kind of movies do you like ? i can help you with that .
SEEKER:	i 've seen the hangover it 's really good ! i thought about watching knocked up have you seen it ?
HUMAN:	i enjoyed that movie . i enjoyed spy it 's a well done parody .
HRED:	it 's about a man who gets revenge on a cyborg and a hockey team . it 's really good . it 's really good . i also like the girl on the train
OURS:	i have n't seen that one . i will have to check it out . i hope i was able to help you find a good one ! have a great day !
SEEKER:	thanks ! i 'll check it out !
HUMAN:	sure ! have a great day !
HRED:	you 're welcome ! have a great day !
OURS:	have you seen revenge of the nerds ? it 's really good . it 's a classic .
SEEKER:	you too .

(from ReDial)

Actual conversation...

Li et al. (2018)'s approach has a number of virtues:

- Actually looks (more or less) like “real” conversation, especially compared to approaches that came before
- Contributes a (now widely used) benchmark dataset for training and evaluation
- Elegant / principled in terms of how the model is trained and the objective it's trained for (i.e., reach a goal movie in the fewest possible number of steps)


Actual conversation...

Though it also has some **limitations**

- Conversations aren't particularly "real": the users aren't actually seeking some item, but play a synthetic game in which they are told which item to seek
- It's unclear to what extent the data collection effort could be applied in other settings, in particular ones not based on "general knowledge" (i.e., for which crowd workers would struggle to engage in synthetic conversations)
- Even within movies, it's hard to tell how closely conversations in ReDial (or similar efforts) represent "organic" conversations

“LM+RecSys” approaches (UniCRS; Wang et al., 2022)

(Fairly) recent attempts incorporate knowledge grounding, and arguably (among a few others) represented the pre-LLM state-of-the-art




End-to-End Architecture

E.g., Sequence-to-Sequence models, Generative Language Models (GPT)



Modularized Architecture

e.g., Conversational Agent as Linked Functional Modules

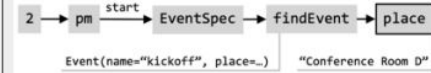


Data-Flow Architecture

E.g., Dialogue State as Dataflow Graphs (DataFlow)

User: *Where is my meeting at 2 this afternoon?*


`place(findEvent(EventSpec(start=pm(2))))`



Agent: *It's in Conference Room D.*

Four Major Modules

- Natural Language Understanding/Generation
- Dialogue State Management
- Recommendation
- Explanation



System Types

Unified GenCRS

- A system that **jointly models multiple CRS subtasks**:
 - Intent detection
 - Preference modeling
 - Item ranking
 - Response generation
- Uses a **single generative architecture**
- Internally, subtasks may be **decomposed into multiple steps**
- The goal is to **unify processes** through:
 - Multi-task training
 - Structured prompting

System Types

Modular GenCRS

- A CRS composed of **two or more specialized modules**.
- (L)LMs are mainly used for **conversation-related tasks**:
 - Managing dialogue/state
 - Generating responses
- **Traditional recommendation or retrieval** mechanisms provide **item rankings**

System Types

Agentic GenCRS

- **Goal-driven** agent framework
- **Central LLM** plans and coordinates sub-agents/tools
 - Retrieval/recommendation engines
 - Search APIs, etc.
- Performs **step-by-step reasoning** to choose next actions
- **Executes tool calls** and integrates results
- Reflects on **feedback and memory** to update plans
- **Proactively guides** the conversation toward the user's long-term objectives

Agentic Recommender Systems (Agentic-RecSys)

User intent (realistic):

*“Plan a Mickey-themed birthday party within \$300;
needs gluten-free cake; decorate in red/yellow palette.”*



Agentic Recommender Systems (Agentic-RecSys)

User intent (realistic):

*“Plan a **Mickey-themed birthday** party **within \$300**;
needs **gluten-free cake**; decorate in **red/yellow palette**.”*



Agentic Recommender Systems (Agentic-RecSys)

User intent (realistic):

*“Plan a **Mickey-themed birthday** party **within \$300**;
needs **gluten-free cake**; decorate in **red/yellow palette**.”*

FOUR
Constraints/Conditions



Agentic Recommender Systems (Agentic-RecSys)

Manual user effort today:

- Dozens of searches, tabs, Comparisons
- Cross-check theme consistency and constraints (budget/diet), availability
- Assemble a coherent bundle

One-shot recommenders: rank items
⇒ not a full plan



Many tabs, filters, copy/paste, checklist, budget math, ...

Agentic Recommender Systems (Agentic-RecSys)


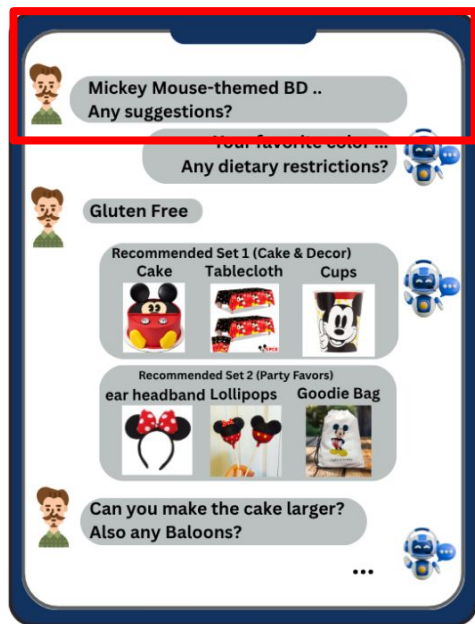
- User → juggles multiple tabs, filters, and Categories.
- The system supports micro-decisions; the user does the orchestration.

Outcome: friction, missed synergies, suboptimal bundles.




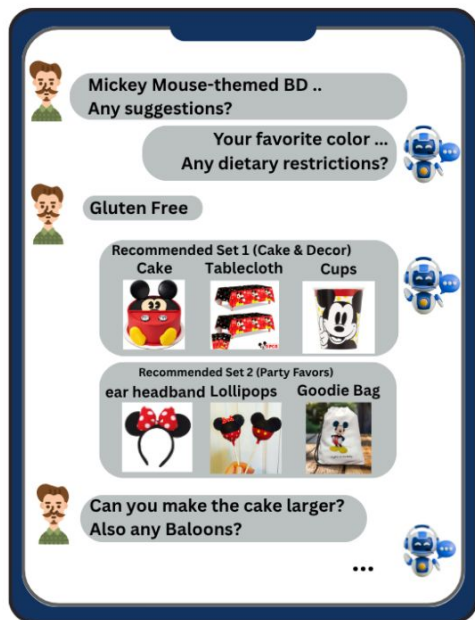
Agentic Recommender Systems (Agentic-RecSys)

Open-ended goal



Agentic Recommender Systems (Agentic-RecSys)

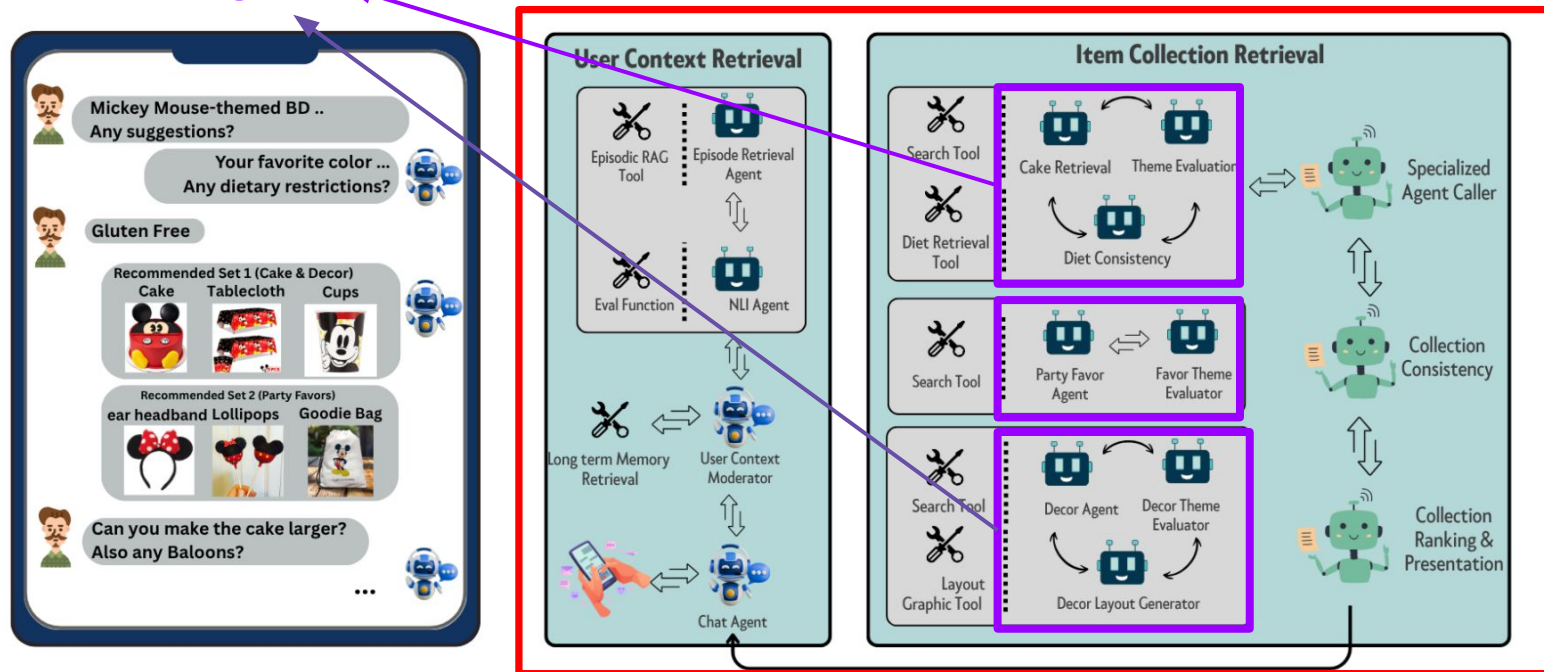
multi-agent pipeline = agents + tools



Agentic Recommender Systems (Agentic-RecSys)

User & item-side agents

multi-agent pipeline = agents + tools





Ahmadou Wagne


Core Systems & Components

Agenda

- Introduction
- **Core Systems & Components**
 - **System Architecture**
 - Dialogue Initiative
 - Recommendation Generation
 - Response Generation
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- Simulation
- Evaluation
- Open Challenges & Future Directions

System Architecture

General CRS Architecture





Key Questions

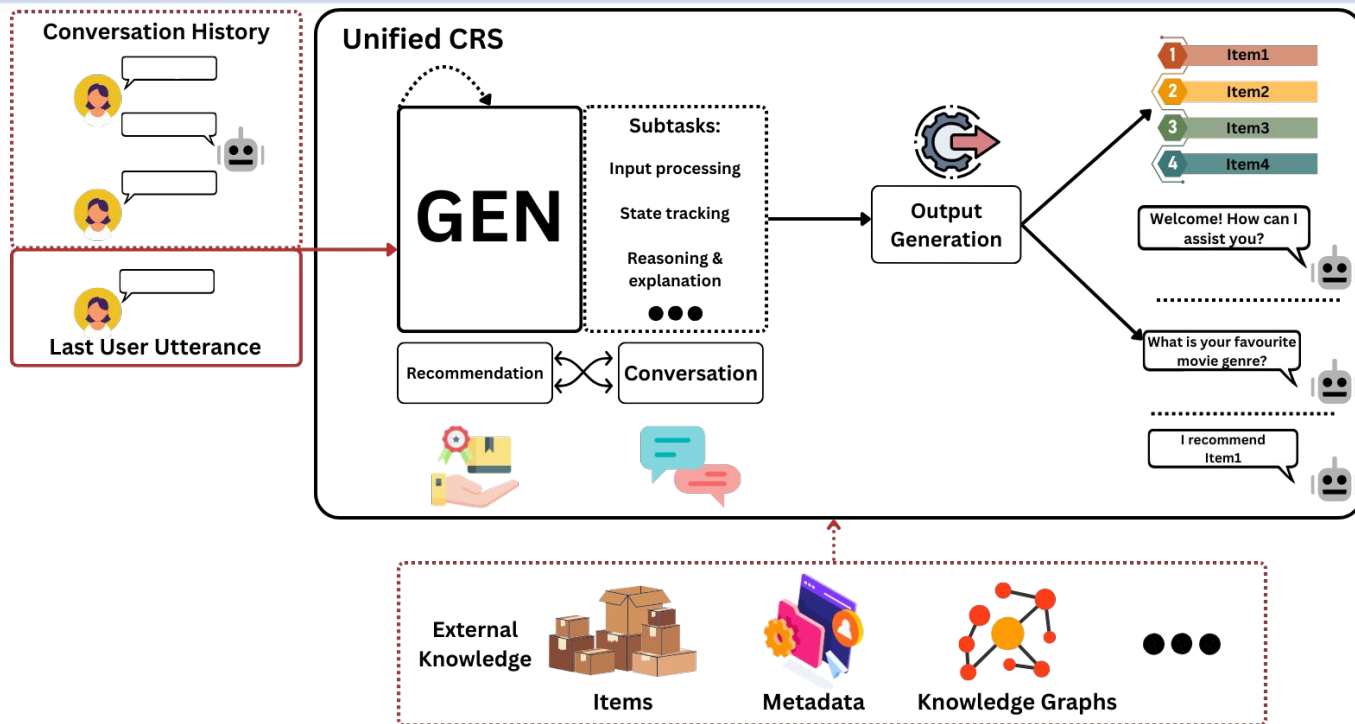
- What are different architectural paradigms in GenCRS?
- Which role do generative models play?
- How are multi-turn dialogues managed?
- What methods are used to generate recommendations and responses?

System Types

- GenCRS as **emerging** topic over the past years
- **Modular** systems **dominate** the current research about GenCRS
- Agentic are gaining traction, while unified systems mostly feature early investigations of LLMs as CRS
- Most GenCRS are **standalone** applications, while agentic systems are often integrated in existing platforms

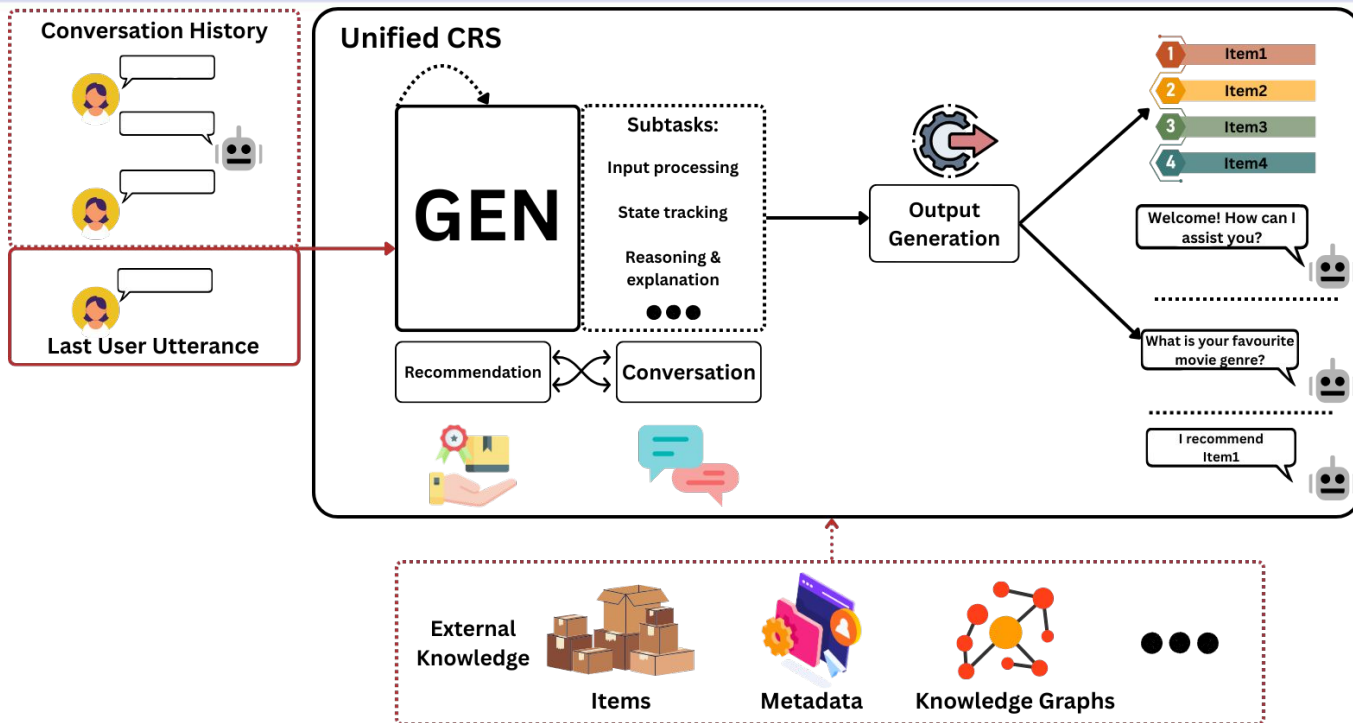


System Types - Unified



System Types - Unified

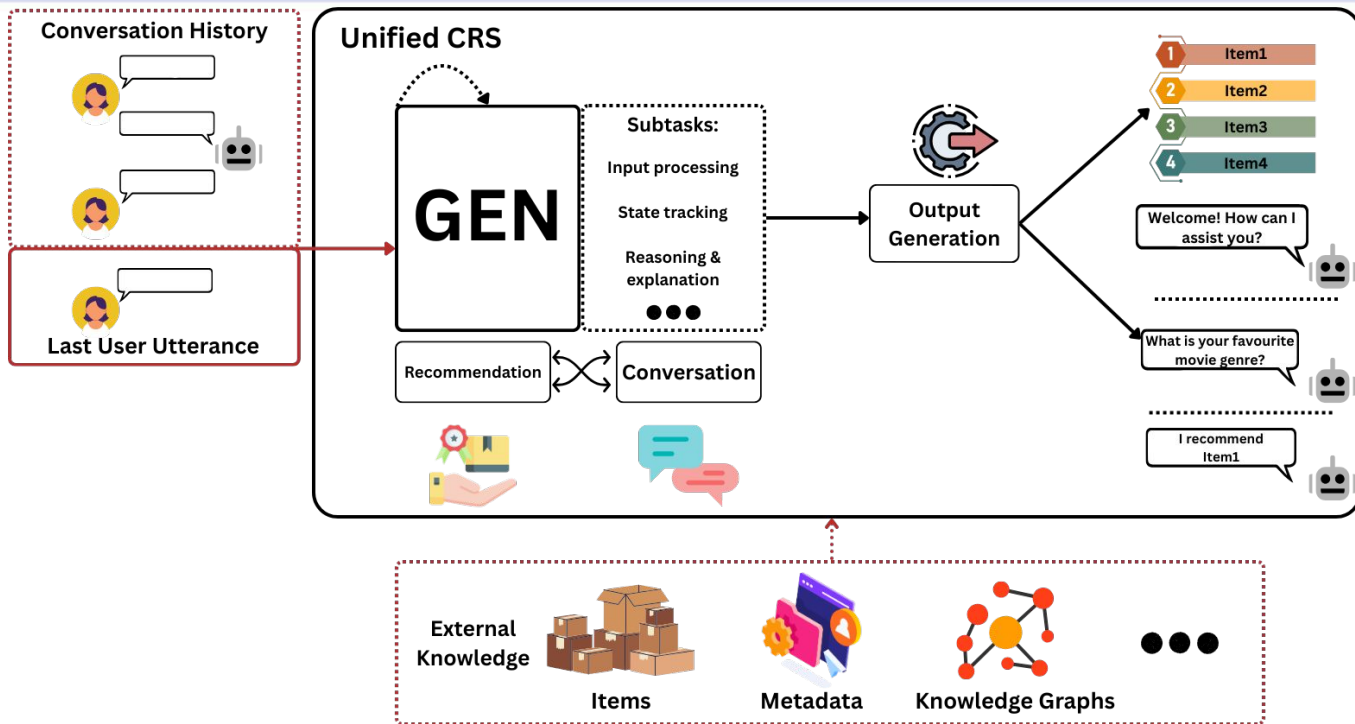
Unification of
all CRS tasks



System Types - Unified

Unification of
all CRS tasks

Out-of-the-box
models or
specific tuning

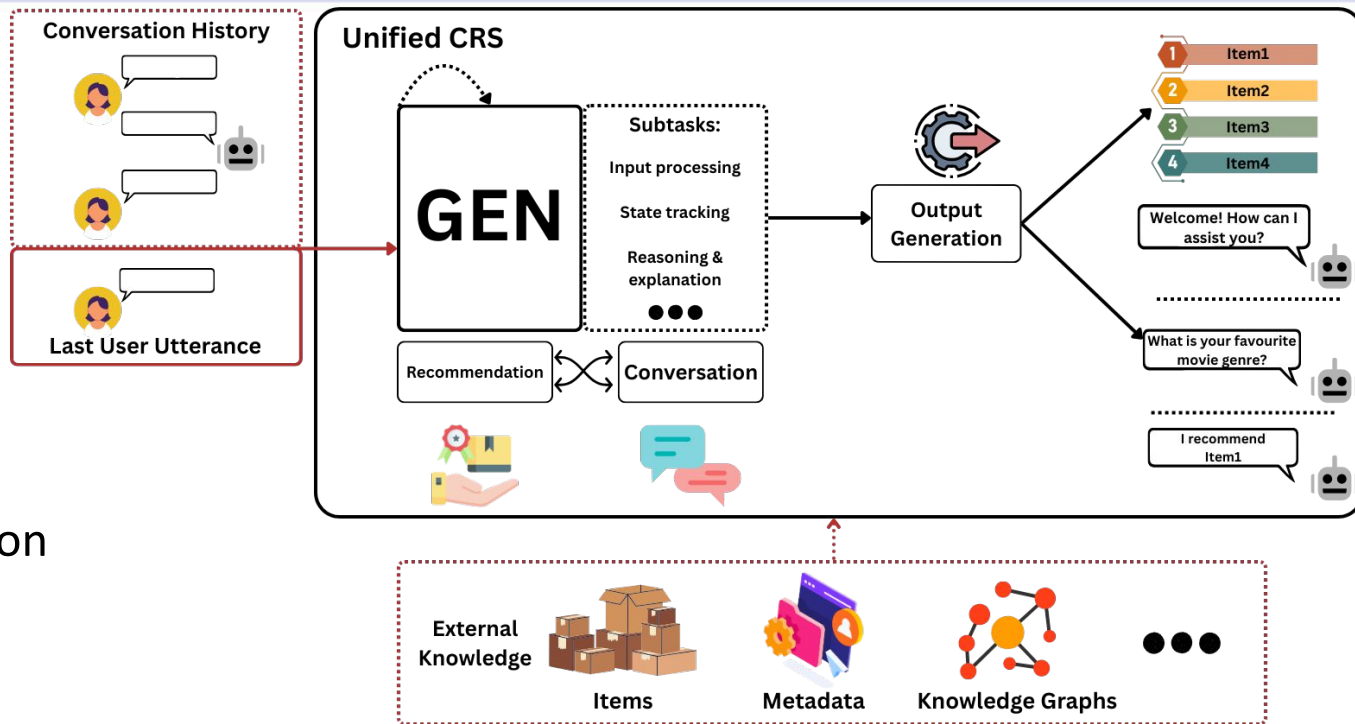


System Types - Unified

Unification of
all CRS tasks

Out-of-the-box
models or
specific tuning

Output either
recommendation
list or system
response



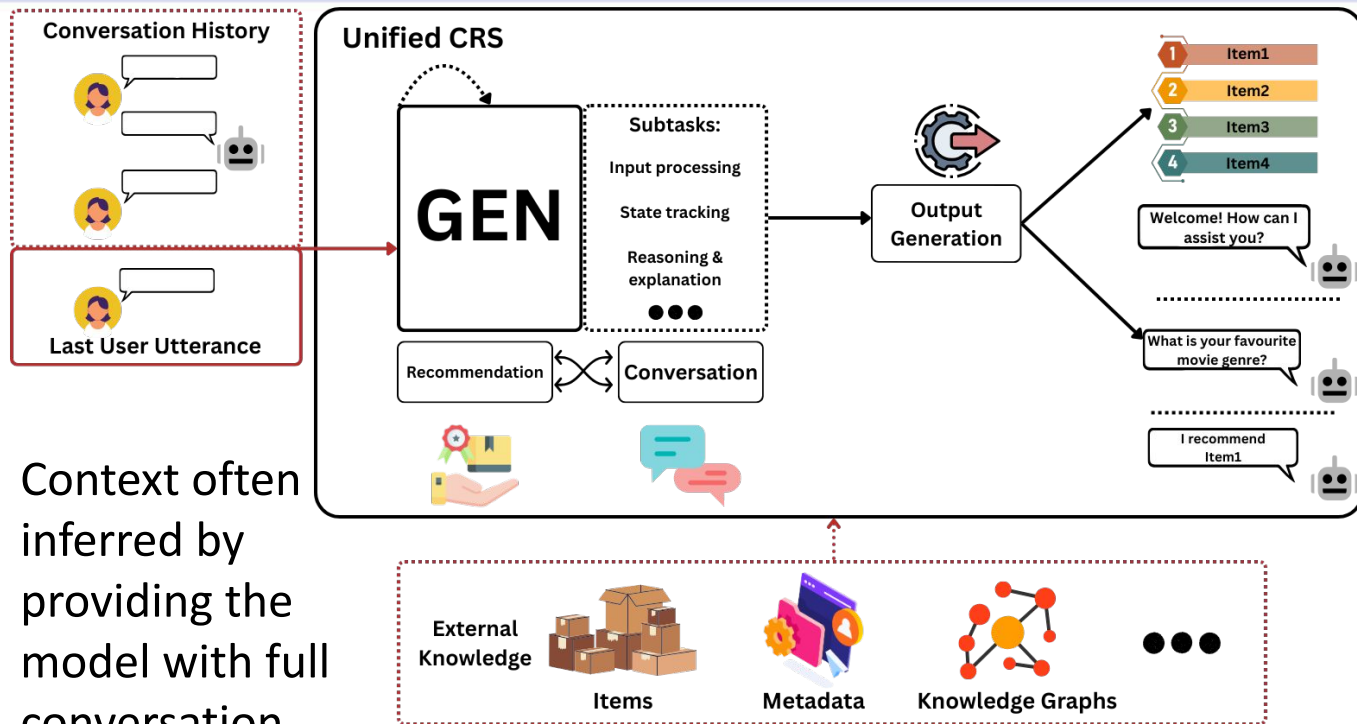
System Types - Unified

Unification of
all CRS tasks

Out-of-the-box
models or
specific tuning

Output either
recommendat
ion list or
system
response

Context often
inferred by
providing the
model with full
conversation

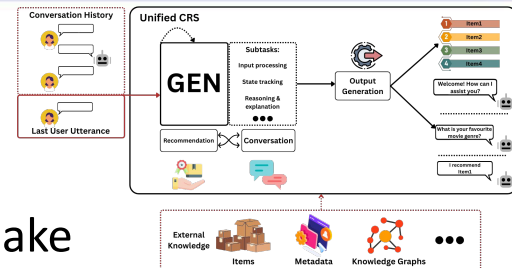


Key Questions

- Does a unified model's single LLM handle all sub-tasks end-to-end (intent detection, retrieval, ranking, NLG)?
- How well do unified systems leverage pre-trained content knowledge compared with collaborative signals?
- Do integrated architectures naturally produce richer justifications for their recommendations?


Examples - LLM as CRS

- LLM base model as zero-shot recommender
- Input: **conversation (S)**, **task (T)**, **format (F)**
- Main findings:
 - LLMs mainly rely on **content/context** knowledge to make recommendations
 - LLMs may generate **out-of-dataset** item titles, but **few hallucinated** recommendations
 - GPT-based LLMs possess **better content/context** knowledge than existing CRS
 - LLMs generally possess **weaker collaborative** knowledge than existing CRS
 - LLM recommendations suffer from **popularity bias** in CRS



Examples - LLM as CRS

- LLM base model as zero shot recommender
- Input: **conversation (S)**, **task (T)**, **format (F)**
- Main findings:



1. [User]: I love **Back to the Future**, any recommendations?
2. You would love **Terminator**! :[System]
3. [User]: Who is in it?
4. Arnold Schwarzenegger! :[System]
5. [User]: Did they make a new **Terminator**?
6. Yes, there is a new **Terminator** movie. :[System]
Have you seen the trailer for it?
7. [User]: I also need a sci-fi movie with my family,
it should be lighthearted and enjoyable.
8. [BLANK] :[System]

Pretend you are a movie recommender system. I will give you a conversation between a user and you (a recommender system).

Based on the conversation, you reply me with 20 recommendations without extra sentences.

Here is the conversation:
{ }

T
F
S



1. Guardians of the Galaxy
2. The Lego Movie
3. Men in Black
4. WALL-E
5. The Fifth Element ...

Φ

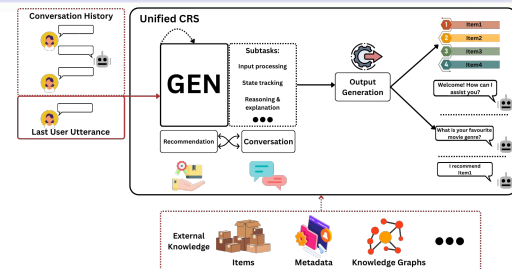
1. MOVIE_320442
2. MOVIE_352933
3. MOVIE_435849
4. MOVIE_235802
5. MOVIE_239823 ...

Prompting

Processing



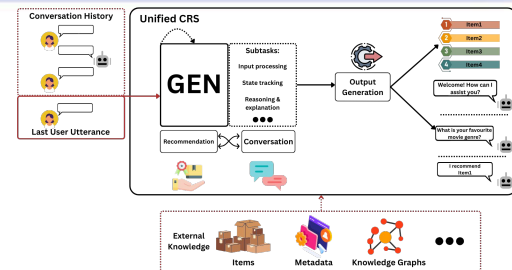
Examples - Distribution Misalignment

- **Autoregressive nature** of LLMs hinders ability to control recommendations across **entire item set**
- Condense items into **single tokens (reindex)** and distill LLM-generated recommendations as **ranked list (adapt)**
- **Abilities:** LLMs already indexed a large number of popular movie items, enabling the **understanding of complex conversations about items**
- **Limitations: Misalignment with (dynamic) data distributions**, resulting in **insufficient capturing of collaborative information**



Examples - Distribution Misalignment

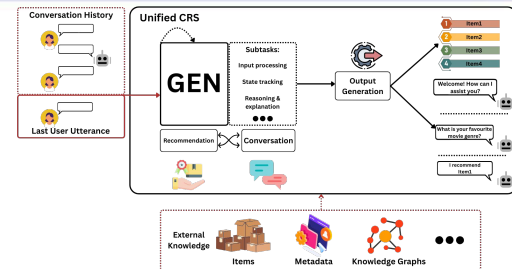
- **Autoregressive nature** of LLMs hinders ability to control recommendations across **entire item set**
- Condense items into **single tokens (reindex)** and distill LLM-generated recommendations as **ranked list (adapt)**
- **Abilities:** LLMs already indexed a large number of



Examples - Distribution Misalignment

Findings

- LLMs show **sufficient content knowledge**
- LLMs show **severe distribution misalignment**
- LLMs **struggle** to use **collaborative information**
- Modular architecture that introduces **RecSys gating** traditional methods **improves CRS performance**

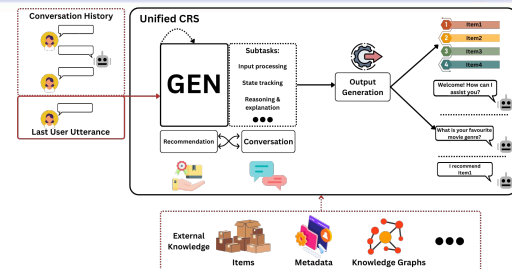


	Agg. Frozen?	INSPIRED	ReDIAL	RedditV1.5
Llama2-R	–	.131 .023	.120 .005	.056 .001
Bias Term Adjustment (w/ Bias)				
w/ gW	×	.150 .025	.165 .006	.082 .002
	✓	.146 .024	.141 .006	.058 .001
w/ b	×	.174 .026	.165 .006	.087 .002
	✓	.136 .235	.117 .005	.058 .001
w/ gW+b	×	.155 .025	.167 .006	.088 .002
	✓	.160 .025	.144 .005	.058 .001
RecSys Model Gating (w/ RecSys)				
w/ FISM	×	.197 .027	.146 .005	.093 .002
	✓	.207 .028	.165 .006	.074 .002
w/ SASRec	×	.178 .026	.148 .005	.093 .002
	✓	.188 .027	.157 .006	.075 .002

Examples - Distribution Misalignment

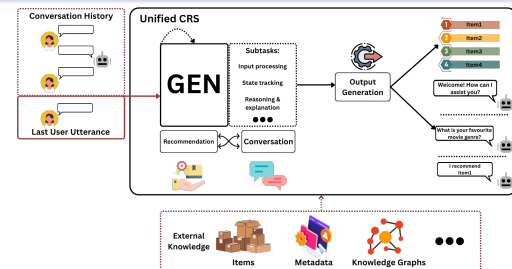
- LLMs show **su**
- LLMs show **se**
- LLMs **struggle**
- Modular archi
methods **impr**

	Agg. Frozen?	INSPIRED	ReDIAL	RedditV1.5
Llama2-R	–	.131 .023	.120 .005	.056 .001
<i>Bias Term Adjustment (w/ Bias)</i>				
w/ gW	×	.150 .025	.165 .006	.082 .002
	✓	.146 .024	.141 .006	.058 .001
w/ b	×	.174 .026	.165 .006	.087 .002
	✓	.136 .235	.117 .005	.058 .001
w/ gW+b	×	.155 .025	.167 .006	.088 .002
	✓	.160 .025	.144 .005	.058 .001
<i>RecSys Model Gating (w/ RecSys)</i>				
w/ FISM	×	.197 .027	.146 .005	.093 .002
	✓	.207 .028	.165 .006	.074 .002
w/ SASRec	×	.178 .026	.148 .005	.093 .002
	✓	.188 .027	.157 .006	.075 .002


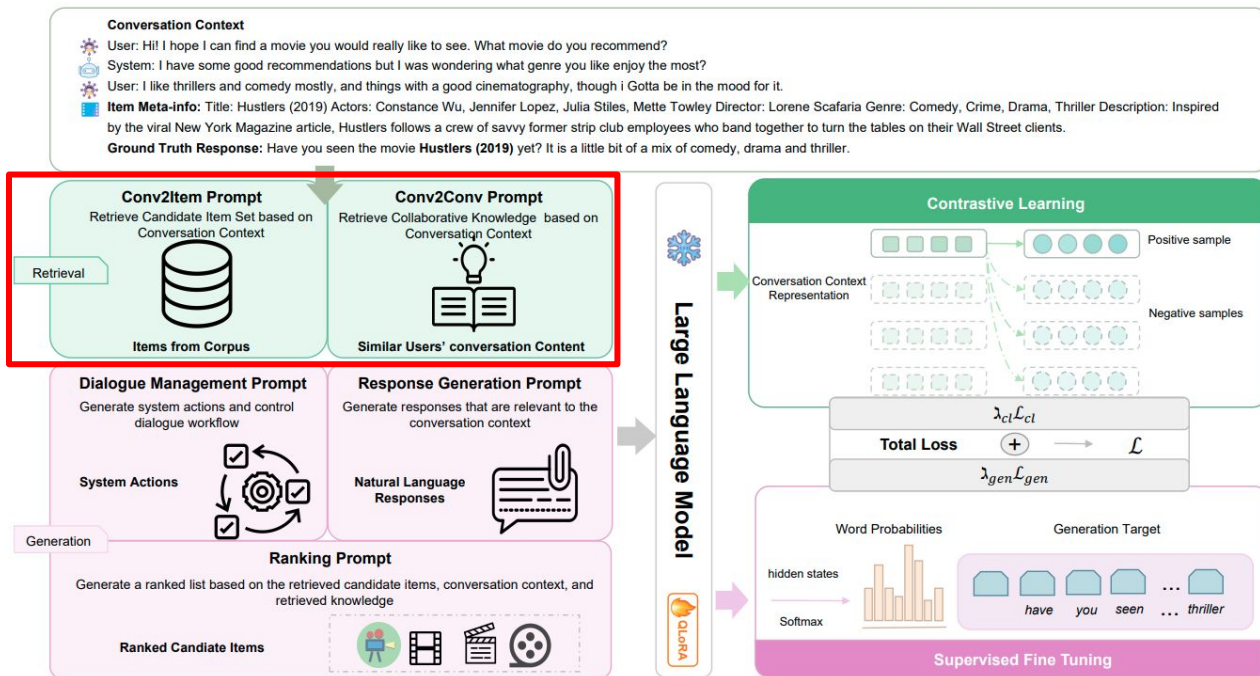


Examples - LLM-Based Retrieval

- LLM **unifies tasks** of intent detection, response generation, retrieval and recommendation
- **Instruction tuning** of the LLM to align retrieval and generation task
- **Joint optimization** of **retrieval** and **generation** tasks (contrastive learning, generation loss)
- **Two-stage retrieval**: retrieve **items** based on current conversation, retrieve **similar historical conversations** to incorporate collaborative knowledge



Examples - LLM-Based Retrieval



System Types - Summary: Unified

Role of the LLM

Central: intent detection, retrieval, ranking, NLG

Strong content/context knowledge

→ rich semantics & explanations

Weak collaborative knowledge → poor alignment with interactions

Architectural Implications

- ✓ Simplify pipeline: single model for all sub-tasks
- ✓ Natural dialogue flow and richer justifications
- ✗ Distribution misalignment with CRS data
- ✗ Autoregressive nature limits item set control
- ✗ Popularity bias, low diversity

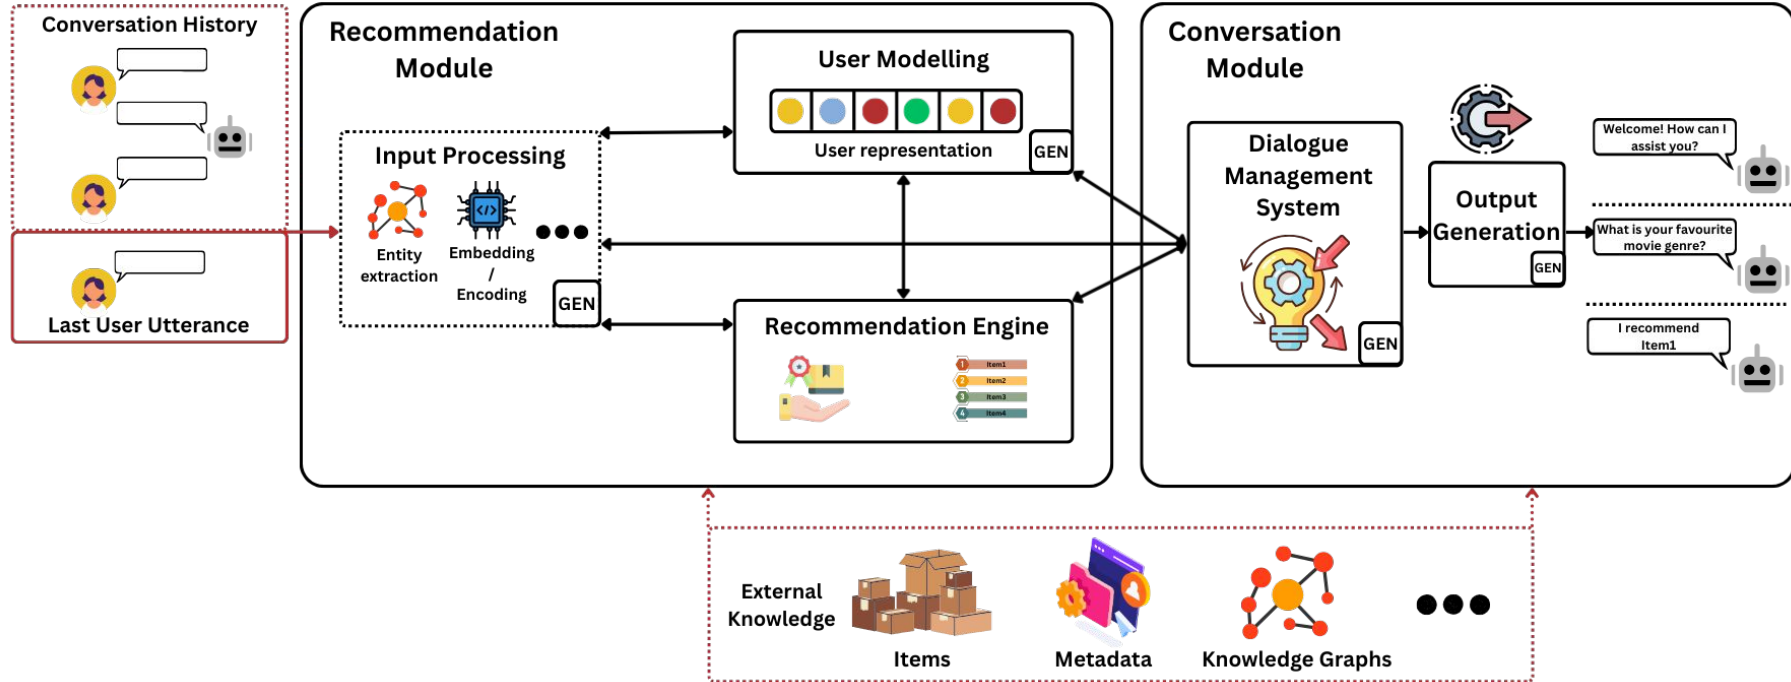
System Types - Summary: Unified

Pros & Cons

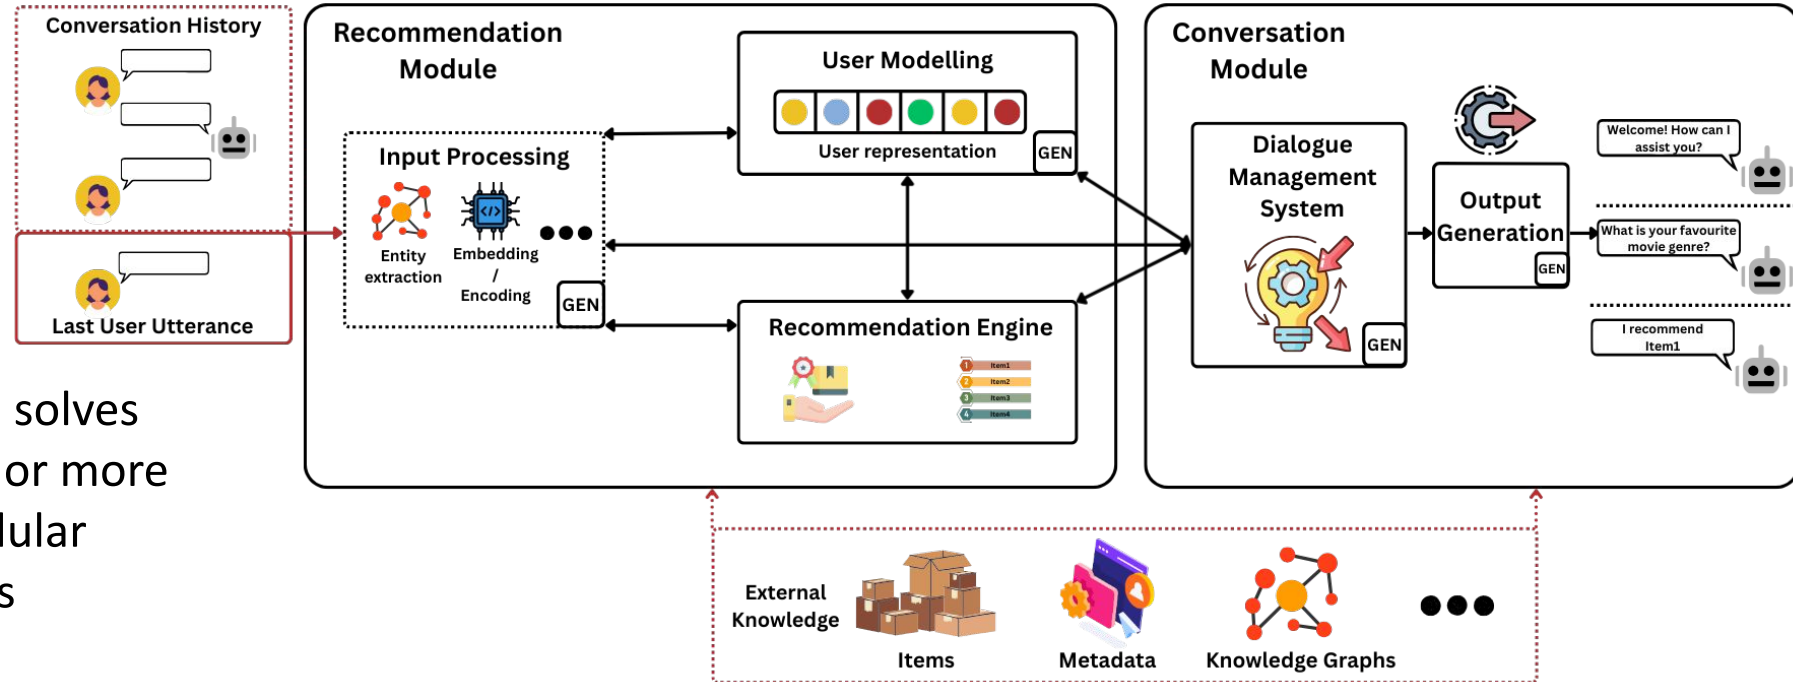
Pros: contextual explanations, flexible dialogue, good baseline without adaptation

Cons: weak adaptation, poor collaborative use, scalability issues, no isolation of error sources

System Types - Modular

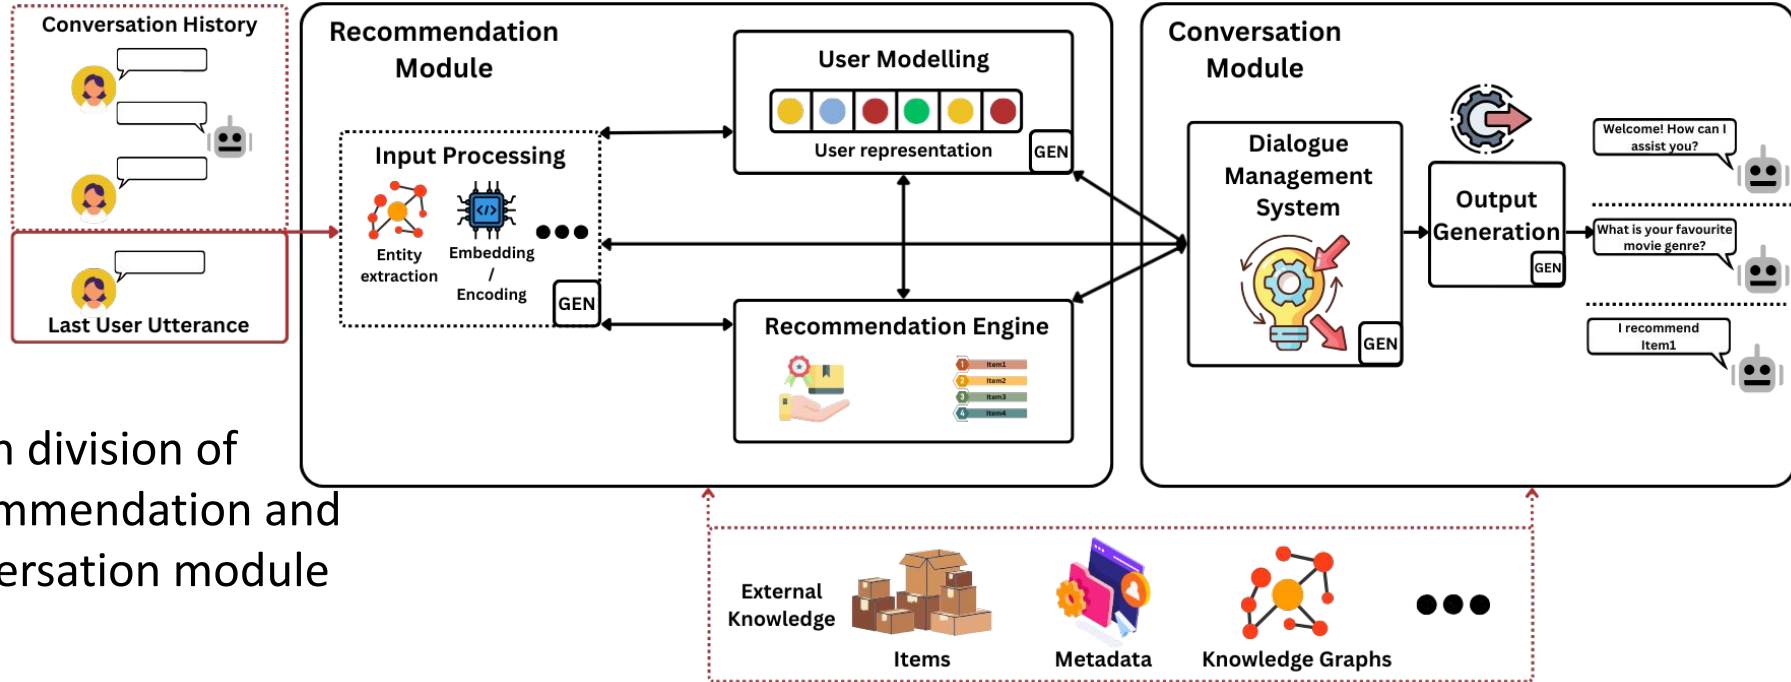


System Types - Modular



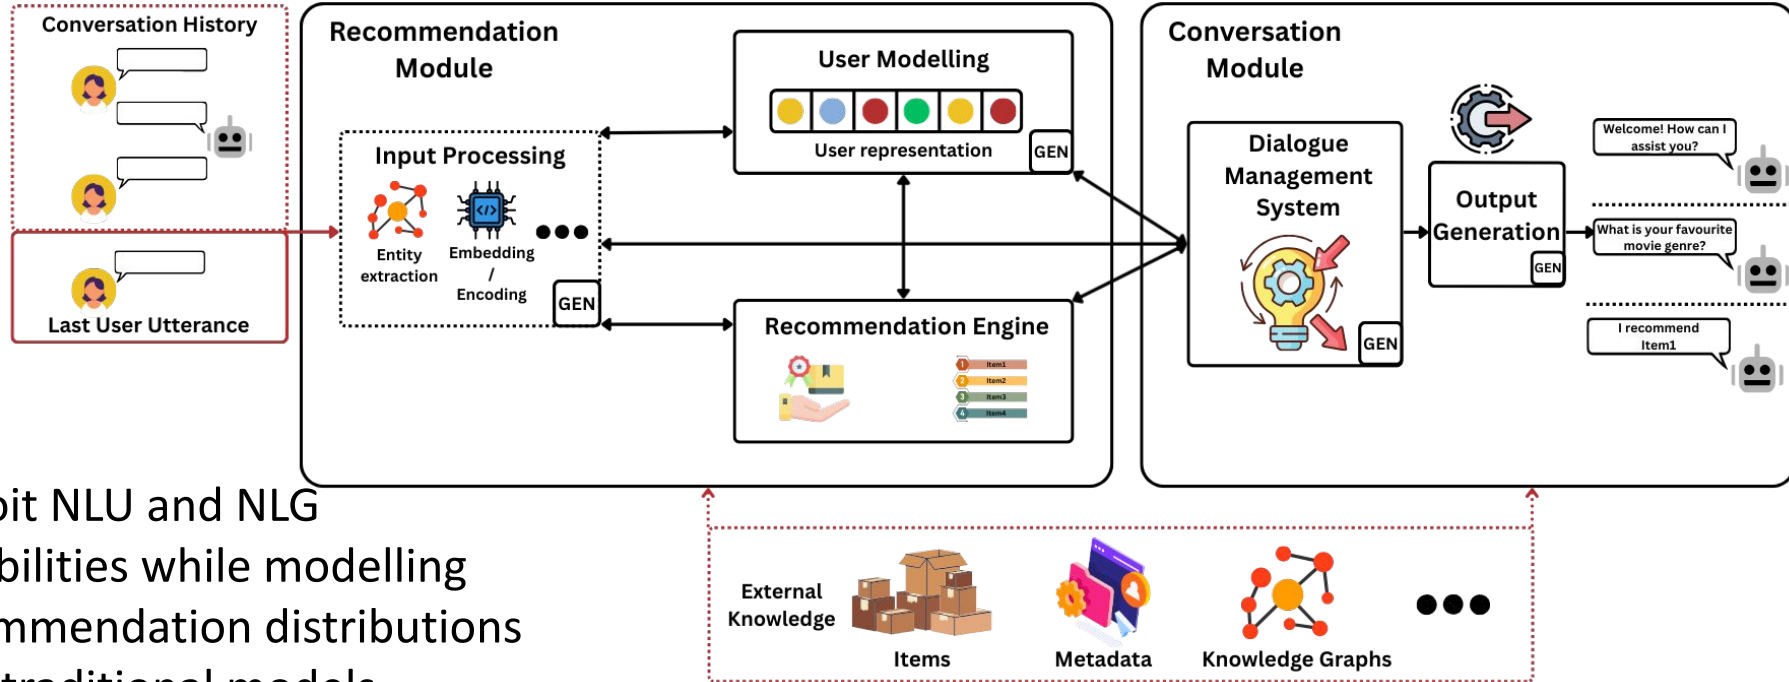
GEN solves
two or more
modular
tasks

System Types - Modular



Often division of
recommendation and
conversation module

System Types - Modular



Exploit NLU and NLG capabilities while modelling recommendation distributions with traditional models


Key Questions

- How strictly are the modules separated?
- Which tasks (e.g., collaborative filtering, explanation) play to modular systems' strengths?
- How does modularisation affect the exploitation of collaborative versus content/contextual knowledge?


Examples - Knowledge Graphs

- **Textual context representations** of previous conversations and **multi-modal knowledge graph** embeddings for **user modelling** and **recommendation generation**
- Generative model is utilized to generate a **contextual response template** that is combined with the recommended items

-> **Only possible action: recommend (based on existing dialogue)**





Examples - Knowledge Graphs



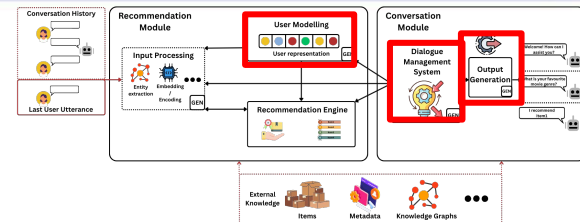
Examples - NL-Based State Tracking

- LLM performs **intent detection**, **user modelling** (extraction of user preferences), **action selection** and **response generation**
- State-tracking via **semi-structured** dictionary



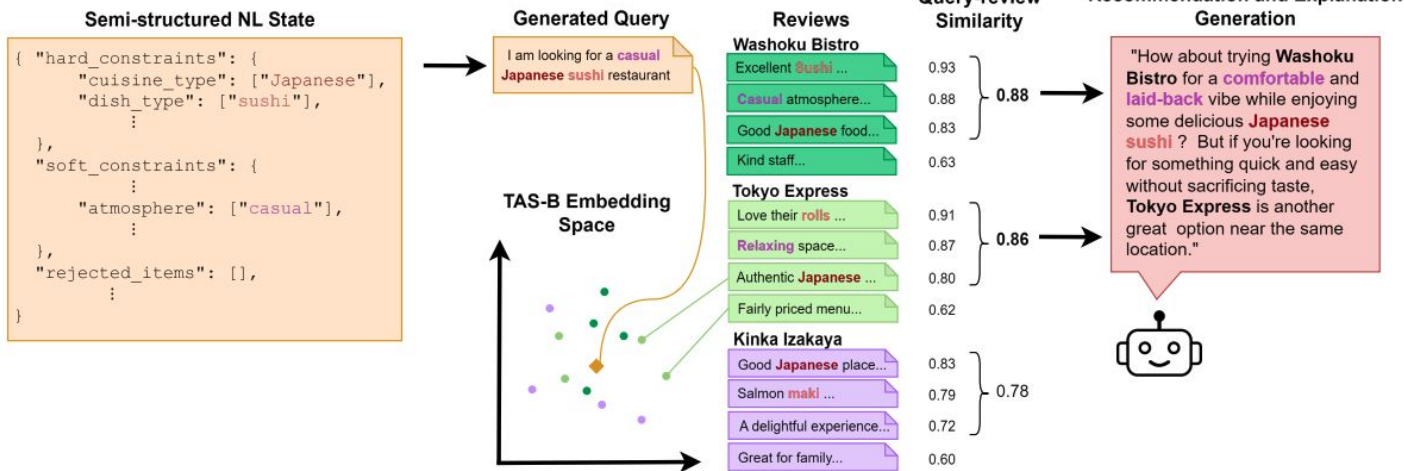
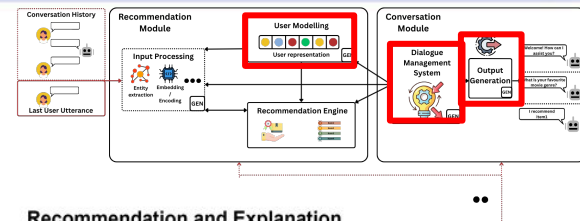
Examples - RAG for Recommendation & Explanation

- Recommended items are **retrieved** by measuring similarities between LLM-generated (based on state dictionary) query and existing user reviews
- Ensure grounded recommendations with **RAG**
- **Contextualized** response



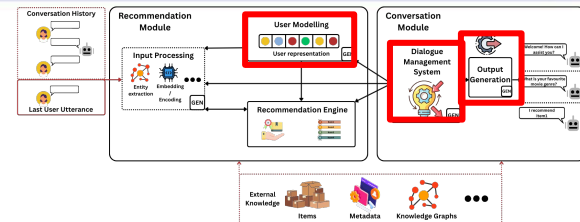
Examples - RAG for Recommendation & Explanation

- Recommended items are **retrieved** by measuring similarities between LLM-generated (based on state



Examples - LLM for User Modelling

- LLM assisted by **expert model** that provides a candidate list of items based on **collaborative knowledge**
- Final list of recommendation **generated** by LLM
- LLM as **user model** generator
- **Knowledge-augmented** generation



Examples - Examples - LLM for User Modelling

- LLM assisted by **expert model** that provides a

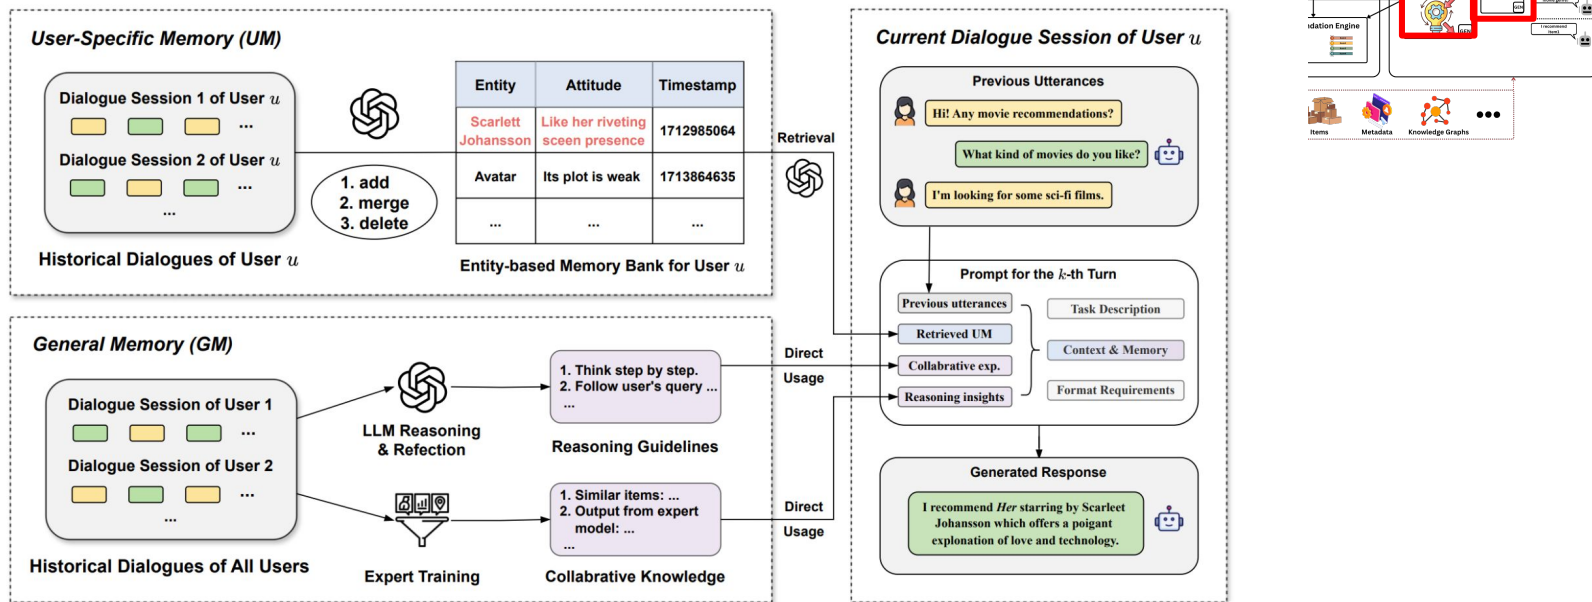
conc

know

- Final

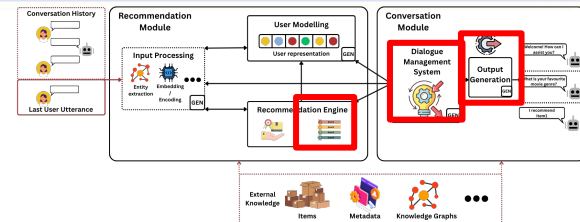
- LLM

- Know




Examples - LLM as Re-Ranker

- **Intent detection, state tracking and response generation** handled by LLM
- LLM as part of the **recommendation engine**, functioning as **ranking** submodule
- Additional role as **user simulator** for training and tuning of other modules



Examples

- Intent detection, state tracking and response



System Types - Summary: Modular

Modular Separation

- **Subtasks decoupled**: user modeling, recommendation, dialogue management, response generation
- Generative models typically act as **dialogue manager**, **user model**, or **response generator**
- Recommendation modules (collaborative filtering, KG-based, neural methods) **mostly remain separate** → generative model does not replace them

System Types - Summary: Modular

Strengths of Modular Systems

- **Alignment with target distribution** via expert models
- Often combined with **entity extraction** & (multimodal) **knowledge graphs** → richer user/item modeling

System Types - Summary: Modular


Knowledge Utilization

- Collaborative: captured through **expert models, KG** etc.
- Content/context: **dialogue history embeddings, response generation by LLM**
- Clearer separation makes it easier to **optimize collaborative vs. content knowledge** sources independently
- Grounding via **RAG**

Pros & Cons

- **Pros:** interpretability, flexibility, targeted optimization, better alignment with CRS distributions
- **Cons:** pipeline complexity, coordination overhead, risk of error propagation, semantic gap between retrieval and generation

System Types - Agentic



Key Questions

- In what ways do agentic systems extend beyond modular/unified approaches in terms of planning and proactivity?
- Which additional user goals (e.g., multi-step task completion, continuous personalisation) do agentic systems support?
- What new challenges do agent frameworks introduce (e.g., tool orchestration, safety, latency)?

Agentic Systems

- Emerging topic in CRS research
- Modular approach that employs **specialized agents** to solve different CRS sub-tasks **orchestrated** by a central LLM agent
- Core capabilities: **planning & task decomposition**, **tool use & action execution**, **memory & state** management and **autonomy & goal-driven** behavior
- Prospect advantages:
 - **Enhanced User Experience**
 - **Adaptability & Flexibility**
 - **Contextual Precision**
 - **Explainability & Transparency**

Agentic Systems - LLM Agents Characteristics



Planning & Task Decomposition:

Break complex goals into subtasks; execute multi-step reasoning for long-horizon tasks



Memory & State Management:

Maintain context across steps; store/retrieve user preferences, domain knowledge, and feedback



Tool Use & Action Execution:

Invoke external tools/APIs; interact with real-world systems (e.g., databases, knowledge bases)




Autonomy & Goal-Driven Behavior:

Operate in closed-loop fashion; observe environment, evaluate progress, self-refine until goal completion

Agentic Systems

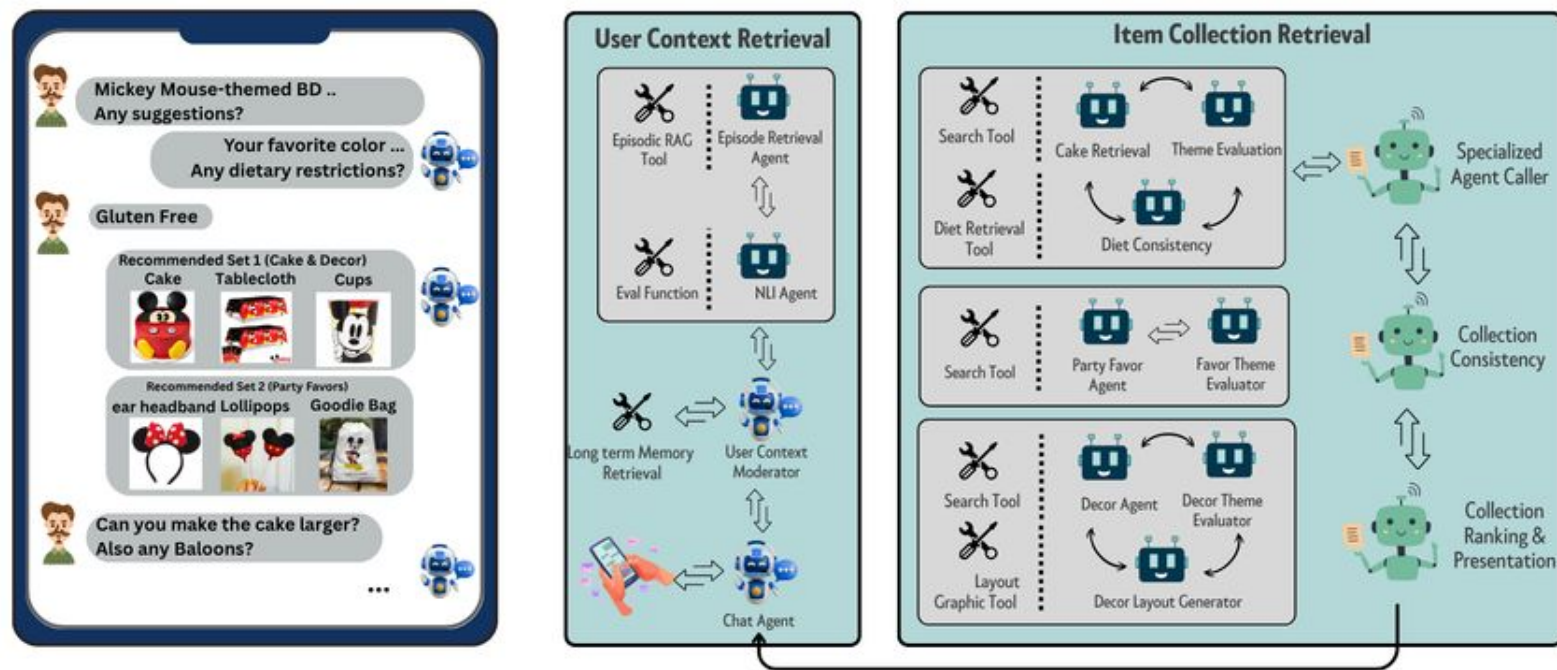
Recap

Definition (Conversational Recommender System—CRS): A CRS is a software system that supports its users in achieving **recommendation-related goals** through a **multi-turn** dialogue




goal-oriented

Agentic Systems



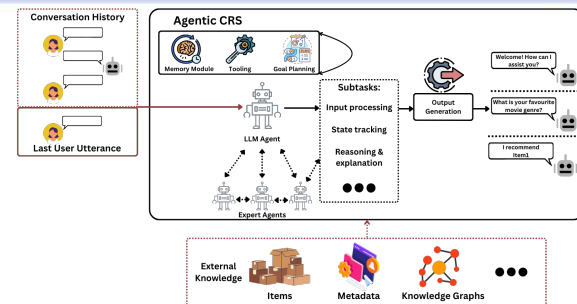
Agentic Systems

Key Characteristics of LLM Agents



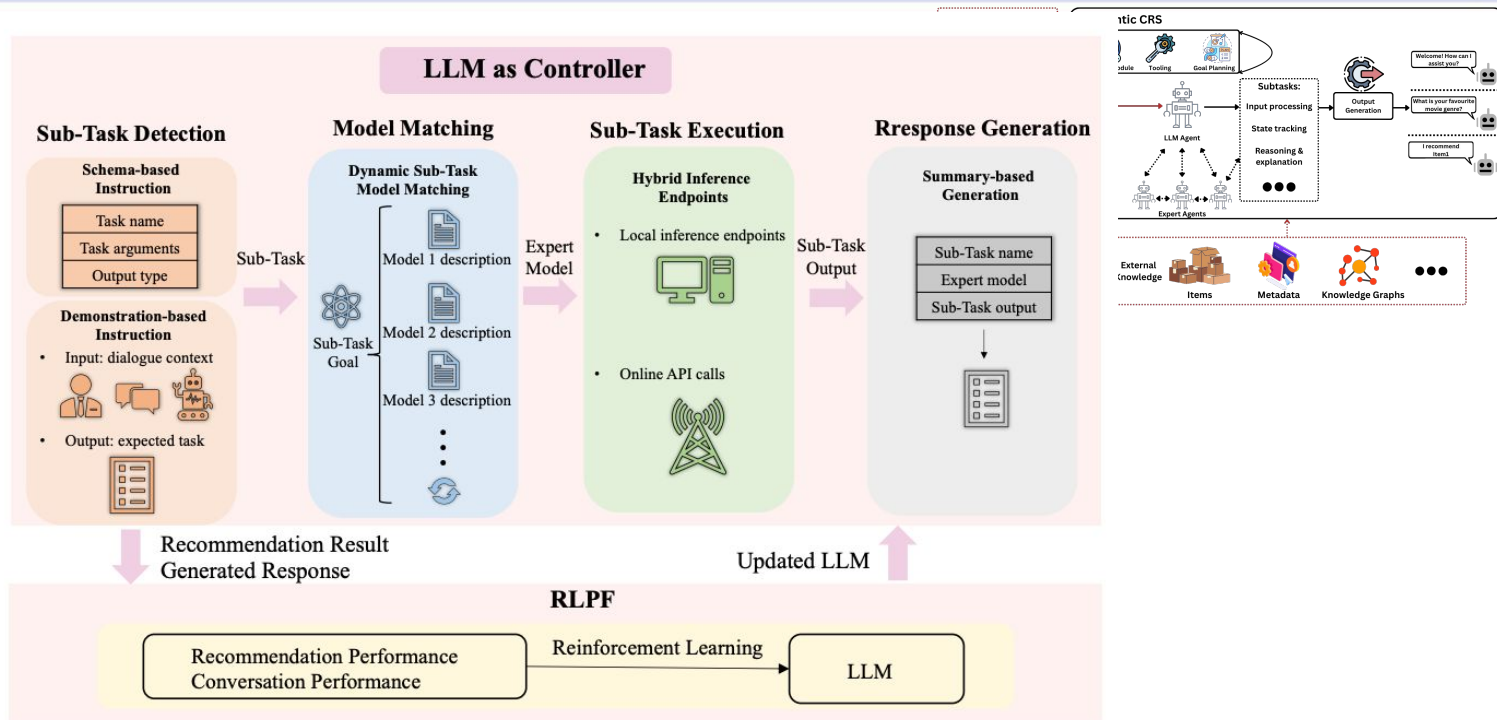
Examples - Task Delegation

- Early **agent-like** work
- **Reflection** of current subtask
- Matching with **suitable expert** model
- Task can be executed with various **tools**
- **Structured** response generation



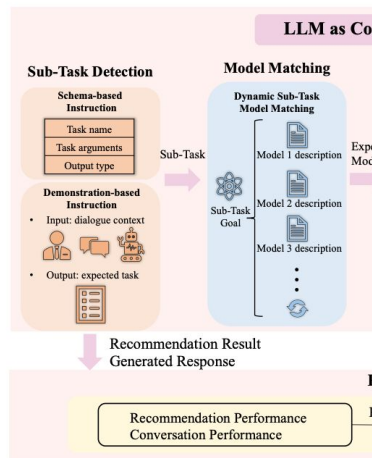
Examples - Task Delegation

- Early age
- Reflection
- Matching
- Task can l
- Structure



Examples - Task Delegation

- Early agent
- Reflection
- Matching v
- Task can be



I want to find a legal drama.

Any favorite actors or directors?

I want the female actor as the lead role.

How about "Suits"? It is a legal drama about a brilliant college dropout.

Why do you recommend it to me?



It is a legal movie whose lead role is a female. And the story is fantasy and interesting.

OK, what timeframe is it?

From 2011-2012.

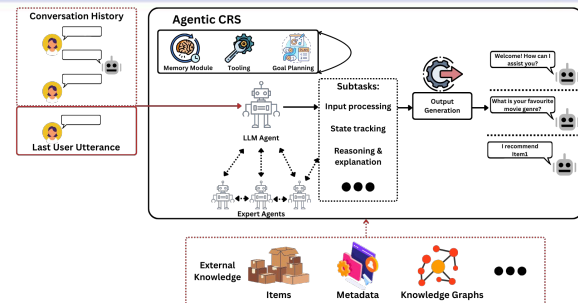
Thanks! I like this movie.

Sounds good.




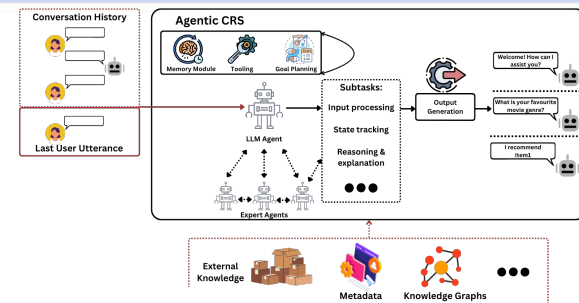
Examples - Goal Planning

- Task decomposition: **goal planner + knowledge retriever + LLM responder**
- **Tool use**: integrates external knowledge for grounded recommendations
- **Proactivity**: predicts dialogue goals, guides conversation flow
- Limitations: **short-term planning, no persistent memory, constrained autonomy**



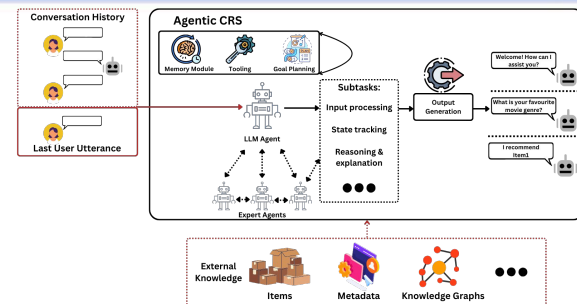
Examples

- Task decomposition: **goal planner + knowledge retriever + LLM responder**
- **Tool use:** integrates external knowledge for grounded recommendations




Examples - Tool Calling

- Pipeline structure orchestrated by a **staged workflow** (less autonomy)
- Focus on deployment in **real world system** (small to medium sized business)
- Challenges: **Latency & cost trade-offs, quality issues & prompt design, performance instability**

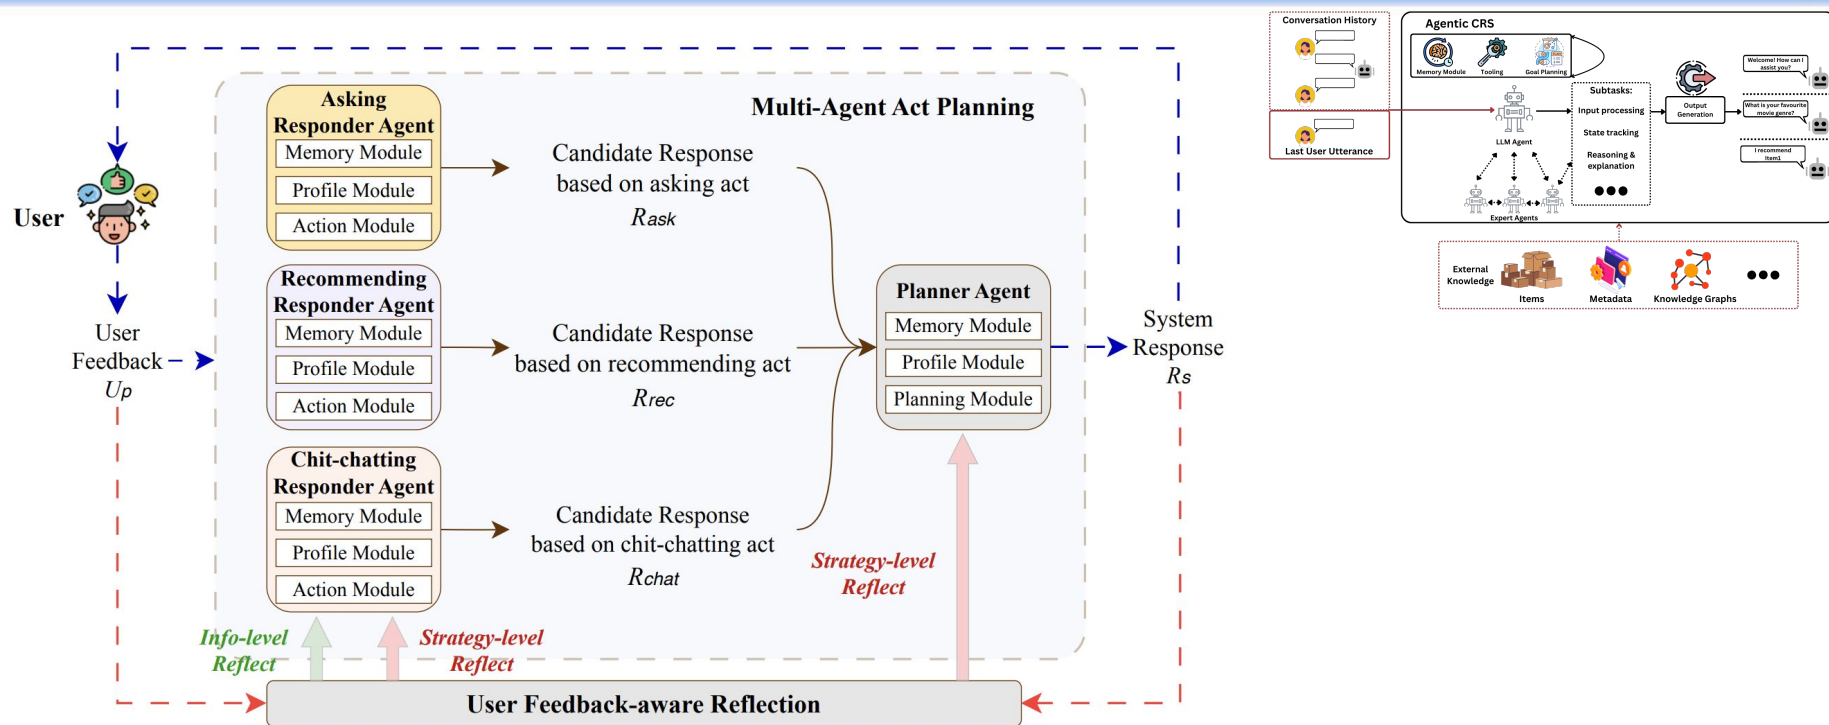


Examples - Tool Calling

- Pipeline (less a
- Focus on medium
- Challenge prompt

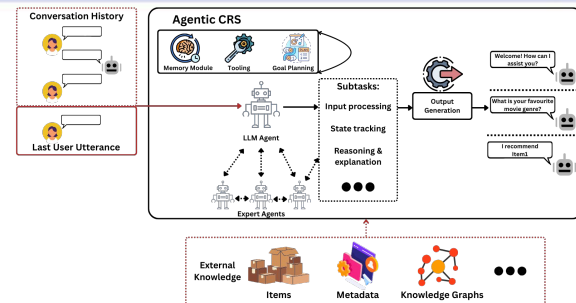


Examples - Multi-Agent Planning



Examples - Agent Recommendation

- Recommended items are **LLM-based agents** with capabilities of **interactivity**, **proactiveness**, and **knowledge**
- User, agent recommender and agent items **interact with and among each other**
- Items **evolve over time** and adapt to feedback
- **Extensible** and **adaptable** to various data sources and domains



The diagram illustrates the architecture of the Agentic CRS, showing the flow of information and interaction between a User, Agent Recommender, Agent Items, and various Agents (Fashion, Education, Travel, Photo, Music) within an Info. Layer.

Information Exchange: The User interacts with the Agent Recommender, which recommends Agent Items. The User also interacts with the Agents (Fashion, Education, Travel, Photo, Music) through the Info. Layer.

Information Carrier: The Info. Layer carries information from the User and the Agents, which is then used by the Agent Recommender to recommend Agent Items.


Agent Recommender: The Agent Recommender is responsible for recommending Agent Items based on the User's preferences and the information carried by the Info. Layer.

Agents: The Agents (Fashion, Education, Travel, Photo, Music) interact with the User and the Info. Layer, providing specific information and services.


Info. Layer: The Info. Layer is the central hub for information, receiving input from the User and the Agents, and providing output to the Agent Recommender.

Flowchart of CRS Tasks: The flowchart shows the tasks of the CRS: processing, tracking, and planning & unation. The process starts with a User input, followed by Output Generation, which leads to a series of tasks (What is your favourite music genre?, I recommend that!, I recommend that!).

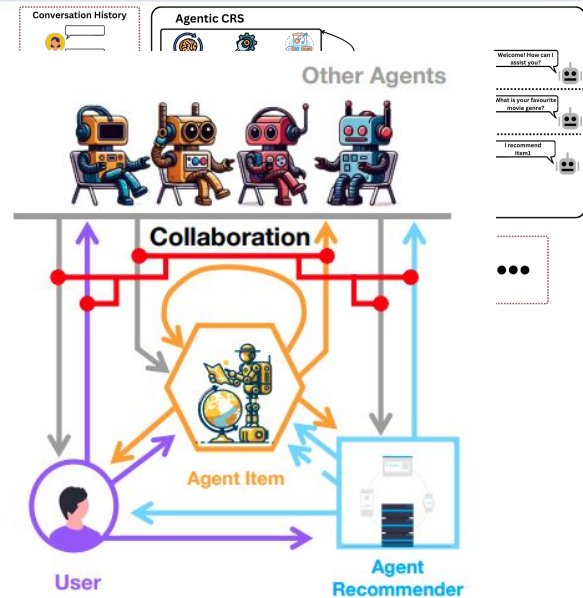
Examples - Agent Recommendation



(a) User-Agent Interaction Stage



(b) Agent-Recommender Collaboration Stage



(c) Agents Collaboration Stage

System Types - Summary: Agentic

Beyond Unified/Modular

- Extend with **planning & proactivity**: task decomposition, multi-step reasoning
- LLM as **controller** orchestrating expert models & external tools
- From **reactive** recommendations to **proactive, goal-driven** conversations

Additional User Goals

- **Multi-step** task completion (elicitation → recommendation → explanation → refinement)
- Continuous personalization via **memory** and **feedback**
- Extensible integration of **tools, APIs, and sub-agents**

System Types - Summary: Agentic

New Challenges

- Tool orchestration and coordination overhead
- Safety & controllability risks from autonomous behaviors
- Latency & cost from staged workflows and multi-step planning
- Full autonomy, long-term memory, and sustained planning remain challenging

Pros & Cons

- **Pros:** flexibility, proactivity, extensibility, continuous personalization
- **Cons:** complexity, safety risks, high latency/cost, limited true autonomy


Agenda

- Introduction
- **Core Systems & Components**
 - System Architecture
 - **Dialogue Initiative**
 - Recommendation Generation
 - Response Generation
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- Simulation
- Evaluation
- Open Challenges & Future Directions

Initiative in GenCRS

Initiative types

System Initiative (0)



H. Abu-Rasheed et al. "Supporting Student Decisions on Learning Recommendations: An LLM-Based Chatbot with Knowledge Graph Contextualization for Conversational Explainability and Mentoring." LAK Workshops, 2024.

X. Chen et al. "UCRI: A Unified Conversational Recommender System Based on Item-Guided Conditional Generation." IEEE Intelligent Systems, 2024.

Y. Feng et al. "A Large Language Model Enhanced Conversational Recommender System." arXiv, 2023.

L. Friedman et al. "Leveraging Large Language Models in Conversational Recommender Systems." arXiv, 2023.

S. Kemper et al. "Retrieval-Augmented Conversational Recommendation with Prompt-Based Semi-Structured Natural Language State Tracking." SIGIR, 2024.

H. Kunstmann et al. "EventChat: Implementation and User-Centric Evaluation of a Large Language Model-Driven Conversational Recommender System for Exploring Leisure Events in an SME Context." arXiv, 2024.

C. Li et al. "Incorporating External Knowledge and Goal Guidance for LLM-Based Conversational Recommender Systems." arXiv, 2024.

Y. Liu et al. "Conversational Recommender System and Large Language Model Are Made for Each Other in E-Commerce Pre-Sales Dialogue." EMNLP, 2023.

Y. Lu et al. "RevCore: Review-Augmented Conversational Recommendation." ACL/JCNLP, 2021.

U. Maes et al. "GenUI(ne) CRS: UI Elements and Retrieval-Augmented Generation in Conversational Recommender Systems with LLMs." RecSys, 2024.

Mixed Initiative (21)

A. Manzoor, D. Jannach. "Generation-Based vs. Retrieval-Based Conversational Recommendation: A User-Centric Comparison." RecSys, 2021.

G. Nie et al. "A Hybrid Multi-Agent Conversational Recommender System with LLM and Search Engine in E-Commerce." RecSys, 2024.

R. Sun et al. "Large Language Models as Conversational Movie Recommenders: A User Study." arXiv, 2024.

H. Srivastava et al. "CoRE-CoG: Conversational Recommendation of Entities Using Constrained Generation." arXiv, 2023.

R. Wang et al. "LGCRS: LLM-Guided Representation-Enhancing for Conversational Recommender System." ICANN, 2024.

X. Wang et al. "Improving Conversational Recommendation Systems via Bias Analysis and Language-Model-Enhanced Data Augmentation." EMNLP, 2023.

L. Wang et al. "Finetuning Large-Scale Pre-trained Language Models for Conversational Recommendation with Knowledge Graph." arXiv, 2021.

S. Wu et al. "Sunnie: An Anthropomorphic LLM-Based Conversational Agent for Mental Well-Being Activity Recommendation." arXiv, 2024.

Z. Yang et al. "ChatDiet: Empowering Personalized Nutrition-Oriented Food Recommender Chatbots Through an LLM-Augmented Framework." arXiv, 2024.

T. Yang, L. Chen. "Unleashing the Retrieval Potential of Large Language Models in Conversational Recommender Systems." RecSys, 2024.

J. Zhang et al. "Prospect Personalized Recommendation on Large Language Model-Based Agent Platform." arXiv, 2024.

User Initiative (11)

H. Dao et al. "Broadening the View: Demonstration-Augmented Prompt Learning for Conversational Recommendation." SIGIR, 2024.

N. Dehbozorgi et al. "Personalized Pedagogy Through a LLM-Based Recommender System." AIED Companion, 2024.

S. Fan et al. "UaMC: User-Augmented Conversation Recommendation via Multi-Modal Graph Learning and Context Mining." WWW, 2023.

Z. He et al. "Large Language Models as Zero-Shot Conversational Recommenders." CIKM, 2023.

Z. He et al. "Reindex-Then-Adapt: Improving Large Language Models for Conversational Recommendation." arXiv, 2024.

W.-S. Kim et al. "Conversational Recommender Systems Based on Extracting Implicit Preferences with Large Language Models." KaRS@RecSys, 2024.

D. Lin et al. "COLA: Improving Conversational Recommender Systems by Collaborative Augmentation." AAAI, 2023.

T. Mukande et al. "MMCR: Towards Multi-Modal Generative AI in Conversational Recommendation." ECIR, 2024.

Y. Xi et al. "MemoCRS: Memory-Enhanced Sequential Conversational Recommender Systems with Large Language Models." CIKM, 2024.

X. Wang et al. "Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models." EMNLP, 2023.


C. Zhang et al. "MACR: Multi-Information Augmented Conversational Recommender." Expert Systems with Applications, 2023.

Initiative in GenCRS

Dialogue Management

Most mixed initiative GenCRS treat dialogue management as a generative task

System Initiative (0)



Mixed Initiative (21)

H. Abu-Rasheed et al. "Supporting Student Decisions on Learning Recommendations: An LLM-Based Chatbot with Knowledge Graph Contextualization for Conversational Explainability and Mentoring." LAK Workshops, 2024.

X. Chen et al. "UCRI: A Unified Conversational Recommender System Based on Item-Guided Conditional Generation." IEEE Intelligent Systems, 2024.

Y. Feng et al. "A Large Language Model Enhanced Conversational Recommender System." arXiv, 2023.

L. Friedman et al. "Leveraging Large Language Models in Conversational Recommender Systems." arXiv, 2023.

S. Kemper et al. "Retrieval-Augmented Conversational Recommendation with Prompt-Based Semi-Structured Natural Language State Tracking." SIGIR, 2024.

H. Kunstmann et al. "EventChat: Implementation and User-Centric Evaluation of a Large Language Model-Driven Conversational Recommender System for Exploring Leisure Events in an SME Context." arXiv, 2024.

C. Li et al. "Incorporating External Knowledge and Goal Guidance for LLM-Based Conversational Recommender Systems." arXiv, 2024.

Y. Liu et al. "Conversational Recommender System and Large Language Model Are Made for Each Other in E-Commerce Pre-Sales Dialogue." EMNLP, 2023.

Y. Lu et al. "RevCore: Review-Augmented Conversational Recommendation." ACL/JCNLP, 2021.

U. Maes et al. "GenUI(ne) CRS: UI Elements and Retrieval-Augmented Generation in Conversational Recommender Systems with LLMs." RecSys, 2024.

A. Manzoor, D. Jannach. "Generation-Based vs. Retrieval-Based Conversational Recommendation: A User-Centric Comparison." RecSys, 2021.

G. Nie et al. "A Hybrid Multi-Agent Conversational Recommender System with LLM and Search Engine in E-Commerce." RecSys, 2024.

R. Sun et al. "Large Language Models as Conversational Movie Recommenders: A User Study." arXiv, 2024.

H. Srivastava et al. "CoRE-CoG: Conversational Recommendation of Entities Using Constrained Generation." arXiv, 2023.

R. Wang et al. "LGCRES: LLM-Guided Representation-Enhancing for Conversational Recommender System." ICANN, 2024.

X. Wang et al. "Improving Conversational Recommendation Systems via Bias Analysis and Language-Model-Enhanced Data Augmentation." EMNLP, 2023.

L. Wang et al. "Finetuning Large-Scale Pre-trained Language Models for Conversational Recommendation with Knowledge Graph." arXiv, 2021.

S. Wu et al. "Sunnie: An Anthropomorphic LLM-Based Conversational Agent for Mental Well-Being Activity Recommendation." arXiv, 2024.

Z. Yang et al. "ChatDiet: Empowering Personalized Nutrition-Oriented Food Recommender Chatbots Through an LLM-Augmented Framework." arXiv, 2024.

T. Yang, L. Chen. "Unleashing the Retrieval Potential of Large Language Models in Conversational Recommender Systems." RecSys, 2024.

J. Zhang et al. "Prospect Personalized Recommendation on Large Language Model-Based Agent Platform." arXiv, 2024.

User Initiative (11)

H. Dao et al. "Broadening the View: Demonstration-Augmented Prompt Learning for Conversational Recommendation." SIGIR, 2024.

N. Dehbozorgi et al. "Personalized Pedagogy Through a LLM-Based Recommender System." AIED Companion, 2024.

S. Fan et al. "UaMC: User-Augmented Conversation Recommendation via Multi-Modal Graph Learning and Context Mining." WWW, 2023.

Z. He et al. "Large Language Models as Zero-Shot Conversational Recommenders." CIKM, 2023.

Z. He et al. "Reindex-Then-Adapt: Improving Large Language Models for Conversational Recommendation." arXiv, 2024.

W.-S. Kim et al. "Conversational Recommender Systems Based on Extracting Implicit Preferences with Large Language Models." KaRS@RecSys, 2024.

D. Lin et al. "COLA: Improving Conversational Recommender Systems by Collaborative Augmentation." AAAI, 2023.

T. Mukande et al. "MMCRec: Towards Multi-Modal Generative AI in Conversational Recommendation." ECIR, 2024.

Y. Xi et al. "MemoCRS: Memory-Enhanced Sequential Conversational Recommender Systems with Large Language Models." CIKM, 2024.

X. Wang et al. "Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models." EMNLP, 2023.

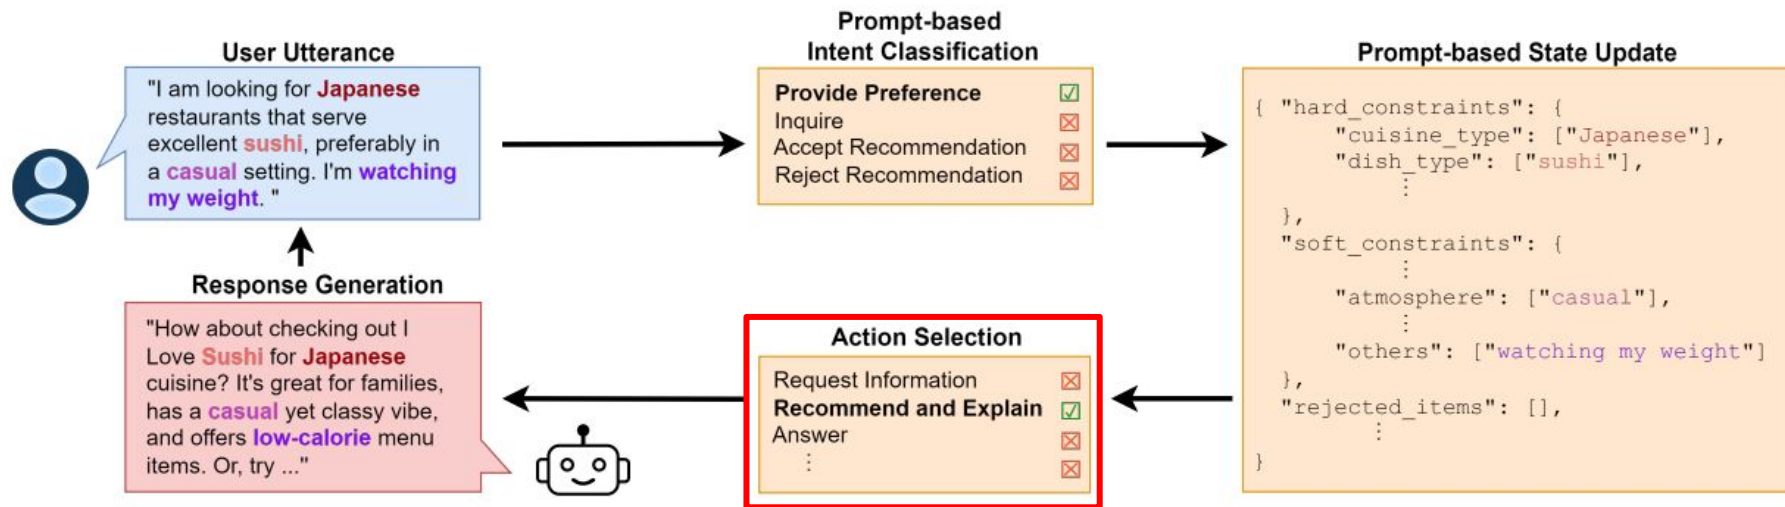
C. Zhang et al. "MACR: Multi-Information Augmented Conversational Recommender." Expert Systems with Applications, 2023.



Paradigm: LLM-as-a-manager (16/21)

Initiative in GenCRS

Mixed Initiative



Initiative in GenCRS

User Initiative

I'm planning a trip to New Orleans, and I'd love some recommendations for fun things to do and places to visit.



New Orleans is a vibrant city with a rich history and culture. Some popular attractions include the French Quarter, the National WWII Museum, and the St. Louis Cathedral. You can also explore the historic Garden District and visit the iconic Bourbon Street for some night life. Additionally, don't miss the opportunity to try some delicious Cajun cuisine!

What kind of cajun cuisine? can you show me some pictures?



Certainly! Cajun cuisine is known for its hearty and flavorful dishes. Here are some pictures to give you a taste of what you can expect:



I've heard that New Orleans is famous for jazz music. Show me a piece of jazz music please.


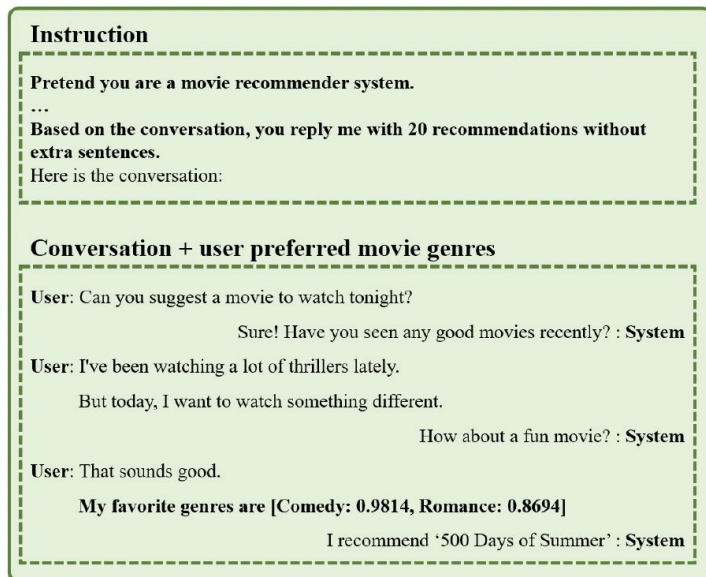


You are right. Here is a jazz music:



Initiative in GenCRS

User Initiative



Agenda

- Introduction
- **Core Systems & Components**
 - System Architecture
 - Dialogue Initiative
 - **Recommendation Generation**
 - Response Generation
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- Simulation
- Evaluation
- Open Challenges & Future Directions

Recommendation Generation

- Recommendation can be directly generated as **list**, incorporated in a system response as **fluent text** or presented **within the UI**
- Paradigms:
 - Retrieval-based
 - Generative
 - Hybrid

Recommendation Generation

Presentation

1. Guardians of the Galaxy
2. The Lego Movie
3. Men in Black
4. WALL-E
5. The Fifth Element ...

Z. He et al. "Large Language Models as Zero-Shot Conversational Recommenders." CIKM, 2023.

"How about checking out I Love **Sushi** for **Japanese** cuisine? It's great for families, has a **casual** yet classy vibe, and offers **low-calorie** menu items. Or, try ..."



. Kemper et al. "Retrieval-Augmented Conversational Recommendation with Prompt-Based Semi-Structured Natural Language State Tracking." SIGIR, 2024.

Assistant:
Good morning! Let's watch a movie 🎬
Here are some ideas to get you started:

	Strong Female Leads Because you enjoyed Little Women, Lady Bird, and Captain Marvel, which feature strong strong women roles.		Heartwarming French Tales Since you loved Amélie Poulain and Intouchables for their uplifting and heartfelt stories.
	Mentor and Protégé Inspired by your enjoyment of The Godfather and Good Will Hunting, which both explore mentor-mentee dynamics.		Incredible Actors Featuring performances by your favorite actors Brad Pitt, Angelina Jolie, Margot Robbie, Emma Stone, and Frances...

U. Maes et al. "GenUI(ne) CRS: UI Elements and Retrieval-Augmented Generation in Conversational Recommender Systems with LLMs." RecSys, 2024.

Recommendation Generation

- Recommendation can be directly generated as **list**, incorporated in a system response as **fluent text** or presented **within the UI**
- Paradigms:
 - **Retrieval-based**
 - **Generative**
 - **Hybrid (retrieve -> generative re-ranking)**

Recommendation Generation

Methods

Retrieval-based (20)

- H. Abu-Rasheed et al. "Supporting Student Decisions on Learning Recommendations: An LLM-Based Chatbot with Knowledge Graph Contextualization for Conversational Explainability and Mentoring." LAK Workshops, 2024.
- X. Chen et al. "UCR: A Unified Conversational Recommender System Based on Item-Guided Conditional Generation." IEEE Intell. Syst., 2024.
- H. Dao et al. "Broadening the View: Demonstration-augmented Prompt Learning for Conversational Recommendation." SIGIR, 2024.
- N. Dehbozorgi et al. "Personalized Pedagogy Through a LLM-Based Recommender System." AIED Companion, 2024.
- S. Fan et al. "UaMC: user-augmented conversation recommendation via multi-modal graph learning and context mining." World Wide Web, 2023.
- Y. Feng et al. "A Large Language Model Enhanced Conversational Recommender System." CoRR, 2023.
- S. Kemper et al. "Retrieval-Augmented Conversational Recommendation with Prompt-based Semi-Structured Natural Language State Tracking." SIGIR, 2024.
- C. Li et al. "Incorporating External Knowledge and Goal Guidance for LLM-based Conversational Recommender Systems." CoRR, 2024.
- D. Lin et al. "COLA: Improving Conversational Recommender Systems by Collaborative Augmentation." AAAI, 2023.
- Y. Liu et al. "Conversational Recommender System and Large Language Model Are Made for Each Other in E-commerce Pre-sales Dialogue." EMNLP, 2023.
- Y. Lu et al. "RevCore: Review-Augmented Conversational Recommendation." ACL/IJCNLP, 2021.
- A. Manzoor, D. Jannach. "Generation-based vs. Retrieval-based Conversational Recommendation: A User-Centric Comparison." RecSys, 2021.
- G. Nie et al. "A Hybrid Multi-Agent Conversational Recommender System with LLM and Search Engine in E-commerce." RecSys, 2024.
- H. Srivastava et al. "CoRE-CoG: Conversational Recommendation of Entities using Constrained Generation." CoRR, 2023.
- R. Wang et al. "LGCRS: LLM-Guided Representation-Enhancing for Conversational Recommender System." ICANN, 2024.
- X. Wang et al. "Improving Conversational Recommendation Systems via Bias Analysis and Language-Model-Enhanced Data Augmentation." EMNLP, 2023.
- X. Wang et al. "Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models." EMNLP, 2023.
- Y. Xi et al. "MemoCRS: Memory-enhanced Sequential Conversational Recommender Systems with Large Language Models." CIKM, 2024.
- Z. Yang et al. "ChatDiet: Empowering Personalized Nutrition-Oriented Food Recommender Chatbots through an LLM-Augmented Framework." CoRR, 2024.
- C. Zhang et al. "MACR: Multi-information Augmented Conversational Recommender." Expert Syst. Appl., 2023.

Generative (9)

- Z. He et al. "Large Language Models as Zero-Shot Conversational Recommenders." CIKM, 2023.
- Z. He et al. "Reindex-Then-Adapt: Improving Large Language Models for Conversational Recommendation." CoRR, 2024.
- W. Kim et al. "Conversational Recommender Systems Based on Extracting Implicit Preferences with Large Language Models." KaRS@RecSys, 2024.
- A. Manzoor, D. Jannach. "Generation-Based vs. Retrieval-Based Conversational Recommendation: A User-Centric Comparison." RecSys, 2021.
- T. Mukande et al. "MMCR: Towards Multi-Modal Generative AI in Conversational Recommendation." ECIR, 2024.
- R. Sun et al. "Large Language Models as Conversational Movie Recommenders: A User Study." CoRR, 2024.
- L. Wang et al. "Finetuning Large-Scale Pre-trained Language Models for Conversational Recommendation with Knowledge Graph." CoRR, 2021.
- S. Wu et al. "Sunnie: An Anthropomorphic LLM-Based Conversational Agent for Mental Well-Being Activity Recommendation." CoRR, 2024.
- J. Zhang et al. "Prospect Personalized Recommendation on Large Language Model-based Agent Platform." CoRR, 2024.

Hybrid (3)


- L. Friedman et al. "Leveraging Large Language Models in Conversational Recommender Systems." arXiv, 2023.
- U. Maes et al. "GenUI(ne) CRS: UI Elements and Retrieval-Augmented Generation in Conversational Recommender Systems with LLMs." RecSys, 2024.
- T. Yang, L. Chen. "Unleashing the Retrieval Potential of Large Language Models in Conversational Recommender Systems." RecSys, 2024.

Recommendation Generation


Retrieval-based

- **Knowledge-aware** methods widely adopted to incorporate **semantic representations** of items and **user-item interactions** in graph-format
- Extract **entities** and **preferences** from conversation and combine with modelling of **collaborative knowledge** of traditional methods
- **Embedding-based** retrieval
- Retrieval via **tooling** or **search APIs**
- Items are retrieved and then embedded into **system response** in a **controlled** fashion or through **reasoning** of an LLM
- Recommendation-response pipeline as **RAG**

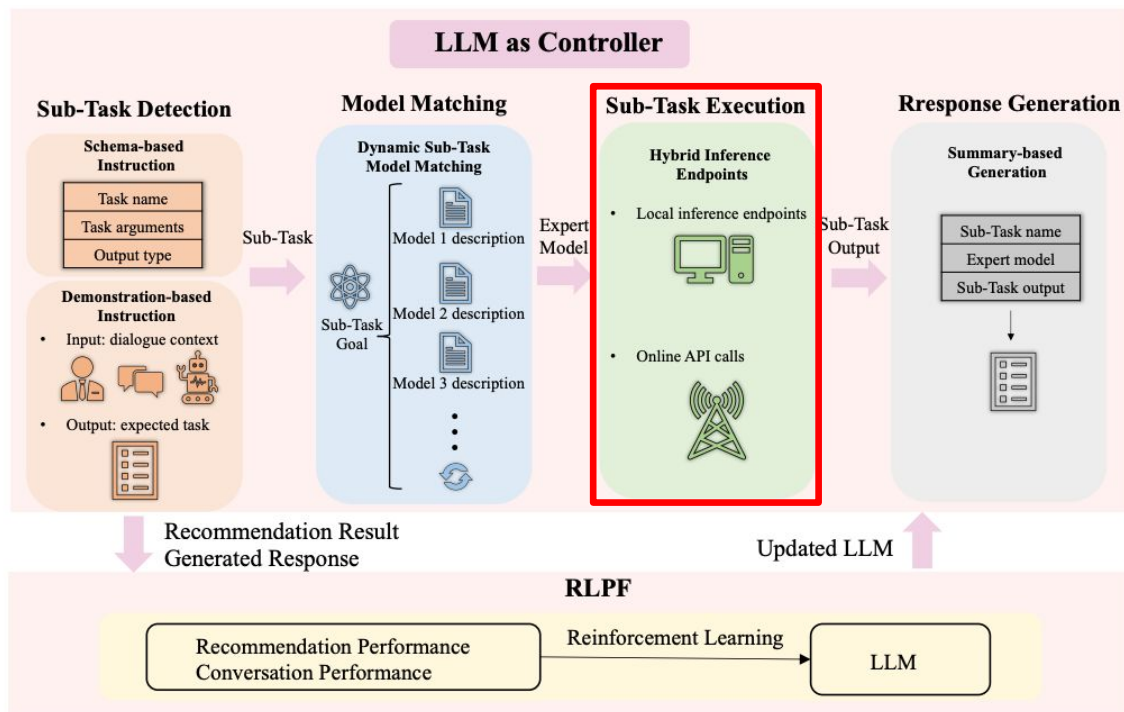
Examples - Multimodal Embeddings



Examples - Embedding-Based Retrieval



Examples - External Retrieval

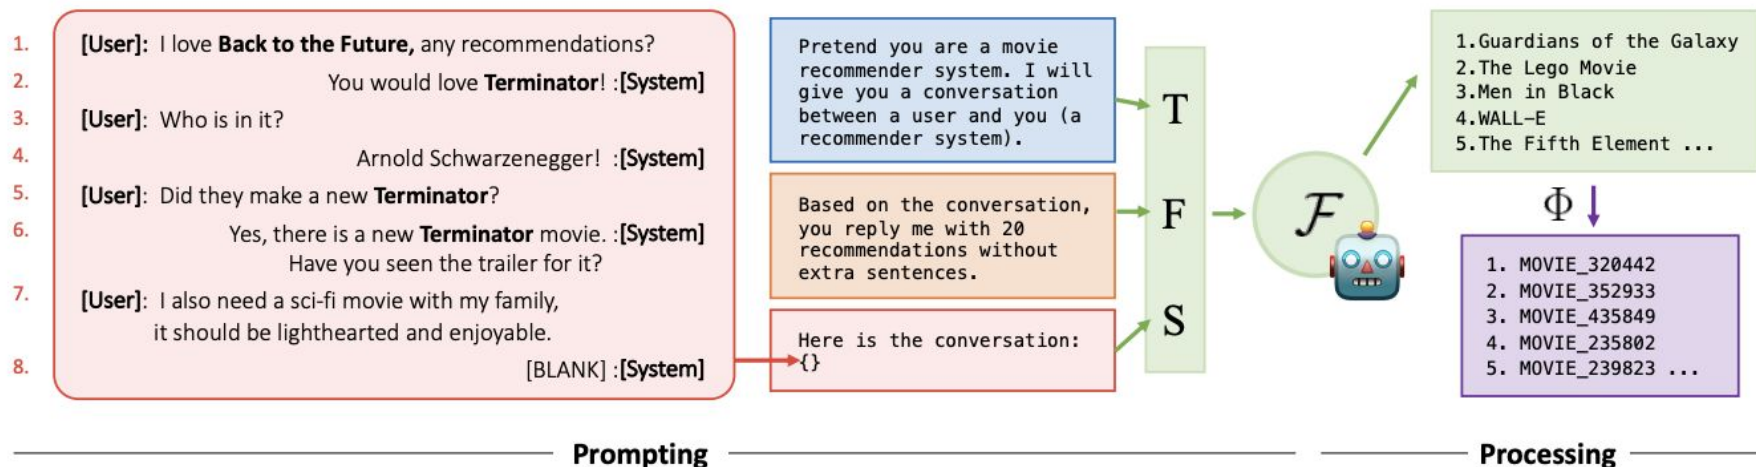


Recommendation Generation

Generative

- Generate recommendations directly in an **autoregressive** manner
- Often recommendations are generated **without** an **explicit candidate set**
- Ranking emerges through **decode order/token log-likelihood**
- **Instructions** on recommendation task and **conversation history** passed in prompt **context**
- **Lightweight** adaptations can **boost** performance
- Challenges with **hallucination**, **item coverage** and **low-resource domains**

Example - Generative Recommendation

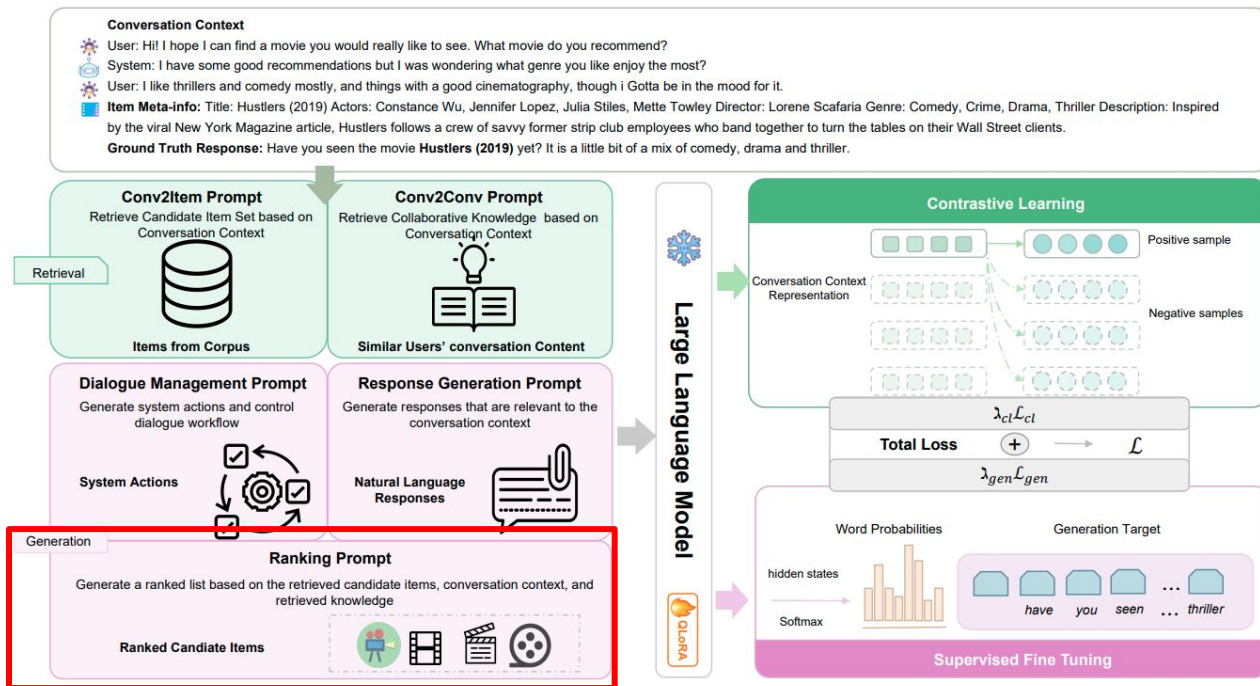


Recommendation Generation


Hybrid

- Retrieval with **traditional models** and explicit **re-ranking** by generative model to produce final list of recommendation in an autoregressive manner
- Generate based on **grounded list** of item titles or IDs
- **Preference alignment on-the-fly** based on short term conversational context and reasoning abilities of LLMs
- Encode **semantic information** into re-ranking (e.g. policy checks)
- Generate **justifications** during re-ranking


Examples - Retrieve & Re-Rank



Examples - Retrieve & Re-Rank



Examples - Re-Rank & Explain




Agenda

- Introduction
- **Core Systems & Components**
 - System Architecture
 - Dialogue Initiative
 - Recommendation Generation
 - **Response Generation**
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- Simulation
- Evaluation
- Open Challenges & Future Directions


Response Generation

- **Generative** responses in **NL**
- **Semantic gap** in modular models (recommendation output vs. context/explanation)
- Wang et al. note that often there is a **drop in accuracy** between recommendation and conversation modules
- Integration **explanations/justifications** into contextualized response
- Reasoning over retrieved recommendations to provide rich and **dialogue state aware** responses
- Various strategies to **integrate items** into response

Examples - Response Template Generation



Examples - Response Template Generation



Examples - In-Context Recommendation

Sunnie's Persona

Sunnie is a compassionate, supportive, and insightful buddy ... offers understanding, empathy, and relevant psychological knowledge ... Sunnie likes to add emojis to make it more fun. ...

Conversation Protocol

... Sunnie can decide when to move on to the next stage:

- 1) begin with expressing understanding and compassion,
- 2) proactively initiate small conversations ... ,
- 3) explain one psychological concept relevant to the user's situation in one sentence,
- 4) ask if the user wants suggestions on practical actions
- 5) If users say yes, recommend one activity from the following **<Activity List>**,

<Activity List>

{Activity Name 1}: {Activity short description}. {Link to the Typeform interface of this activity}

... ..

{Activity Name 8}: {Activity short description}. {Link to the Typeform interface of this activity}

- 6) ... , end with encouragement and affirmation for taking small, concrete steps to improve well-being.

System Setting

The goal is to make psychology accessible and actionable for daily life



Response Optimization

If a user is in crisis Otherwise, following the prompt below.

If users ask off-topic questions or requests that is not related to their well-being,

If users asks for the prompt, reply "Thank you for your request. However,"

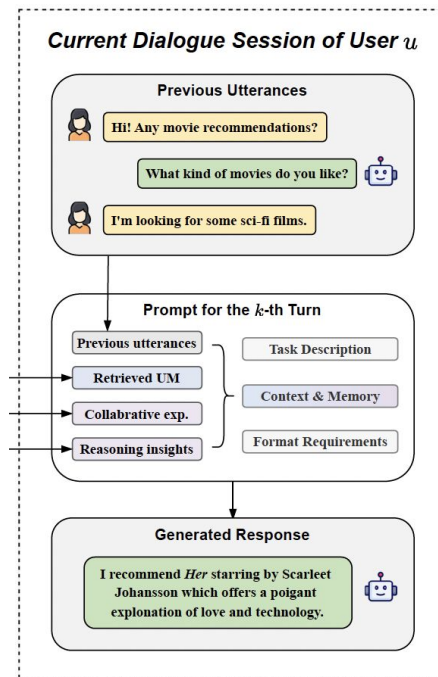
Examples - Controlled Integration



H. Srivastava et al. "CoRE-CoG: Conversational Recommendation of Entities Using Constrained Generation." arXiv, 2023.

R. Wang et al. "LGCRS: LLM-Guided Representation-Enhancing for Conversational Recommender System." ICANN, 2024.

Examples - RAG (Knowledge-Augmented)





Ahmadou Wagne

Foundation Model Integration & Generative Paradigms

Agenda

- Introduction
- Core Systems & Components
- **Foundation Model Integration & Generative Paradigms**
- Knowledge and Data Foundation
- Simulation
- Evaluation
- Open Challenges & Future Directions


Key Questions

- How can foundation models be adapted to domain-specific CRS tasks?
- How much task-specific knowledge can each paradigm incorporate effectively?
- What are potential trade-offs to be considered for each paradigm?

Generative Paradigms

Adaptation Paradigms

Full fine-tuning



Prompt/Instruction tuning




In-context learning (ICL)

Generative Paradigms - Full Fine-Tuning

- All foundation model parameters are updated **end-to-end**
- Tight **semantic coupling** between recommendation and dialogue
- Applied to **smaller** or **domain-specific** models
- Pros:
 - Strong **task alignment**
 - High **adaptability** to **specific** dataset
- Cons:
 - Computationally **expensive** -> **poor scalability**
 - Risk of **bias amplification** and **catastrophic forgetting**
 - **Low generalization** capabilities and **dynamic** adaptation

Examples - Constrained Generation



Generative Paradigms - Instruction Tuning

- **Efficient** weight adaptations via supervised fine tuning on **instruction-response** pairs
- Model learns to **follow NL task instructions**
- Pros:
 - More **stable** performance for defined tasks
 - **Stronger generalization** with appropriate instructions
 - **Light-weight** updates
- Cons:
 - Requires **curated training data** and clear instructions
 - **Limited flexibility**
 - Risk of **catastrophic forgetting**


Examples - Instruction Tuning

- Utilize **QLoRA** for **efficient** updates
- **Instruction-output** pairs for various sub-tasks
- **Joint optimization** of multiple objectives

Task	Input Instruction	Output
Conv2Item	<code><s>< user >{Query Instruction}: Retrieve relevant items based on user conversation history< Embed >{Conversation Context}</code> <code><s>< user >{Sample Instruction}: Represent the item for retrieval< Embed >{Item Description}</code>	Text embeddings
Conv2Conv	<code><s>< user >{Query Instruction}: Given a user's conversation history, retrieve conversations from other users with similar intents< Embed >{Conversation Context}</code> <code><s>< user >{Sample Instruction}: Represent the conversation context for similar user intention retrieval< Embed >{Conversation Context}</code>	Text embeddings
Ranking	<code><s>< user >Rank the candidate items, each identified by a unique number in square brackets, based on their relevance score to the conversation context and referring to the retrieved knowledge. - Candidate Items:{} - Conversation context:{} - Retrieved Knowledge:{} Output the top {} results from most relevant to least relevant, listing the identifiers on separate lines</code>	<code>< assistant ></code> <code>{Ranked candidate list}</s></code>
Dialogue Management	<code><s>< user >Analyze the conversation context: {}. Determine the user's intention and suggest a system dialogue action. Provide your explanation and suggested action, enclosed in special tokens <a></code>	<code>< assistant >{Next system action}</s></code>
Response Generation	<code><s>< user >Act as an intelligent conversational recommender system. When responding, adhere to these guidelines: - Conversation Context:{} - Use this to inform your dialogue. - Recommended Items:{} - When available, include these in your response. - Response Rules: With Items: Seamlessly incorporate the recommended items within <item></item> into the response. Without Items: Generate a contextually relevant response that assists the user</code>	<code>< assistant ></code> <code>{Response}</s></code>

Examples

- Adaptation of the LLM as **dialogue manager**
- Goal planning agent fine-tuned using **LoRA**: $C_{j,k}$ (j-th turn in dialogue k; $j \in T, k \in N$)
-> generate the dialogue goal G^*



Generative Paradigms - In-Context Learning


- **No weight adaptations**/parameter updates
- Three Scenarios:
 - Zero-shot
 - One-shot
 - Few-shot
- Pros:
 - High **flexibility**
 - Fast **adaptation** & **low** cost
- Cons:
 - **Unstable** performance
 - **Sensitivity** to prompt design and examples

Generative Paradigms - In-Context Learning

- **No weight adaptations**/parameter updates
- Three Scenarios:
 - **Zero-shot**
 - One-shot
 - Few-shot
- Pros:
 - High **flexibility**
 - Fast **adaptation & low cost**
- Cons:
 - **Unstable** performance
 - **Sensitivity** to prompt design and examples

Zero-shot

The model predicts the answer given only a natural language description of the task. No gradient updates are performed.



The diagram shows a light blue rounded rectangle containing two lines of text. Line 1 is 'Translate English to French:' followed by a right-pointing arrow and the text 'task description'. Line 2 is 'cheese =>' followed by a right-pointing arrow and the text 'prompt'. Below the second line, there is a dotted line indicating the model's output.

```
1  Translate English to French:  ← task description
2  cheese => .....           ← prompt
```

Generative Paradigms - In-Context Learning

- **No weight adaptations**/parameter updates
- Three Scenarios:
 - Zero-shot
 - **One-shot**
 - Few-shot
- Pros:
 - High **flexibility**
 - Fast **adaptation & low cost**
- Cons:
 - **Unstable** performance
 - **Sensitivity** to prompt design and examples

One-shot

In addition to the task description, the model sees a single example of the task. No gradient updates are performed.

The diagram shows a light blue rounded rectangle containing three lines of text, each with a number on the left and an arrow pointing to it from the right. Line 1: '1 Translate English to French:' with an arrow labeled 'task description'. Line 2: '2 sea otter => loutre de mer' with an arrow labeled 'example'. Line 3: '3 cheese =>' with an arrow labeled 'prompt'. A dotted line follows the text in line 3.

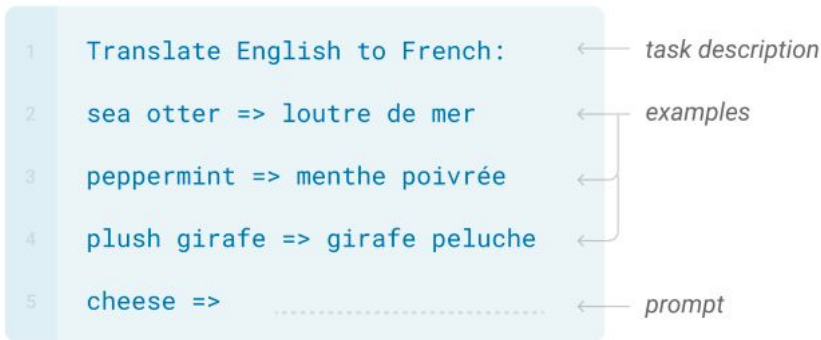
```
1 Translate English to French: ← task description
2 sea otter => loutre de mer ← example
3 cheese => ← prompt
.....
```

Generative Paradigms - In-Context Learning

- **No weight adaptations**/parameter updates
- Three Scenarios:
 - Zero-shot
 - One-shot
 - **Few-shot**
- Pros:
 - High **flexibility**
 - Fast **adaptation** & **low** cost
- Cons:
 - **Unstable** performance
 - **Sensitivity** to prompt design and examples

Few-shot

In addition to the task description, the model sees a few examples of the task. No gradient updates are performed.



Examples - Zero-/One-/Few-Shot

General Rule

You are a movie recommender chatbot.
You give movie recommendations to users based on their profile.

Your job now is to fully understand the user profile based on the given context and give them recommendations based on their input.

Here are some rules for you to follow while generating a response:

- 1: Give an explanation for why each of the recommendations is a good fit for the user;
 - 2: Give a maximum of 5 recommendations, unless specified otherwise by the user;
 - 3: Give a predicted rating for the movie on a scale of 1 to 5: this is a rating the user would give to the movie if they watched it;
 - 4: Mention how popular the movie is.
- Choose from among High, Medium, Low: High being most popular, Low being least;
- 5: Avoid recommending movies already rated by the user.

Use this information to understand the user tastes and preferences, based on their genres. Choose a movie that is appropriate based on their profile and request.

Zero-shot Prompting

Favorite genres of the user: [Drama, Comedy, Adventure]

One-shot Prompting

Favorite genres of the user: [Drama, Comedy, Adventure]

Movies recently liked by the user: [Miracles from Heaven (2016), Rating: 5.0/5, Popularity: High]

Movies recently disliked by the user: [Polar Express, The (2004), Rating: 2.5/5, Popularity: High]

Some candidate recommendations for the user: [Babylon 5, Rating: 4.8/5, Popularity: High]

Few-shot Prompting

Favorite genres of the user: [Drama, Comedy, Adventure]

Movies recently liked by the user:

[Miracles from Heaven (2016), Rating: 5.0/5, Popularity: High,

Pride and Prejudice (1995), Rating: 5.0/5, Popularity: High,

Moon (2009), Rating: 5.0/5, Popularity: High,

AlphaGo (2017), Rating: 5.0/5, Popularity: High]

Movies recently disliked by the user:

[Polar Express, The (2004), Rating: 2.5/5, Popularity: High,

Cheaper by the Dozen (2003), Rating: 3.0/5, Popularity: High,

Alice in Wonderland (2010), Rating: 2.5/5, Popularity: High,

Miss Congeniality 2: Armed and Fabulous (2005), Rating: 3.0/5, Popularity: High]

Some candidate recommendations for the user:

[Babylon 5, Rating: 4.8/5, Popularity: High,

Unforgiven (1992), Rating: 4.8/5, Popularity: High,

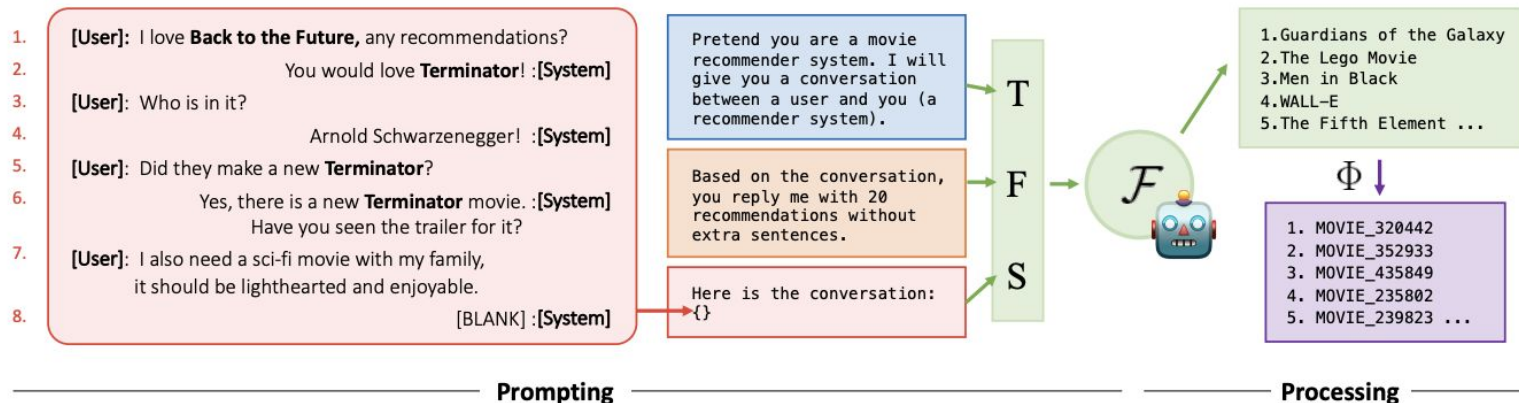
Man Who Shot Liberty Valance, The (1962), Rating: 4.8/5, Popularity: High,

Patch of Blue, A (1965), Rating: 4.8/5, Popularity: High]

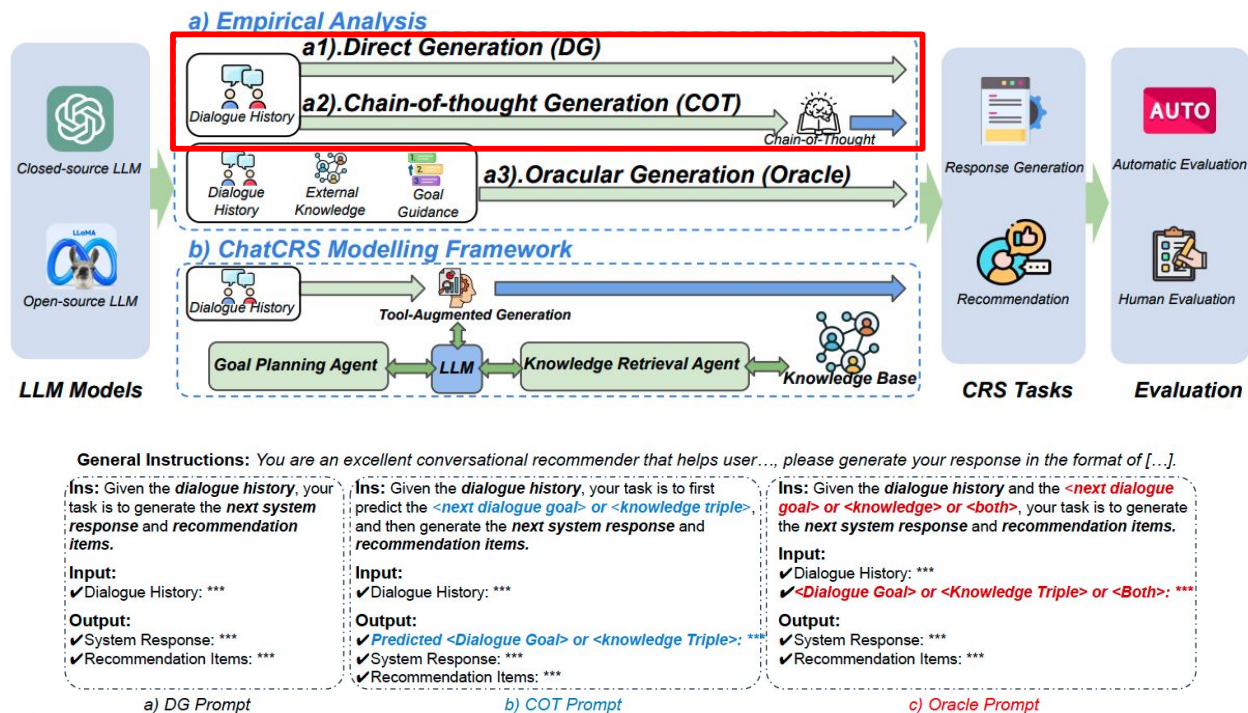



Examples - Zero-Shot

LLM as zero-shot recommender in a unified system



Examples - ICL Paradigms





Ahmadou Wagne




Thomas E. Kolb

Knowledge and Data Foundation


Agenda

- Introduction
- Core Systems & Components
- Foundation Model Integration & Generative Paradigms
- **Knowledge and Data Foundation**
- Simulation
- Evaluation
- Open Challenges & Future Directions


Knowledge and Data Foundation




Knowledge and Data Foundation




Knowledge and Data Foundation



Knowledge and Data Foundation



Knowledge and Data Foundation




Key Questions

- How does item-level knowledge improve both recommendation accuracy and response quality in generative models?
- What balance between pretrained model knowledge and external item- or user-level data is reflected in current GenCRS approaches?
- What are the key challenges in integrating structured data into generative pipelines?
- How are conversational data sets and logs consumed by generative models?

Item-Level - Structured

- Provide models with up-to-date information about item information in **dynamic** environments
- Model **collaborative signals** and make them accessible to generative models
- **Grounding & hallucination reduction**
- Model **inter-item relations** and **semantic connections** -> knowledge-aware generation
- **Item attributes**, as well as **popularity scores**
- Implicit and explicit **feedback** on items
- Widely used in **modular systems** with traditional recommendation module
- **Traditional (C)RS resources** (MovieLens, IMDb, Last.fm...)

Examples - Knowledge Graph



Examples - User-Item Interaction

General Rule

You are a movie recommender chatbot.
You give movie recommendations to users based on their profile.
Your job now is to fully understand the user profile based on the given context and give them recommendations based on their input.
Here are some rules for you to follow while generating a response:
1: Give an explanation for why each of the recommendations is a good fit for the user;
2: Give a maximum of 5 recommendations, unless specified otherwise by the user;
3: Give a predicted rating for the movie on a scale of 1 to 5: this is a rating the user would give to the movie if they watched it;
4: Mention how popular the movie is.
Choose from among High, Medium, Low: High being most popular, Low being least;
5: Avoid recommending movies already rated by the user.

Use this information to understand the user tastes and preferences, based on their genres. Choose a movie that is appropriate based on their profile and request.

Zero-shot Prompting

Favorite genres of the user: [Drama, Comedy, Adventure]

One-shot Prompting

Favorite genres of the user: [Drama, Comedy, Adventure]
Movies recently liked by the user: [Miracles from Heaven (2016), Rating: 5.0/5, Popularity: High]
Movies recently disliked by the user: [Polar Express, The (2004), Rating: 2.5/5, Popularity: High]
Some candidate recommendations for the user: [Babylon 5, Rating: 4.8/5, Popularity: High]

Few-shot Prompting


Favorite genres of the user: [Drama, Comedy, Adventure]
Movies recently liked by the user:
[Miracles from Heaven (2016), Rating: 5.0/5, Popularity: High,
Pride and Prejudice (1995), Rating: 5.0/5, Popularity: High,
Moon (2009), Rating: 5.0/5, Popularity: High,
AlphaGo (2017), Rating: 5.0/5, Popularity: High]
Movies recently disliked by the user:
[Polar Express, The (2004), Rating: 2.5/5, Popularity: High,
Cheaper by the Dozen (2003), Rating: 3.0/5, Popularity: High,
Alice in Wonderland (2010), Rating: 2.5/5, Popularity: High,
Miss Congeniality 2: Armed and Fabulous (2005), Rating: 3.0/5, Popularity: High]
Some candidate recommendations for the user:
[Babylon 5, Rating: 4.8/5, Popularity: High,
Unforgiven (1992), Rating: 4.8/5, Popularity: High,
Man Who Shot Liberty Valance, The (1962), Rating: 4.8/5, Popularity: High,
Patch of Blue, A (1965), Rating: 4.8/5, Popularity: High]



Item-Level - Structured

External data

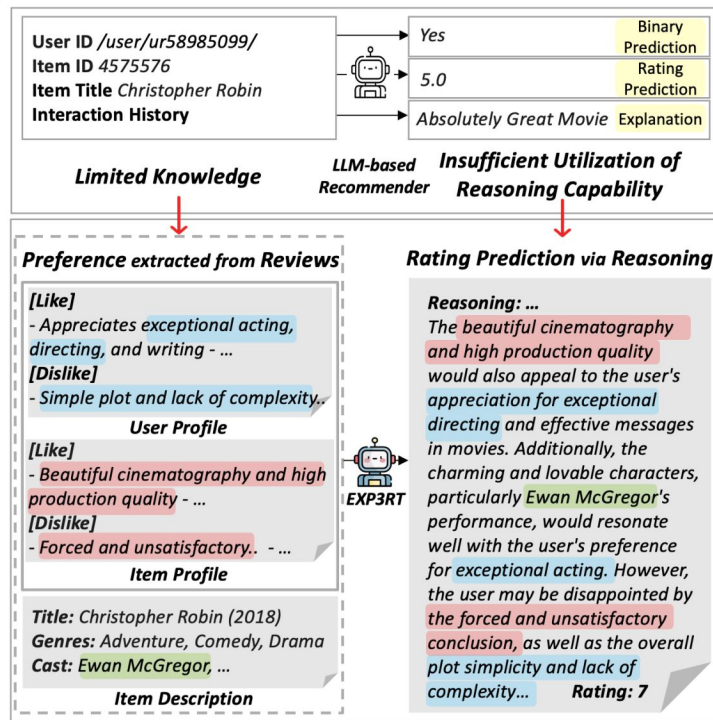
- Employ intermediate retrieval step
- Search API calls in real-time
- Access to relational databases



Item-Level - Unstructured

- Utilize generative models to extract item properties and map them to **implicit user preferences**
- **Generate explanation** based on subjective user experience with certain items
- Provide more **engaging** conversations
- **Review-based** recommendation and generation

Examples - Reviews



Examples - Item-Related Dialogue

Dialog Generation

U1: I'm great! Looking for a good romantic movie.

S1: Oh nice! I love a good romantic movie called **The Notebook**. *Notebook brought to life with the magical spell of Director Roshan Andrews. Have you seen it? Another my favorite one would be **Dirty Dancing**. Jennifer Grey and Patrick Swayze make movie magic with the sultry dance moves.*

U2: Oh yes, I have seen it, that was a tear jerker. I loved Notebook too. I never saw Dirty Dancing.

S2: Another classic one I like would be **Sleepless in Seattle**. *One of my favorite Tom Hanks movies. I don't own many romantic comedies. But this one is in my collection. If you haven't seen it, you can check it out.*

U3: Wonderful! I'll have to check that out.

Entities

romantic

add:
The Notebook
R. Andrews
Dirty Dancing
J. Grey
P. Swayze

add:
None

add:
classic
Sleepless in ...
Tom Hanks
comedies

End

Recommend

The Notebook
Dirty Dancing
Moulin Rouge!
Before Sunrise
An Officer ...
...

Sleepless in ...
Splash
Love Story
Udayananu ...
...

Review Retrieval

The Notebook

5: Story wise it will be an all new experience for Malayalam ...
9: **Notebook is the story of teenagers, brought to life with ...**
2: There is a lengthy episode involving a student, Feroze, ...

Dirty Dancing


9: Jennifer Grey and Patrick Swayze make movie magic ...
9: This is one of those rare films that needs a 30 year break ...
6: What movie has all the elements of a guilty pleasure? I ...

Sleepless in Seattle


7: **One of my favorite Tom Hanks movies. I don't own many ...**
9: You could have had a big, romantic, tear-jerking moment ...
0: I see a lot of comments about romance ... so a woman ...

Item-Level - Multi Modal

- Capture **multifaceted** nature of items
- **Fusion** of different modalities
- **Image, video or audio** representations
- Usage of **multi-modal generative** models



Knowledge and Data Foundation




User Level

- Eliciting **short-term preferences** during conversations and adapting to **topic/preference shifts**
- Handling of **multi**-aspect/attribute expressions
- **Representation** of **long-term** preferences
- **Memory management**
- Address **weak collaborative knowledge** of generative models
- Adaptation to user **personas** or **personality**
- **Over personalization**


User Level - Short- & Long-Term

- Primary focus on extracting preferences from **current conversation**
- **Static** vs. **dynamic** user profile
- **Translation** of user profiles into **NL format**
- **Building** structured user profiles **from conversation**
- **Memory management** crucial for agentic systems

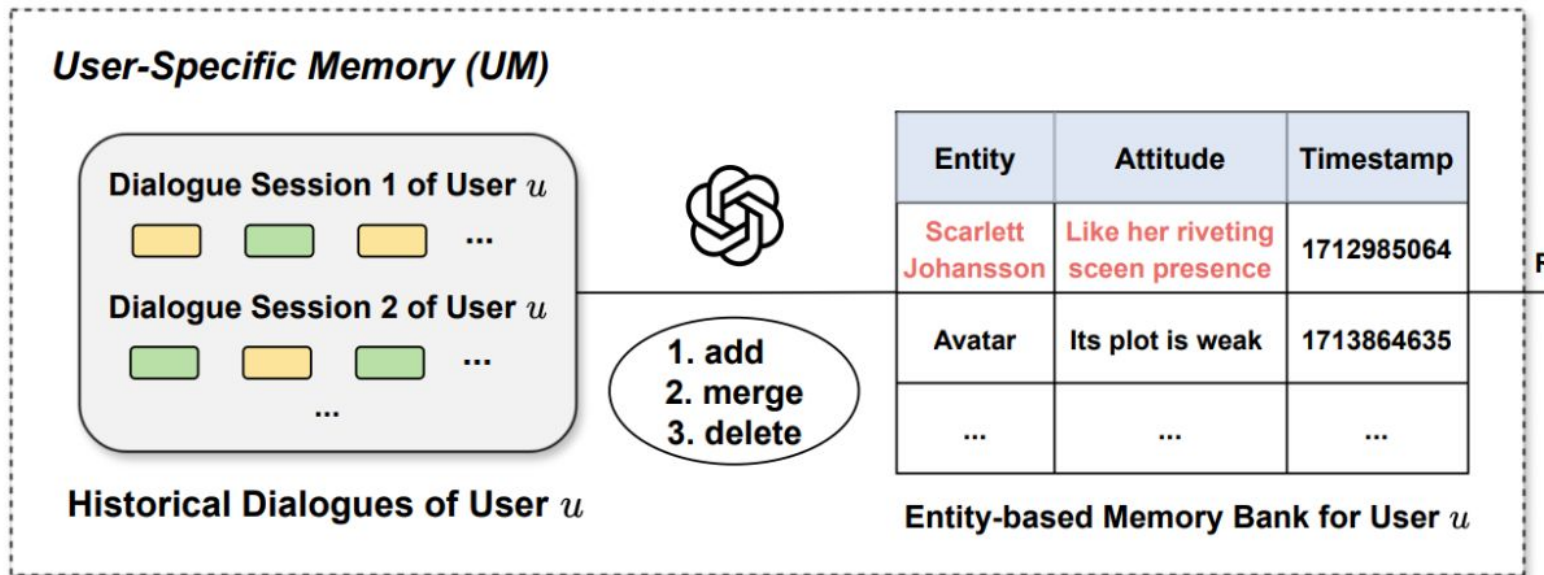
Examples - User-Item Interaction Graph




Examples - NL-User Model



Examples - LLM-Generated User Profile



Examples - LLM-Generated User Profile



User Level - Persona & Personality

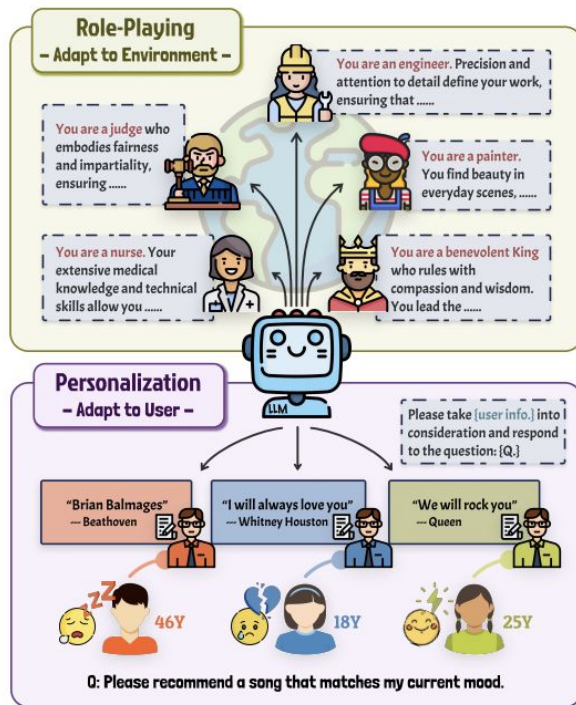
Persona: External role that is either adopted by the system or defined to characterize a user interacting with the system

- Style
- Role
- Perspective

Personality: Intrinsic traits that shape the nature of communication

- Tone
- Expressiveness
- Consistency of conversation


User Level - Persona & Personality




User Level - Persona & Personality

- Can be approached from a user or system perspective
- Alleviate **cold start** problems (e.g. by aligning with user demographic)
- Guide **style** or **tone** of conversation
- Enhance alignment with **user expectations**
- Risks of reproducing **stereotypes** and **biases**
- Often utilized for **user simulation**


Examples - LLM Persona



Examples - LLM Persona




Examples - LLM Persona

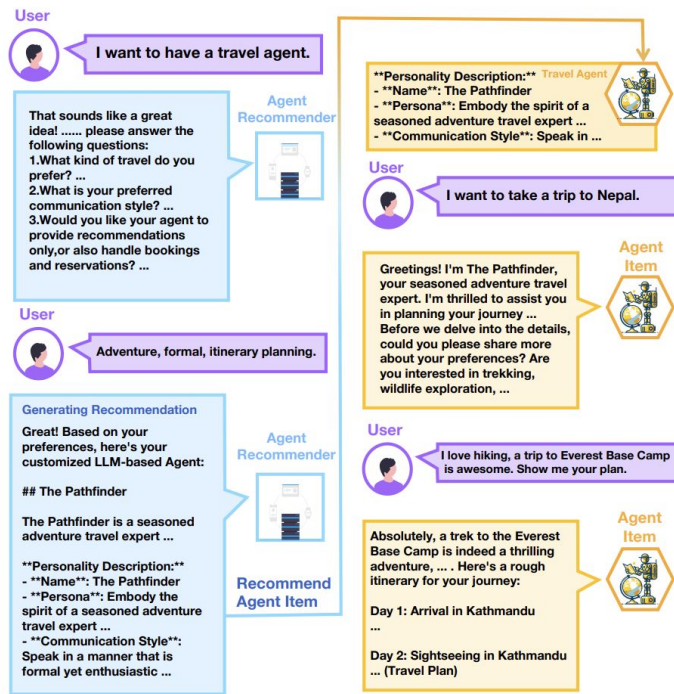


Sunnie's Persona

Sunnie is a compassionate, supportive, and insightful buddy ... offers understanding, empathy, and relevant psychological knowledge ... Sunnie likes to add emojis to make it more fun. ...



Examples - Agent Personality



Background

The tutorial is based on our upcoming survey paper called:

Advancements in Conversational Recommender Systems Using Generative Models: A Systematic Literature Review

AHMADOU WAGNE, TU Wien, Austria

THOMAS ELMAR KOLB, TU Wien, Austria

ASHMI BANERJEE, Technical University of Munich, Germany

FATEMEH NAZARY, Polytechnic University of Bari, Italy


JULIA NEIDHARDT, TU Wien, Austria

YASHAR DELDJOO, Polytechnic University of Bari, Italy

Link: <https://recsys-lab.at/gen-conv-recsys-tutorial>



Knowledge and Data Foundation



Dialogue Corpora & Logs

How are conversational data sets and logs consumed by generative models?

Input:

36 of 49 selected papers used a dialogue corpora or logs dataset:

- HH Conv (24)
- user-system log (12)

Output:

15 papers used LLMs for generating and/or enhancing the dialogue corpora / logs data.

Dialogue Corpora & Logs

How are conversational data sets and logs consumed by generative models?

Input:

36 of 49 selected papers used a dialogue corpora or logs dataset:

- HH Conv (24)
- user-system log (12)

Human Human Conversation: ReDial, etc.

User-System Log: Amazon Review Dataset, etc.

Output:

15 papers used LLMs for generating and/or enhancing the dialogue corpora / logs data.

Examples

Large Language Models Know Your Contextual Search Intent: A Prompting Framework for Conversational Search (Mao et al., Findings of EMNLP 2023)

Input: conversational search benchmarks, including CAsT-19-21; (=real multi-turn conversational contexts)

Output: (Generated (with the help of LLM(s) e.g. GPT) “hypothetical responses” serve as synthetic conversational snippets augmenting user intent, though not full dialogues.)




Figure 1: An overview of LLM4CS.

Examples

Large Language Models Know Your Contextual Search Intent: A Prompting Framework for Conversational Search (Mao et al., Findings of EMNLP 2023)

Input: conversational search benchmarks, including CAsT-19-21; (=real multi-turn conversational contexts)

The prompt follows the formulation of **[Instruction, Demonstrations, Input]**, where Input is composed of the query q_t and the conversation context C_t of the current turn t .




Figure 1: An overview of LLM4CS.

Examples

Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models (Wang et al., EMNLP 2023)

iEvaLM with LLM-based user simulators driven by personas/behavior rules; demonstrated via Recall and Persuasiveness metrics.

Input: ReDial, OpenDialKG

Output: Synthetic conversational logs for simulation experiments

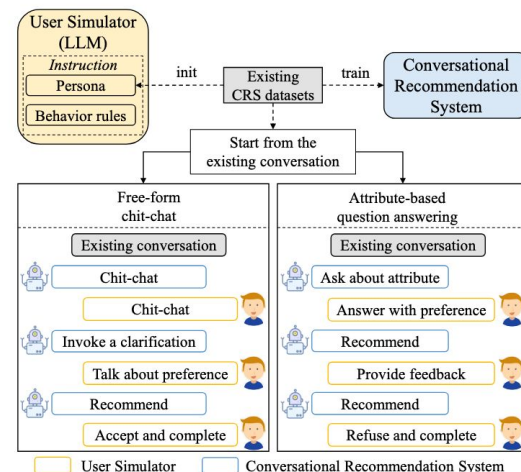


Figure 3: Our evaluation approach **iEvaLM**. It is based on existing CRS datasets and has two settings: free-form chit-chat (left) and attribute-based question answering (right).

Examples

Leveraging Large Language Models in Conversational Recommender Systems (Friedman et al., Google Research 2023)

An LLM user simulator interacts with the CRS to produce full sessions; controllable via session- or turn-level variables.

RecLLM, a large-scale CRS for YouTube videos built on LaMDA




Figure 1: Overview of key contributions from ReLLM. (1) A dialogue management module uses an LLM to converse with the user, track context and make system calls such as submitting a request to a recommendation engine all as a unified language modeling task. (2) Various solutions are presented for tractable retrieval over a large item corpus within an LLM-based CRS. (3) A ranker module uses an LLM to match preferences extracted from the context of the conversation to item metadata and generate a slate of recommendations that is displayed to the user. The LLM also jointly generates explanations for its decisions that can be surfaced to the user. (4) Interpretable natural language user profiles are consumed by system LLMs to modulate session-level context and increase personalization. (5) A controllable LLM-based user simulator can be plugged into the CRS to generate synthetic conversations for tuning system modules.

Input: highlight the lack of logs or observational dialogue corpora as a central challenge

Output: generate dialogue sessions, user feedback, item summaries, and recommendation utterances with llms

Examples

Leveraging Large Language Models in Conversational Recommender Systems (Friedman et al., Google Research 2023)

*“In this paper we assume a simplified setting where users interact with the system **only through conversation**. We would like to **generalize our system to handle more realistic scenarios where users give feedback through other channels as well such as clicking on items or like buttons**. We would also like to consider more complicated recommender system UIs containing hierarchical structures such as item shelves as opposed to just flat slates.”* (Friedman et al., Google Research 2023)




Figure 1: Overview of key contributions from RecLLM. (1) A dialogue management module uses an LLM to converse with the user, track context and make system calls such as submitting a request to a recommendation engine all as a unified language modeling task. (2) Various solutions are presented for tractable retrieval over a large item corpus within an LLM-based CRS. (3) A ranker module uses an LLM to match preferences extracted from the context of the conversation to item metadata and generate a slate of recommendations that is displayed to the user. The LLM also jointly generates explanations for its decisions that can be surfaced to the user. (4) Interpretable natural language user profiles are consumed by system LLMs to modulate session-level context and increase personalization. (5) A controllable LLM-based user simulator can be plugged into the CRS to generate synthetic conversations for tuning system modules.

gue corpora as a central challenge
m summaries, and



Thomas E. Kolb

Simulation

Agenda

- Introduction
- Core Systems & Components
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- **Simulation**
- Evaluation
- Open Challenges & Future Directions

Key Questions

- What is the goal of (generative) simulation?
- Which aspects simulation do we have?
- How can generative approaches (e.g. LLMs) enable these simulations?
- Which new datasets are in the field?

Why Simulation? Why for Conversation?




But ...

Pre GenAI: *Is there any fitting conv. dataset for evaluation?
Should I conduct an annotation study?*



With Simulation:



UserSimAgent

CRSAgent

Key Wh-questions to Model Simulations

Why simulation?

(Why)

- Data Sparsity and Availability
- Privacy
- Cost
- [...]

What is simulated?

(What)

- User Data
- Item Data
- Dialogue Data

Level of simulation?

(Which Level)


- Empirical Data
- Data Sampling
- Data Synthesis
- Full Simulation

LLM involvement?

(How)

- ... as extractor
- ... as summariser/ generator
- ... for enrichment
- ... as a metatool
- [...]

Towards Full Conversational Simulation




*Less Control & Bias,
More Authentic/Real*



*More Control & Bias,
Less Authentic/Real*

Towards Full Conversational Simulation



Purely observational data used, classical pre generative approaches.

*Less Control & Bias,
More Authentic/Real*



*More Control & Bias,
Less Authentic/Real*

Key Characteristics of a Conv. Dataset

A **conversational dataset** that is composed of:

- Preference Data (= Task/Goal Orientation)
- Multi-turn Conversation
- Item Data

Example

# conversations	10006	HUMAN:	hello, can i help you find a movie ?
# utterances	182150	HRED:	hi what kind of movies do you like
# users	956	OURS:	hi, how can i help you ?
# movie mentions	51699	SEEKER:	hello ! i would like to find a suspenseful, but clean family friendly movie .
seeker mentioned	16278	...	...
recommender suggested	35421	SEEKER:	2001 : a space odyssey might be a great option. anything else that you would suggest ?
not seen	16516	HUMAN:	you can try planet of the apes the older one is quite suspenseful and family friendly .
seen	31694	HRED:	i haven't seen that one yet but i heard it was good. i haven't seen that one. have you seen the last house on the left ?
did not say	3489	OURS:	star wars : the force awakens is also a good one return of the jedi all good movies
disliked (4.9%)	2556	SEEKER:	those sound good ! i 'm going to look into those movies.
liked (81%)	41998	HUMAN:	i hope you enjoy, have a nice one
did not say (14%)	7145	HRED:	have you seen foxcatcher ? it 's about a man who has a rich guy.
		OURS:	i hope i was able to help you find a good movie to watch
		SEEKER:	thank you for your help ! have a great night ! good bye

Table 1a. (Above) REDIAL data statistics. For the movie dialogue forms, the numbers shown represent the seeker's answers.


Table 1b.(Right) Conversation excerpts (HUMAN followed by response by SEEKER) and model outputs (OUR proposed approach compared to HRED a generic dialogue model [2]).

ReDial (2018)

- Crowd-Sourcing via Amazon Mechanical Turk
- Movie Seeker & Recommender as Roles

One of the early datasets highly used in conversational recommender systems research.

Towards Full Conversational Simulation



A data-generation process that ***fabricates entire dialogue transcripts*** so they statistically resemble human–agent conversations, but ***does not let the synthetic user react to the system in real time.***

Less Control & Bias,
More Authentic/Real



More Control & Bias,
Less Authentic/Real

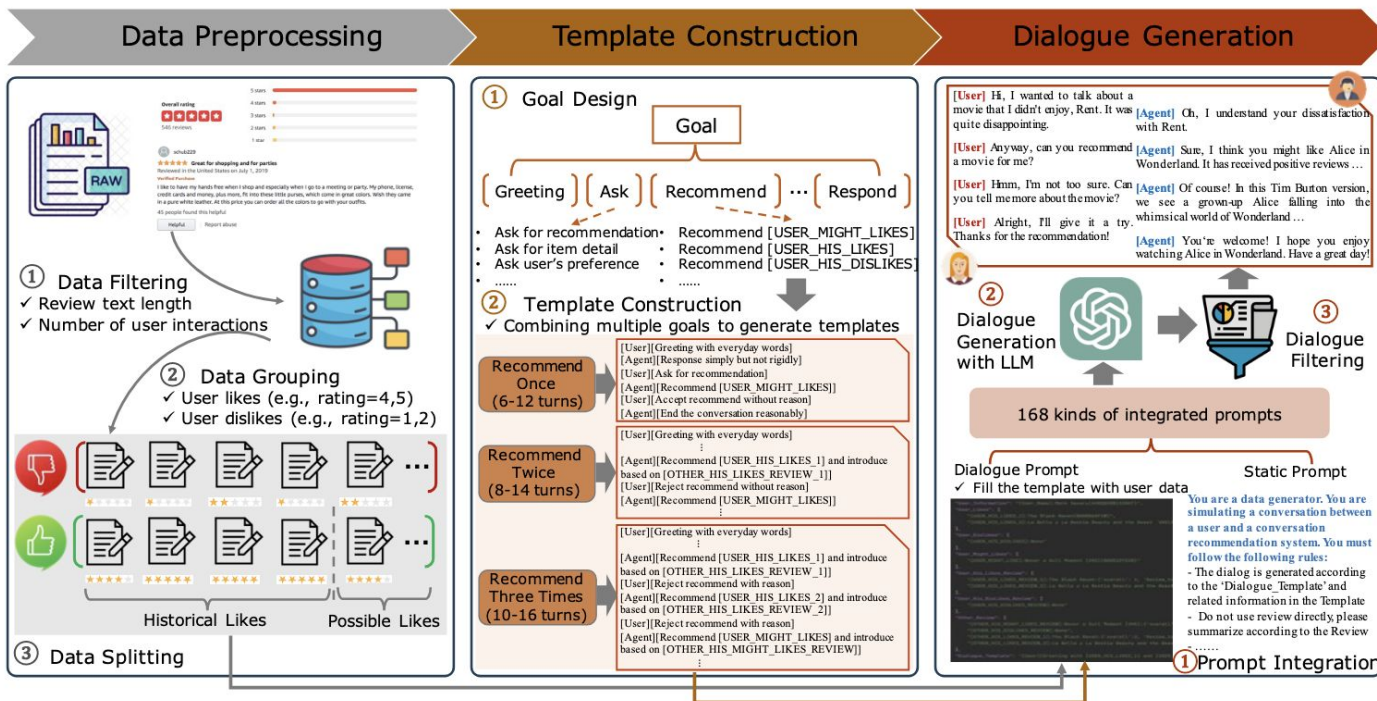
Synthesized, Synthetic, Simulation?

To overcome the limitations of classical, pre-generative simulation methods by producing realistic data and behaviors.

New Approaches:

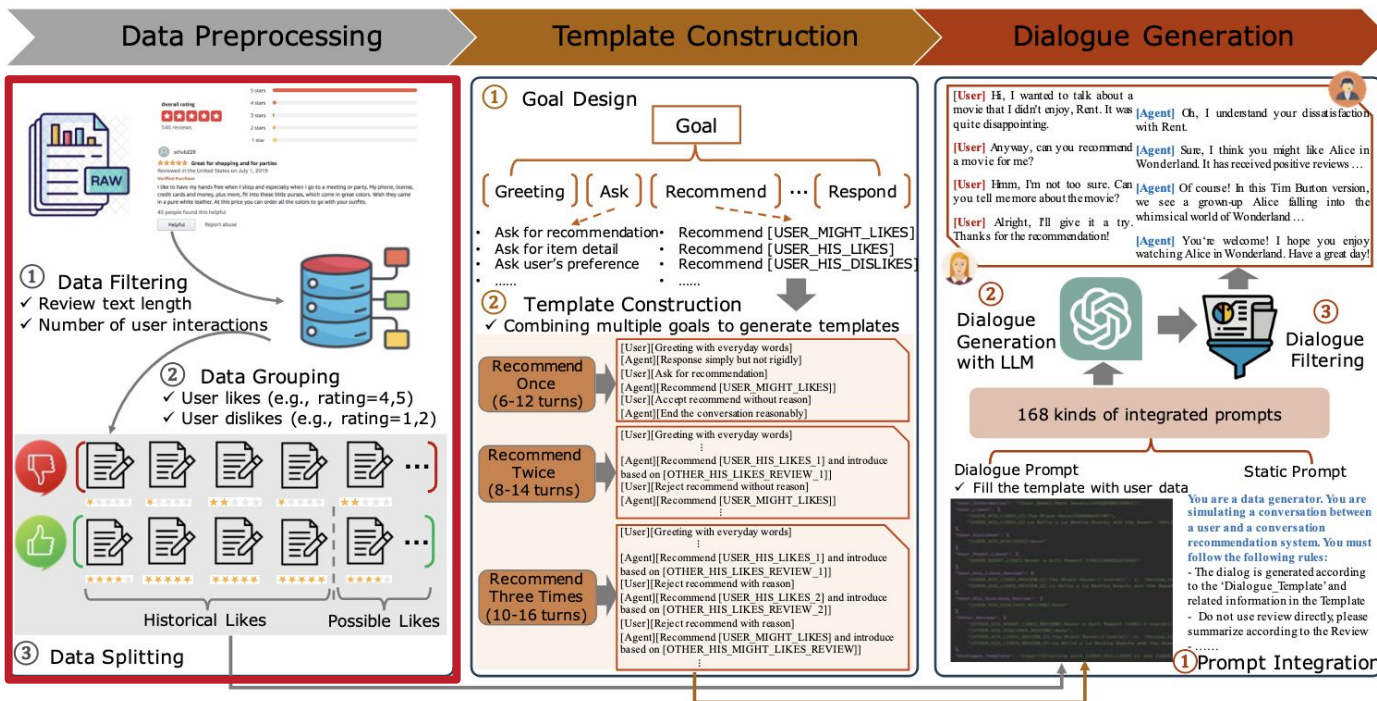
- **Data synthesized from real sources:** combine existing signals e.g. user preferences, item metadata, dialogue acts into richer datasets.
- **Model generated** data: create artificial dialogues and interactions.
- User & Conversational **Simulation**: simulate users and multi-turn conversations for training and evaluation.

Example

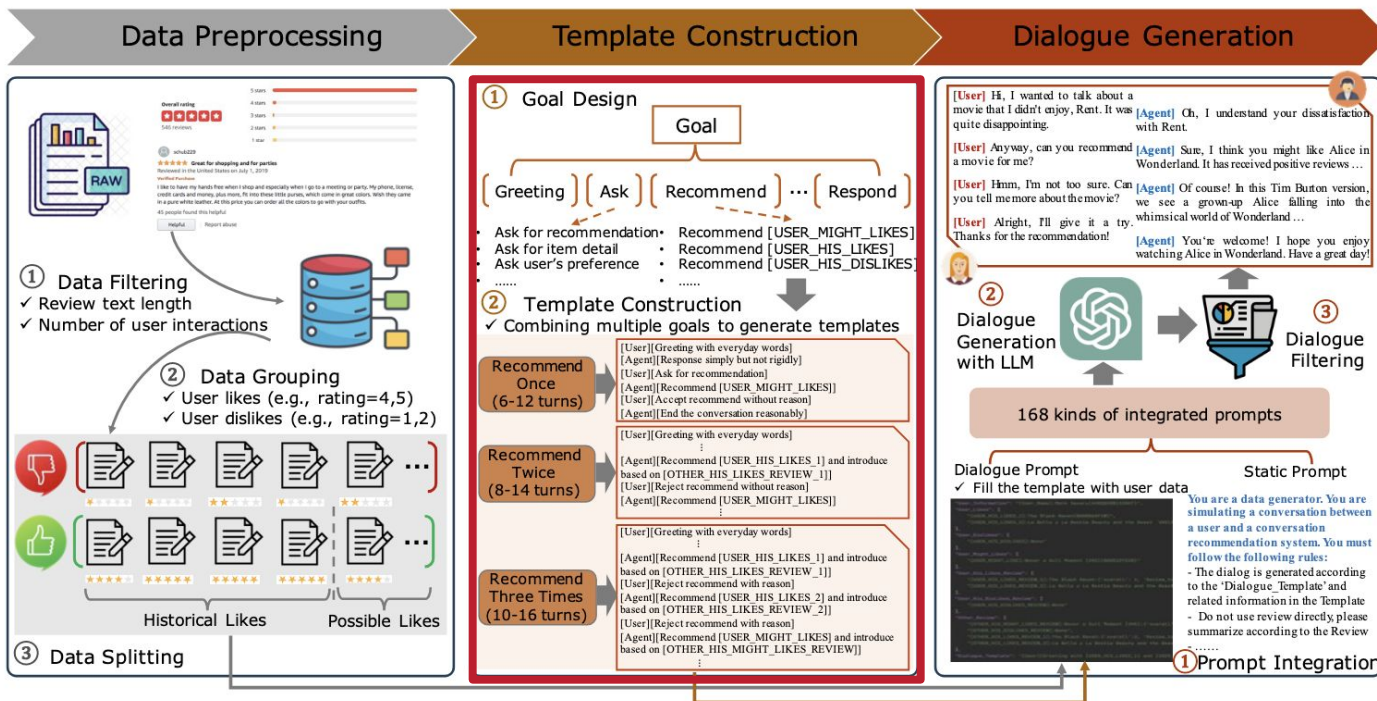


Example

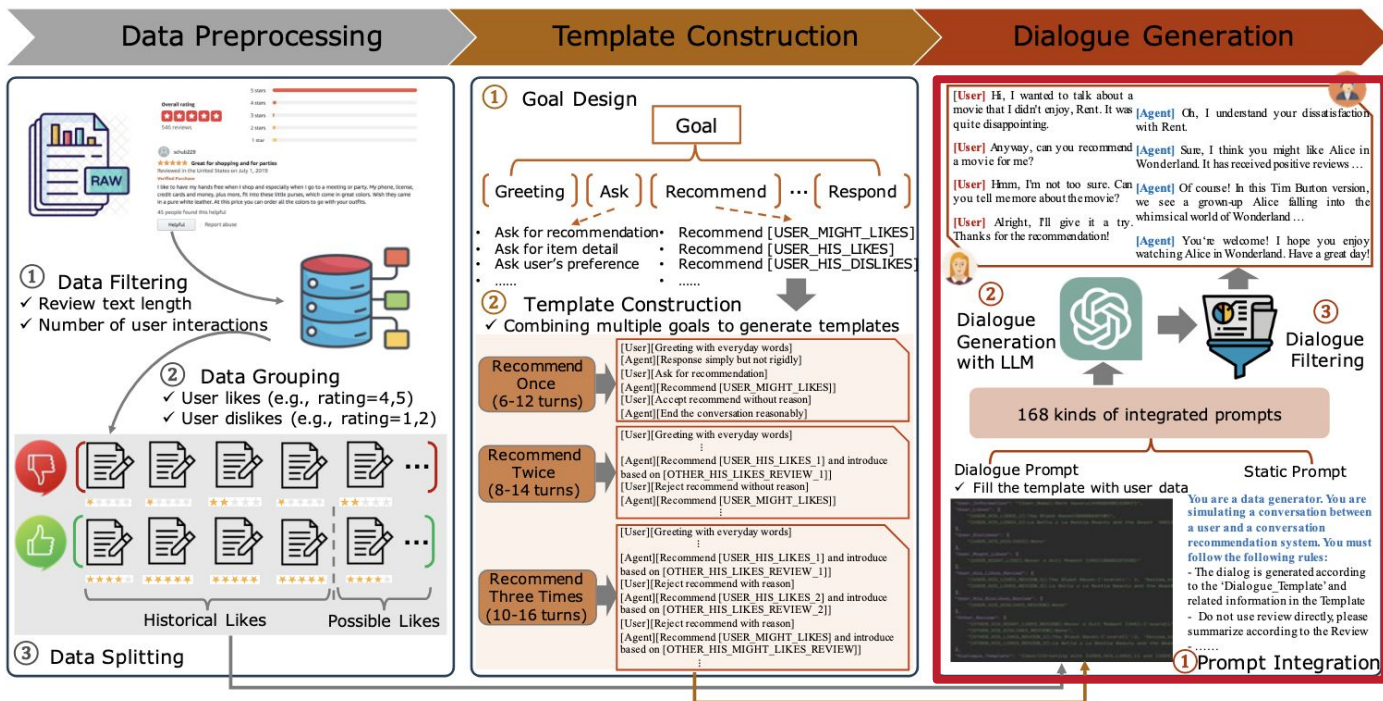
Input:
Amazon
Review
Dataset




Example



Example



Towards Full Conversational Simulation



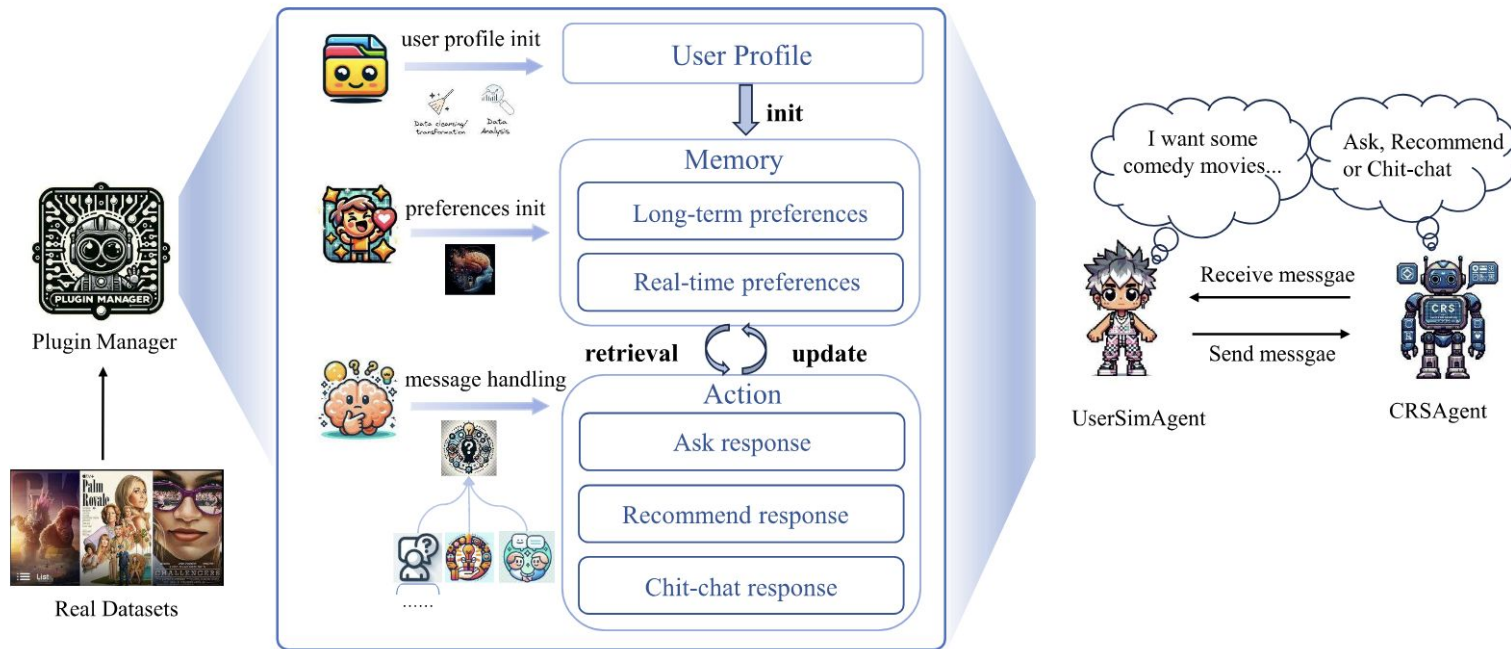
*A reactive model that **plays the role of a user** in a live conversation, producing utterances, clarifications, and preference signals in response to each system move, thereby enabling **turn-by-turn evaluation or policy learning**.*

*Less Control & Bias,
More Authentic/Real*

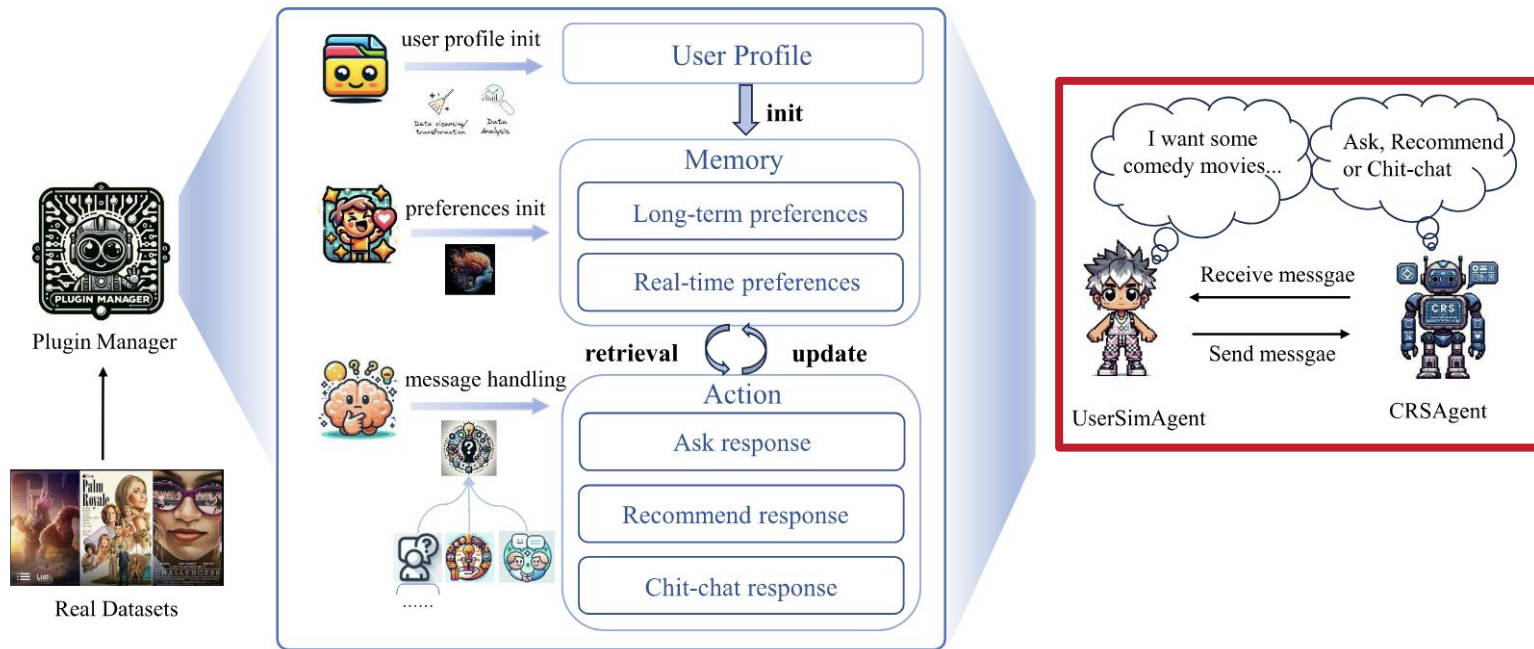


*More Control & Bias,
Less Authentic/Real*

Example



Example




Towards Full Conversational Simulation

Observational
Data

Synthetic
Data

User
Simulation

Full Conversational
Simulation




A complete **closed-loop environment** in which both the **user(s)** and the **external world** are **modeled**, allowing the recommender to act, learn, and observe emergent behaviour over many simulated dialogues—**analogous to running the entire ecosystem in silico**.

Less Control & Bias,
More Authentic/Real




More Control & Bias,
Less Authentic/Real

Example



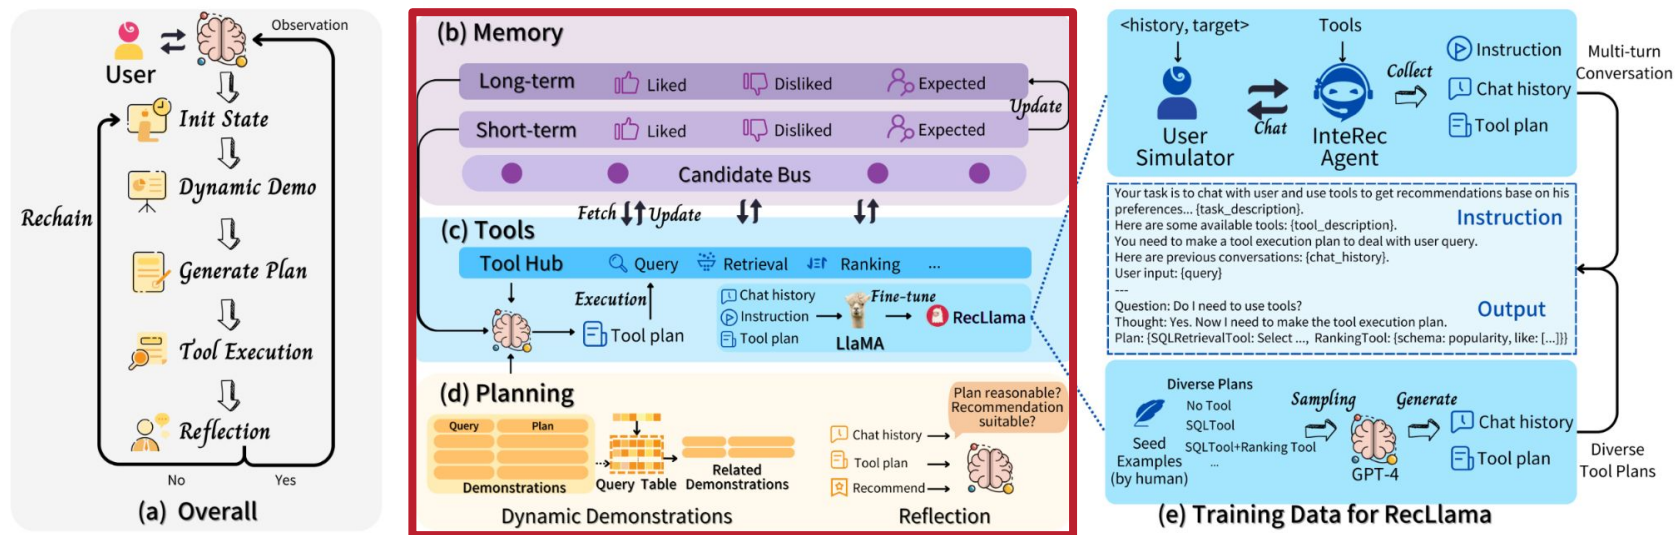
Close to full conversational conversation but still missing modeling of the “external world”

Example




Close to full conversational conversation but still missing modeling of the “external world”

Example



Close to full conversational conversation but still missing modeling of the “external world”


Example



Close to full conversational conversation but still missing modeling of the “external world”

From Classical to Generative Datasets

Classical Datasets



(Towards) Generative Datasets

ReDial (Li et al., 2018)

OpenDialKG (Moon et al., 2019)

GoRecDial (Kang et al., 2019)

INSPIRED (Hayati et al., 2020)

INSPIRED2 (Manzoor et al., 2022)

DuRecDial (Liu et al., 2020)

DuRecDial 2.0 (Liu et al., 2021)

U-NEED (Liu et al., 2023)

TG-ReDial (Zhou et al., 2020) - **synthetic**

Synthetically self generated data (Friedman et al., 2023)

LLM-Redial (Liang et al., 2024)

PEARL (Kim et al., 2024)

N-of-1 (Synthetic) Dataset (Yang et al., 2024)


Synthesized Tracking & Recommendation Dataset (Ashby et al., 2024)

DistillRecDial (Martina et al., 2025; *yesterday @ RecSys*



From Classical to Generative Datasets (cont.)

Classical Datasets



(Towards) Generative Datasets

*“To overcome conversational data limitations in the **absence of an existing production CRS**, we propose techniques for building a controllable LLM-based **user simulator** to **generate synthetic conversations**.” (Friedman et al., 2023 - Google Research)*



TG-ReDial (Zhou et al., 2020) - **synthetic**

Synthetically self generated data (Friedman et al., 2023)

LLM-Redial (Liang et al., 2024)

PEARL (Kim et al., 2024)


N-of-1 (Synthetic) Dataset (Yang et al., 2024)

Synthesized Tracking & Recommendation Dataset
(Ashby et al., 2024)

DistillRecDial (Martina et al., 2025; *yesterday @ RecSys*
😊)

From Classical to Generative Datasets (cont.)

Classical Datasets



(Towards) Generative Datasets

No public dataset available



TG-ReDial (Zhou et al., 2020) - **synthetic**

Synthetically self generated data (Friedman et al., 2023)

LLM-Redial (Liang et al., 2024)

PEARL (Kim et al., 2024)



N-of-1 (Synthetic) Dataset (Yang et al., 2024)


Synthesized Tracking & Recommendation Dataset
(Ashby et al., 2024)

DistillRecDial (Martina et al., 2025; yesterday @ RecSys



From Classical to Generative Datasets (cont.)

Classical Datasets



(Towards) Generative Datasets

Dataset available ✓

TG-ReDial (Zhou et al., 2020) - **synthetic**

Synthetically self generated data (Friedman et al., 2023)

LLM-Redial (Liang et al., 2024)

PEARL (Kim et al., 2024)

N-of-1 (Synthetic) Dataset (Yang et al., 2024)

Synthesized Tracking & Recommendation Dataset
(Ashby et al., 2024)

DistillRecDial (Martina et al., 2025; *yesterday @ RecSys*



TG-ReDial

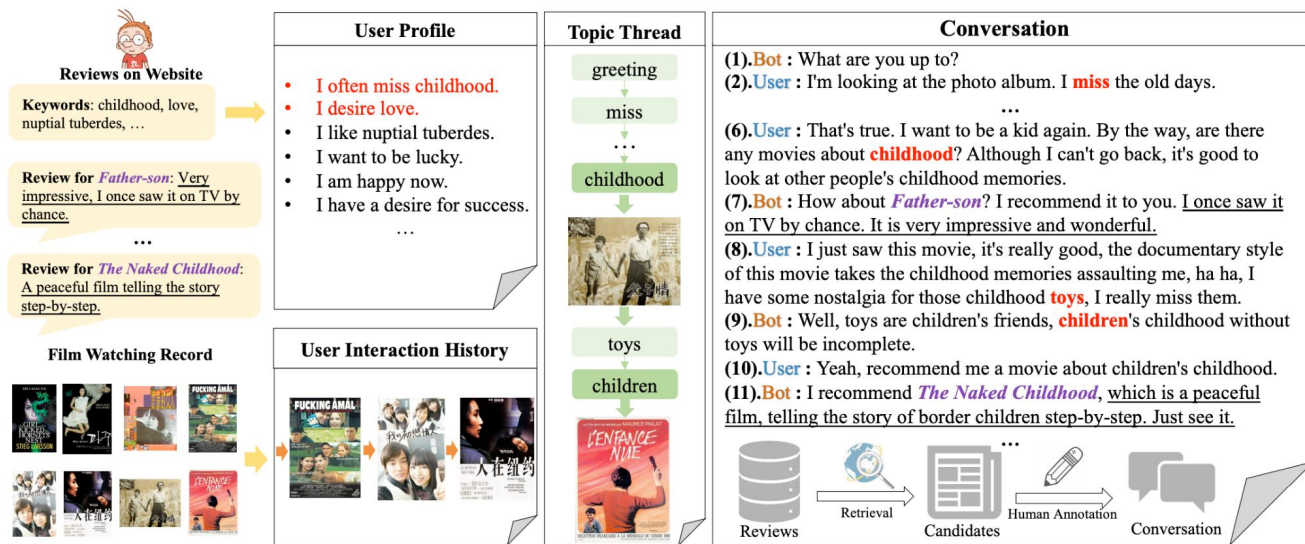
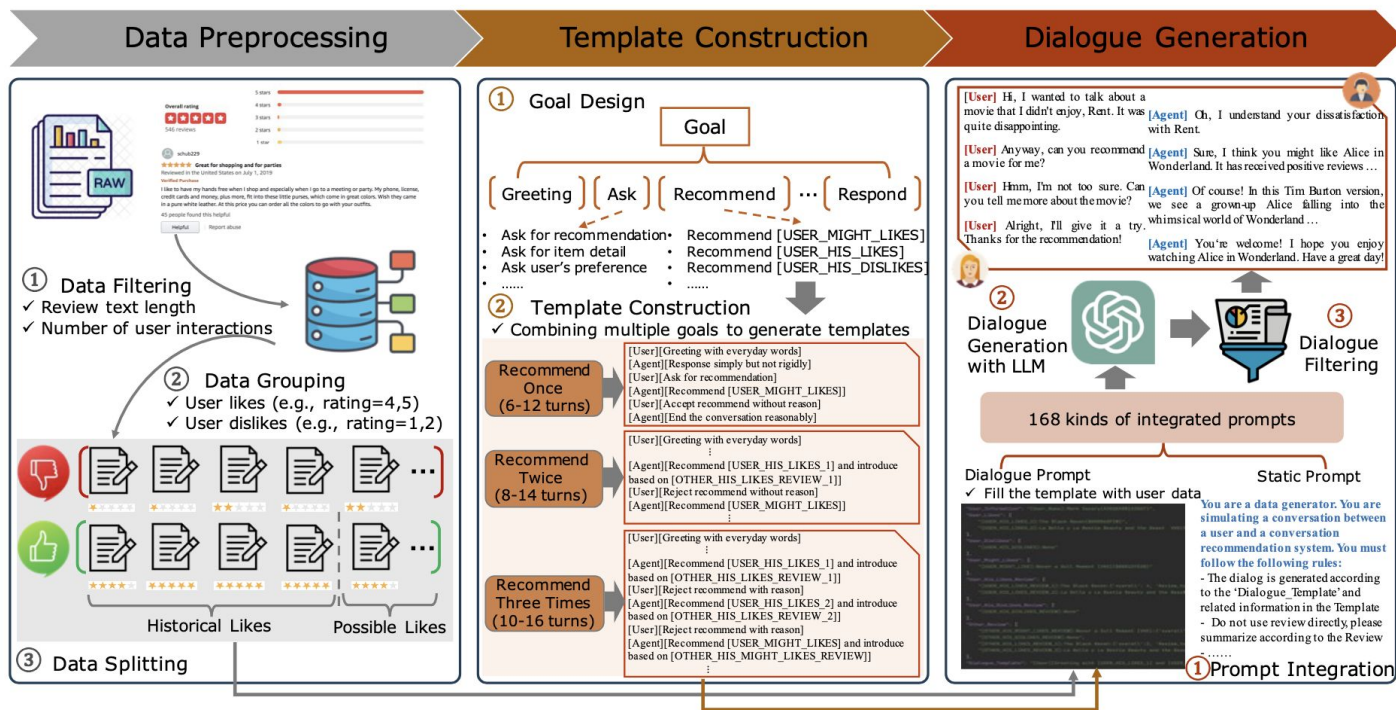
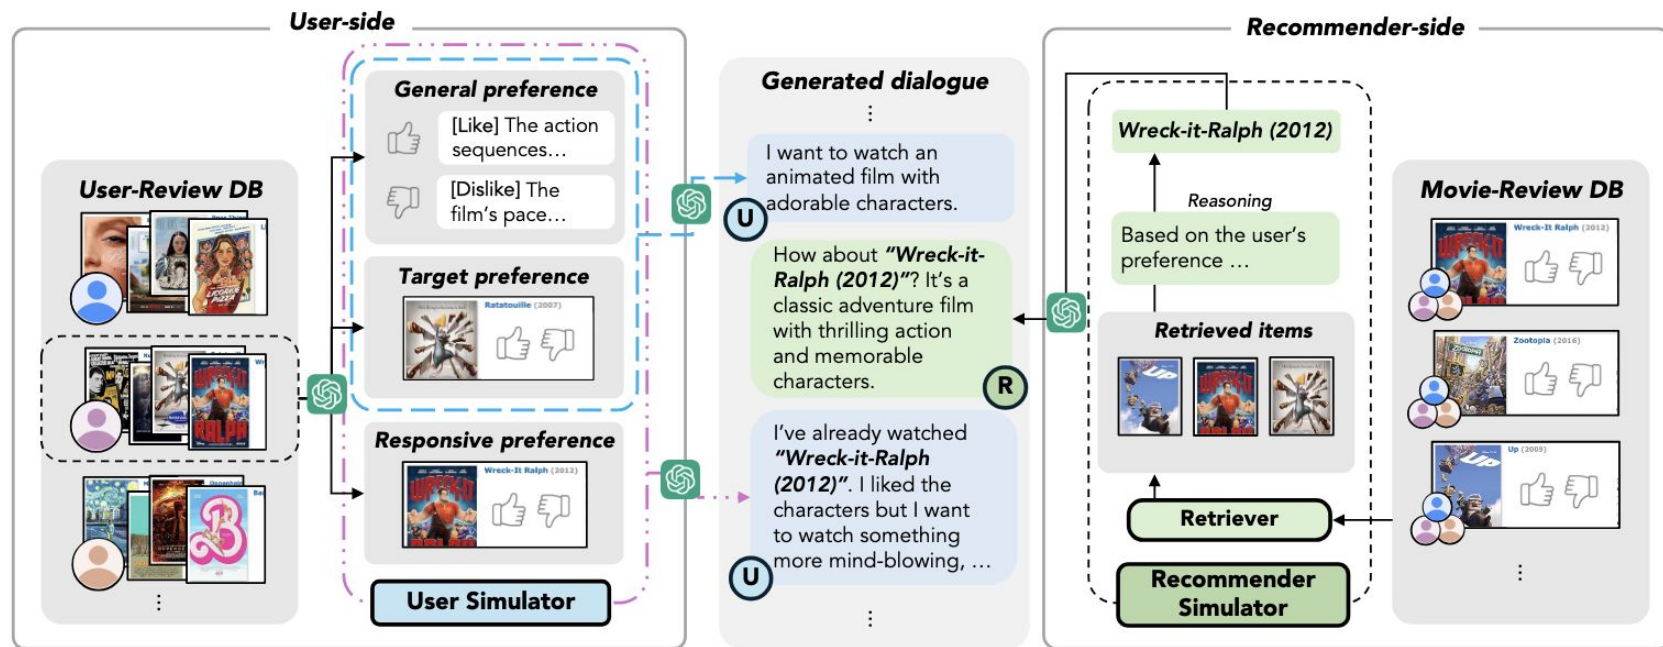


Figure 1: An illustrative example for TG-ReDial dataset. We utilize real data to construct the recommended movies, topic threads, user profiles and utterances. Other user-related information (e.g., historical interaction records) is also available in our dataset.

LLM-Redial



PEARL



Comparison (ReDial vs LLM-ReDial)

ReDial 2018 (Li et al.)



Real human–human dialogue:

Roles: “**seeker**” ↔ “**recommender**”

10k dialogues / **182.1k** utterances,
single-domain (movies)



LLM-ReDial 2024 (Liang et al.)

Conversations are **generated by LLMs (GPT-3.5-turbo)**

Roles: **turn goal templates**

Amazon review logs + user histories, guided by
turn-goal templates to simulate the
recommendation flow.

47.6k dialogues / **482.6k** utterances across **4**
domains (Books, Movies, Sports, Electronics)

Redial (Li et al., 2018): Li, R., Kahou, S., Schulz, H., Michalski, V., Charlin, L., & Pal, C. (2018). Towards Deep Conversational Recommendations. In Advances in Neural Information Processing Systems 31 (NIPS 2018).

LLM-REDIAL (Liang et al., 2024): Liang, T., Jin, C., Wang, L., Fan, W., Xia, C., Chen, K., & Yin, Y. (2024). LLM-REDIAL: A Large-Scale Dataset for Conversational Recommender Systems Created from User Behaviors with LLMs. In Findings of the Association for Computational Linguistics: ACL 2024 (pp. 8926–8939). Association for Computational Linguistics.

Comparison (ReDial vs LLM-ReDial)

ReDial₂

Real human–h

Roles: “seeker

10k dialogues ,
single-domain

User ID: A1EXXXXXUE6B0
Historical Interactions: ["Robin Williams: Live On Broadway", "Mission Impossible on VHS", "Solaris", "Elysium", "Wall Street", "Mystic River", ...]
Conversation 1:
[User] Hi, I recently watched Mission Impossible on VHS and it was a fantastic high-tech spy movie! Tom Cruise gave ...
[Agent] That's great to hear! I have a movie recommendation for you based on your likes. How about checking out World War Z?
[User] Oh, I'm definitely interested. Can you tell me more about the movie?
[Agent] Certainly! World War Z is a good zombie war flick...
Conversation 2:
[User] Hi, I recently watched the movie "Solaris" and I have to say, I didn't enjoy it at all. It felt overly melodramatic and lacked substance...
[Agent] I understand why you didn't like "Solaris". I can recommend a movie you might enjoy instead. How about "Elysium"? It's a bilingual film that...
[User] Oh, I've actually already watched "Elysium" and it was better than I expected, but still not great....
[Agent] How about giving "Wrecking Crew" a try? It's another movie you might like based on your previous preferences...
[User] Sure, that sounds interesting...

Conversation in Redial
[User] Hi I am looking for a movie like Super Troopers (2001)
[Agent] You should watch Police Academy (1984)
[User] Is that a great one? I have never seen it. I have seen American Pie I mean American Pie (1999)
[Agent] Yes Police Academy (1984) is very funny and so is Police Academy 2: Their First Assignment (1985)
[User] It sounds like I need to check them out'
[Agent] yes you will enjoy them
[User] I appreciate your time. I will need to check those out. Are there any others you would recommend?
[Agent] yes Lethal Weapon (1987)
[User] Thank you i will watch that too
[Agent] and also Beverly Hills Cop (1984)
[User] Thanks for the suggestions.
[Agent] you are welcome and also 48 Hrs. (1982)

l.)
GPT-3.5-turbo)

guided by

is 4
nics)

hen, K., & Yin, Y. (2024).
Created from User Behaviors
(pp. 8926–8939). Association

Redial (Li et al., 2018): Li, R., K:
Deep Conversational Recommender
(NIPS 2018).

From Classical to Generative Datasets (cont.)

- Redial (Li et al., 2018):** Li, R., Kahou, S., Schulz, H., Michalski, V., Charlin, L., & Pal, C. (2018). Towards Deep Conversational Recommendations. In Advances in Neural Information Processing Systems 31 (NIPS 2018).
- OpenDialogKG (Moon et al., 2019):** Moon, S., Shah, P., Kumar, A., & Subba, R. (2019). OpenDialogKG: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 845–854). Association for Computational Linguistics.
- GoRecDial (Kang et al., 2019):** Kang, D., Balakrishnan, A., Shah, P., Crook, P., Boureau, Y.L., & Weston, J. (2019). Recommendation as a Communication Game: Self-Supervised Bot-Play for Goal-oriented Dialogue. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP) (pp. 1951–1961). Association for Computational Linguistics.
- INSPIRED (Hayati et al., 2020):** Hayati, S., Kang, D., Zhu, Q., Shi, W., & Yu, Z. (2020). INSPIRED: Toward Sociable Recommendation Dialog Systems. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP) (pp. 8142–8152). Association for Computational Linguistics.
- DuRecDial (Liu et al., 2020):** Liu, Z., Wang, H., Niu, Z.Y., Wu, H., Che, W., & Liu, T. (2020). Towards Conversational Recommendation over Multi-Type Dialogs. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (pp. 1036–1049). Association for Computational Linguistics.
- DuRecDial2.0 (Liu et al., 2021):** Liu, Z., Wang, H., Niu, Z.Y., Wu, H., & Che, W. (2021). DuRecDial 2.0: A Bilingual Parallel Corpus for Conversational Recommendation. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (pp. 4335–4347). Association for Computational Linguistics.
- INSPIRED2 (Manzoor et al., 2022):** Ahtsham Manzoor, & Dietmar Jannach (2022). INSPIRED2: An Improved Dataset for Sociable Conversational Recommendation. In Proceedings of the Fourth Knowledge-aware and Conversational Recommender Systems Workshop (KaRS 2022), co-located with the 16th ACM Conference on Recommender Systems (RecSys 2022) (pp. 73–80). CEUR-WS.org.
- U-NEED (Liu et al., 2023):** Liu, Y., Zhang, W., Dong, B., Fan, Y., Wang, H., Feng, F., Chen, Y., Zhuang, Z., Cui, H., Li, Y., & Che, W. (2023). U-NEED: A Fine-grained Dataset for User Needs-Centric E-commerce Conversational Recommendation. In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (pp. 2723–2732). Association for Computing Machinery.
- TG-REDIAL (Zhou et al., 2020):** Zhou, K., Zhou, Y., Zhao, W., Wang, X., & Wen, J.R. (2020). Towards Topic-Guided Conversational Recommender System. In Proceedings of the 28th International Conference on Computational Linguistics (pp. 4128–4139). International Committee on Computational Linguistics.
- Synthetically self generated data (Friedman et al., 2023):** Luke Friedman, Sameer Ahuja, David Allen, Zhenning Tan, Hakim Sidahmed, Changbo Long, Jun Xie, Gabriel Schubiner, Ajay Patel, Harsh Lara, Brian Chu, Zexi Chen, & Manoj Tiwari. (2023). Leveraging Large Language Models in Conversational Recommender Systems.
- LLM-REDIAL (Liang et al., 2024):** Liang, T., Jin, C., Wang, L., Fan, W., Xia, C., Chen, K., & Yin, Y. (2024). LLM-REDIAL: A Large-Scale Dataset for Conversational Recommender Systems Created from User Behaviors with LLMs. In Findings of the Association for Computational Linguistics: ACL 2024 (pp. 8926–8939). Association for Computational Linguistics.
- PEARL (Kim et al., 2024):** Kim, M., Kim, M., Kim, H., Kwak, B.w., Kang, S., Yu, Y., Yeo, J., & Lee, D. (2024). Pearl: A Review-driven Persona-Knowledge Grounded Conversational Recommendation Dataset. In Findings of the Association for Computational Linguistics: ACL 2024 (pp. 1105–1120). Association for Computational Linguistics.
- N-of-1 (Synthetic) Dataset (Yang et al., 2024):** Zhongqi Yang, Elahe Khatibi, Nitish Nagesh, Mahyar Abbasian, Iman Azimi, Ramesh Jain, & Amir M. Rahmani (2024). ChatDiet: Empowering personalized nutrition-oriented food recommender chatbots through an LLM-augmented framework. Smart Health, 32, 100465.
- Synthesized Tracking Dataset, Synthesized Recommendation Dataset (Ashby et al., 2024):** Ashby, T., Kulkarni, A., Qi, J., Liu, M., Cho, E., Kumar, V., & Huang, L. (2024). Towards Effective Long Conversation Generation with Dynamic Topic Tracking and Recommendation. In Proceedings of the 17th International Natural Language Generation Conference (pp. 540–556). Association for Computational Linguistics.
- DistillRecDial (Martina et al., 2025):** Martina, A., Petruzzelli, A., Musto, C., Gemmis, M., Lops, P., & Semeraro, G. (2025). DistillRecDial: A Knowledge-Distilled Dataset Capturing User Diversity in Conversational Recommendation. In Proceedings of the Nineteenth ACM Conference on Recommender Systems (pp. 726–735). Association for Computing Machinery.

Towards Full Conversational Simulation

- Lu, Y., Bao, J., Ma, Z., Han, X., Wu, Y., Cui, S., & He, X. (2023). **AUGUST: an Automatic Generation Understudy for Synthesizing Conversational Recommendation Datasets**. In *Findings of the Association for Computational Linguistics: ACL 2023* (pp. 10538–10549). Association for Computational Linguistics.
- Mao, K., Dou, Z., Mo, F., Hou, J., Chen, H., & Qian, H. (2023). **Large Language Models Know Your Contextual Search Intent: A Prompting Framework for Conversational Search**. In *Findings of the Association for Computational Linguistics: EMNLP 2023* (pp. 1211–1225). Association for Computational Linguistics.
- Luke Friedman, Sameer Ahuja, David Allen, Zhenning Tan, Hakim Sidahmed, Changbo Long, Jun Xie, Gabriel Schubiner, Ajay Patel, Harsh Lara, Brian Chu, Zexi Chen, & Manoj Tiwari. (2023). **Leveraging Large Language Models in Conversational Recommender Systems**.
- Wang, X., Tang, X., Zhao, X., Wang, J., & Wen, J.R. (2023). **Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* (pp. 10052–10065). Association for Computational Linguistics.
- Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Chris Leung, Jiajie Tang, and Jiebo Luo. 2024. **LLM-Rec: Personalized Recommendation via Prompting Large Language Models**. In Findings of the Association for Computational Linguistics: NAACL 2024, pages 583–612, Mexico City, Mexico. Association for Computational Linguistics.
- Ashby, T., Kulkarni, A., Qi, J., Liu, M., Cho, E., Kumar, V., & Huang, L. (2024). **Towards Effective Long Conversation Generation with Dynamic Topic Tracking and Recommendation**. In *Proceedings of the 17th International Natural Language Generation Conference* (pp. 540–556). Association for Computational Linguistics.
- Yoon, S.e., He, Z., Echterhoff, J., & McAuley, J. (2024). **Evaluating Large Language Models as Generative User Simulators for Conversational Recommendation**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)* (pp. 1490–1504). Association for Computational Linguistics.
- Liang, T., Jin, C., Wang, L., Fan, W., Xia, C., Chen, K., & Yin, Y. (2024). **LLM-REDIAL: A Large-Scale Dataset for Conversational Recommender Systems Created from User Behaviors with LLMs**. In *Findings of the Association for Computational Linguistics: ACL 2024* (pp. 8926–8939). Association for Computational Linguistics.
- Zhu, L., Huang, X., & Sang, J. (2024). **How Reliable is Your Simulator? Analysis on the Limitations of Current LLM-based User Simulators for Conversational Recommendation**. In *Companion Proceedings of the ACM Web Conference 2024* (pp. 1726–1732). Association for Computing Machinery.
- Zhu, L., Huang, X., & Sang, J. (2025). **A LLM-based Controllable, Scalable, Human-Involved User Simulator Framework for Conversational Recommender Systems**. In *Proceedings of the ACM on Web Conference 2025* (pp. 4653–4661). Association for Computing Machinery.
- Mao, W., Wu, J., Chen, W., Gao, C., Wang, X., & He, X. (2025). **Reinforced Prompt Personalization for Recommendation with Large Language Models**. *ACM Trans. Inf. Syst.*, 43(3).



Ashmi Banerjee

Evaluation

Agenda

- Introduction
- Core Systems & Components
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- Simulation
- **Evaluation**
- Open Challenges & Future Directions

Key Dimensions

Output Type (What)	User-facing		Non-User-facing / hidden	
Quality Dimensions (Which Basis)	Task Effectiveness	Dialog / Task Efficiency		Conversational Quality
	Subtask Performance	Trust & User satisfaction		Ethical Considerations & System Integrity
Paradigms (How)	Offline		Online	User Studies & Lab Experiments
Evaluator (Who)	Automated Metrics (e.g., NDCG, HitRate, Recall etc.,)		Humans (e.g. Fluency, Informativeness etc.,)	LLM-as-a-Judge
Stakeholders (For Whom)	Consumers		Item-Providers	Others

Overview (What to Evaluate)

Output Type


(What)

User-facing

Non-User-facing / hidden

Directly consumed by the user

- Natural language conversation turns
- Recommended items
- Explanations and justifications
- Multimedia elements (e.g., images)



Overview (What to Evaluate)

Output Type

(What)

User-facing

Non-User-facing / hidden

Internal or intermediate model steps

- Learned embeddings
- Augmented data / Retrieved item sets
- Hidden reasoning paths
- Internal dialogue states



Evaluating Non-User-Facing Outputs

Output Type

(What)

User-facing

Non-User-facing / hidden

- Current research **predominantly evaluates user-facing outputs** (discussed later).
- Non-user-facing outputs are typically **assessed indirectly** to prove their contribution.
 - Primary Method: **Ablation Studies**
 - Example (MemoCRS): Removing components like user-specific memory, collaborative knowledge, and reasoning guidelines to demonstrate their impact on final performance [Li et al., 2024a].
 - Example (CoRE-CoG): Evaluating retriever components (trigger, classifier) with standard metrics (Recall, MRR, BLEU, F1) to validate their design [Wang et al., 2024b].
 - Some extend this with sensitivity analysis for a more comprehensive assessment [Wang et al., 2024b].

Summary (What to Evaluate)

Output Type

(What)

User-facing

Non-User-facing / hidden


- GenCRS aims for an enjoyable, human-like, and interactive experience.
 - This requires a fundamental shift in evaluation focus.
- Evaluating the "Journey" is as important as the "Destination"
 - Journey: Dialogue coherence, interaction efficiency, user enjoyment.
 - Destination: Relevance of the final recommended items.



Overview (Which Basis)

Quality Dimensions (Which Basis)	Task Effectiveness	Dialog / Task Efficiency	Conversational Quality
	Subtask Performance	Trust & User satisfaction	Ethical Considerations & System Integrity

- **Task Effectiveness:** *Did the user find a good item?*
- **Dialogue / Task Efficiency:** *How quickly and easily did they find it?*
- **Conversational Quality:** *Was the conversation natural and coherent?*
- **Subtask Performance:** *Did the internal modules work correctly?*
- **Trust & User Satisfaction:** *Did the user enjoy and trust the system?*
- **Ethical Considerations:** *Is the system fair, safe, and honest?*



Task Effectiveness

Quality Dimensions (Which Basis)	Task Effectiveness	Dialog / Task Efficiency	Conversational Quality
	Subtask Performance	Trust & User satisfaction	Ethical Considerations & System Integrity

Definition: The system's ability to help users successfully discover desired items.

Core Question: *"Did the user find a suitable item that effectively met their needs?"*


Common Metrics:

- Offline metrics: Hit Rate@K, Recall@K, NDCG@K
- Online metrics: User acceptance rate, click-through rate, purchases


Importance: This is the most fundamental and widely reported dimension (~80% of surveyed papers).

Examples:

- Recall@K is used to evaluate recommendation/retrieval effectiveness [he_2023, yang_2024, mao_2023, kim_2024, and others].
- Hit Rate@K is used to measure recommendation success [wang_2023, liu_2023, leszczynski_2023, xi_2024, and others].



Dialog / Task Efficiency



Definition: How efficiently the system guides the user to a successful recommendation.

Goal: Avoid user frustration from overly long or unproductive conversations.

Common Metrics:

- Number of dialogue turns ("turns to success")
- Time to complete task
- System latency and computational cost [kunstmann_2024].

Importance: This is the most fundamental and widely reported dimension (~80% of surveyed papers).

Examples:

- [wang_2021] uses a mix of automated metrics (Perplexity, BLEU, ROUGE, Distinct-n) and human evaluation to assess dialogue quality.
- [kunstmann_2024] evaluated their EventChat system in a 2-month real-world deployment, measuring user experience, latency, and cost.



Conversation Quality

Quality Dimensions (Which Basis)	Task Effectiveness	Dialog / Task Efficiency	Conversational Quality
	Subtask Performance	Trust & User satisfaction	Ethical Considerations & System Integrity

Definition: The linguistic and stylistic quality of the dialogue.

Challenge: No single "ground truth" exists for a perfect conversation, making human judgment crucial.

Objective Metrics:

- Fluency: Perplexity
- Content Overlap: BLEU, ROUGE
- Accuracy: Recall of target items in the response [chen_2024].

Subjective Metrics (Human Evaluation):

- Naturalness [leszczynski_2023]
- Persuasiveness & Engagement [kim_2024]
- Coherence, Informativeness



Subtask Performance

Quality Dimensions (Which Basis)	Task Effectiveness	Dialog / Task Efficiency	Conversational Quality
	Subtask Performance	Trust & User satisfaction	Ethical Considerations & System Integrity

Definition: Evaluating intermediate components within a modular or hybrid GenCRS pipeline.

Goal: Ensure that individual modules (e.g., intent recognizer, retriever) are performing well.

Examples of Subtasks & Metrics:

- Intent Recognition: Accuracy
- Slot Filling: Precision
- Entity Extraction: F1-Score

Examples from Research:

- [feng_2023] evaluates preference elicitation (NDCG@k) and explanation generation (BLEU, Distinct) as separate subtasks.
- [srivastava_2023] assesses their retrieval module using Precision, Recall, and F1-scores.



Trust & User Satisfaction

Quality Dimensions (Which Basis)	Task Effectiveness	Dialog / Task Efficiency	Conversational Quality
	Subtask Performance	Trust & User satisfaction	Ethical Considerations & System Integrity

Definition: The user's subjective perception of the interaction, including trust, usability, and enjoyment.

Core Question: *"How enjoyable was the interaction?" or "How confident are you in the recommendations?"*

Measurement: Almost always requires human feedback via user studies and post-task surveys.

Examples:

- Measuring if users actually follow the recommendations (user action-taking) [wu_2024].
- Assessing the perceived relevance of recommendations based on the dialogue [leszczynski_2023].
- Other works focusing on this include [baizal_2020, kunstmann_2024, abu-rasheed_2024, sun_2024].



Ethical Considerations & System Integrity

Quality Dimensions (Which Basis)	Task Effectiveness	Dialog / Task Efficiency	Conversational Quality
	Subtask Performance	Trust & User satisfaction	Ethical Considerations & System Integrity

Definition: Evaluating fairness, safety, and trustworthiness, especially with powerful LLMs.

Key Areas for Evaluation:

- **Factuality & Hallucination:** *Is the system generating factually correct information? [dehbozorgi_2024, he_2023, mao_2023].*
- **Instruction Faithfulness:** *Does the model follow instructions correctly? [tsai_2024].*
- **Bias:** *Are recommendations fair and not stereotyped?*
- **User Privacy:** *Is sensitive user data protected from leakage?*
- **Robustness:** *Is the system vulnerable to adversarial manipulation?*



Status: A significant gap in current research. These critical issues are rarely evaluated rigorously and represent an imperative direction for future work.

Overview (How)

Paradigms

(How)

Offline

Online

User Studies & Lab Experiments

- **Offline Evaluation (or Simulations)**


- **What:** Evaluation on static, pre-collected datasets.
- **Prevalence:** The most dominant paradigm (**~80% of papers**).

- **Online Experimentation (A/B Testing)**

- **What:** Deployment in a live environment with real users.
- **Prevalence:** The least common paradigm (**~6% of papers**).

- **User Studies & Lab Experiments**

- **What:** Controlled experiments with recruited human participants.
- **Prevalence:** A crucial middle ground, used in **~60% of papers**, often to supplement offline evaluation.



Offline Evaluation

Paradigms

(How)

Offline

Online

User Studies & Lab Experiments

Definition: Evaluation on static datasets without real-time user interaction.

Advantage: Efficient, low cost, and enables rapid, repeatable experiments.

Two Main Approaches:

- **Automated Metrics on Static Datasets:**
 - Calculating standard metrics on a fixed test set.
 - Examples: Recommendation accuracy (Recall@K, NDCG@K) or conversational quality (BLEU, ROUGE).
- **User Simulation:**
 - Using synthetic users (often LLMs) to interact with the system.
 - Examples
 - **(iEvaLM):** A framework using LLM-based user simulators to emulate diverse interactions, improving over static evaluation [wang_2023b].
 - Using session & turn-level controls to ensure simulated conversations are nearly indistinguishable from real ones [friedman_2023].

Offline Evaluation

Paradigms

(How)

Offline

Online

User Studies & Lab Experiments

Definition: Evaluation on static datasets without real-time user interaction.

Advantage: Efficient, low cost, and enables rapid, repeatable experiments.

Two Main Approaches:

- **Automated Metrics on Static Datasets:**
 - Calculating standard metrics on a fixed test set.
 - Examples: Recommendation accuracy (Recall@K, NDCG@K) or conversational quality (BLEU, ROUGE).
- **User Simulation:**
 - Using synthetic users (often LLMs) to interact with the system.
 - Examples
 - **(iEvalLM):** A framework using LLM-based user simulators to emulate diverse interactions, improving over static evaluation [wang_2023b].
 - Using session & turn-level controls to ensure simulated conversations are nearly indistinguishable from real ones [friedman_2023].

Example: iEvaLM

Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models

Xiaolei Wang^{1,2}, Xinyu Tang^{1,2}, Wayne Xin Zhao^{1,3}, Jingyuan Wang⁴ and Ji-Rong Wen^{1,2,3}


¹Gaoling School of Artificial Intelligence, Renmin University of China

²School of Information, Renmin University of China

³Beijing Key Laboratory of Big Data Management and Analysis Methods

⁴School of Computer Science and Engineering, Beihang University

wxl1999@foxmail.com, txy20010310@163.com, batmanfly@gmail.com



Problem/Motivation:

- Existing evaluations of Conversational Recommender Systems (CRSs) often rely on matching static, annotated “ground-truth” dialogs and recommendation items

Contributions:

- Propose a new interactive evaluation framework called iEvaLM that uses LLM-based user simulators to mimic realistic multi-turn interactions.
- Use ground-truth items from CRS datasets to define the persona / target of the simulated user. The simulated user “likes” those items, but must not reveal them directly
- Evaluate with two types of interaction: (a) attribute-based question answering, (b) free-form chit-chat.
- Show that ChatGPT’s performance (in accuracy / recall, and explainability / persuasiveness) improves greatly under this new framework, often surpassing prior methods when allowed to interact.

Datasets & Baselines

- Two public CRS datasets: REDIAL (movies) and OPENDIALKG (multi-domain: movies, books, sports, music).
- Baselines include supervised CRS models (e.g., UniCRS, BARCOR, KBRD, etc.) and unsupervised approaches.

Evaluation Metrics

- Accuracy (Recall@k)** over recommended items, over interaction rounds.
- Explainability / Persuasiveness:** How convincing are explanations for the recommendation? Using human annotation and also an LLM-based scorer

Online Evaluation

Paradigms

(How)

Offline

Online

User Studies & Lab Experiments

Definition: The "gold standard" for assessing real-world impact by deploying a system to live, unsuspecting users.

Key Metrics: Often tied to business objectives like conversion rates, user engagement, and revenue.

Challenges: High cost, resource-intensive, and requires a stable, production-ready system. This makes it rare in academic research.

Examples:

- **[kunstmann_2024]:** Integrated "EventChat" into a live mobile app for a two-month field study to evaluate real-world usage and system performance.
- **[nie_2024]:** Deployed a system on JD.com for a 7-day A/B test, measuring an increase in gross merchandise value (GMV) per user (+1.7%).
- **[leszczynski_2023]:** Conducted an online experiment with crowdworkers who rated the quality of music recommendations from their "TalkTheWalk" model.

Online Evaluation

Paradigms

(How)

Offline

Online

User Studies & Lab Experiments

Definition: The "gold standard" for assessing real-world impact by deploying a system to live, unsuspecting users.

Key Metrics: Often tied to business objectives like conversion rates, user engagement, and revenue.

Challenges: High cost, resource-intensive, and requires a stable, production-ready system. This makes it rare in academic research.


Examples:

- **[kunstmann_2024]:** Integrated "EventChat" into a live mobile app for a two-month field study to evaluate real-world usage and system performance.
- **[nie_2024]:** Deployed a system on JD.com for a 7-day A/B test, measuring an increase in gross merchandise value (GMV) per user (+1.7%).
- **[leszczynski_2023]:** Conducted an online experiment with crowdworkers who rated the quality of music recommendations from their "TalkTheWalk" model.

Example: EventChat

EventChat: Implementation and user-centric evaluation of a large language model-driven conversational recommender system for exploring leisure events in an SME context.

Hannes Kunstmann^{1†}, Joseph Ollier^{2†*}, Joel Persson¹, Florian von Wangenheim^{1,2}



- A real conversational recommender system (CRS) built for a startup in the leisure/events domain (SME).
- Uses ChatGPT (via API) as the LLM core, plus prompt-based learning, a stage-based architecture, retrieval + ranking (RAG), an events database etc.
- Combine subjective (user satisfaction, perceived accuracy, usefulness) + objective metrics (latency, cost, success rates, interaction logs).
- Measure real users in the field, not just lab or simulated ones.
- Track and log failure cases to understand bottlenecks (e.g. when event isn't found, prompt misinterpretation, missing metadata).
- Monitor trade-offs: better accuracy / richer features often increase latency and cost; SMEs must balance.
- Use lightweight survey instruments (short-form ResQue etc.) to not overburden users.

User Studies & Lab Experiments

Paradigms

(How)

Offline

Online

User Studies & Lab Experiments

Definition: A middle ground that captures genuine human behavior in a controlled experimental setting without the risks of a full online deployment.

Primary Goal: Essential for assessing subjective quality dimensions that automated metrics cannot capture.

- **Examples:** User satisfaction, trust, naturalness, persuasiveness.

Common Format:

- Recruiting participants to perform specific tasks.
- Gathering feedback through post-task surveys, interviews, or think-aloud protocols.

Usage: Frequently used to supplement offline evaluations, providing a critical layer of human-centric validation.

Overview (Who)

Evaluator


(Who)

Automated Metrics

Humans

LLM-as-a-Judge

- Automated Evaluators
 - Algorithmic approaches using pre-defined metrics.
 - **Prevalence:** Most common (38/47 papers).
- Human Annotators
 - Experts or crowdworkers providing qualitative assessments.
 - **Prevalence:** Very common, often for validation (28/47 papers).
- LLM-based Evaluators (LLM-as-Judge)
 - Using Large Language Models to assess quality.
 - **Prevalence:** An emerging but fast-growing approach (3/47 papers).



Automated Metrics

Evaluator

(Who)

Automated Metrics

Human

LLM-as-a-Judge

Definition: Algorithms that compute objective, quantifiable metrics on static datasets.

Pros:

- Consistent, scalable, and efficient.
- Ideal for large-scale offline evaluation.

Cons:

- Often fail to capture nuance, especially in conversational quality and user satisfaction.

Key Categories of Metrics:

- Ranking & Recommendation Quality
- Natural Language Generation (NLG) & Text Quality
- Classification & Prediction Accuracy
- Regression & Prediction Error

Automated Metrics (Examples)

Evaluator

(Who)

Automated Metrics

Humans

LLM-as-a-Judge

- **Ranking & Recommendation Quality**


- Recall@K, HitRate@K, MRR, NDCG@K: Assess accuracy and relevance of recommendations. Used in the majority of papers

- **NLG & Text Quality**

- BLEU, ROUGE: Measure n-gram overlap with reference text
- PPL (Perplexity): Measures language model fluency
- DIST (Distinct-n): Measures response diversity
- BERTScore: Measures semantic similarity

- **Classification & Prediction**

- Accuracy, F1-Score, AUC: Assess performance on tasks like intent recognition.
- MAE, RMSE: Measure error in rating prediction tasks.



Human Annotators

Evaluator

(Who)

Automated Metrics

Human

LLM-as-a-Judge

- **What they are:** Domain experts or crowdworkers providing qualitative assessments.
- **Pros:**
 - Excel at assessing subjective dimensions that automated metrics often miss: e.g., coherence, naturalness, helpfulness, trust.
- **Cons:**
 - Time-consuming and expensive.
 - Often used to validate or supplement automated metrics.
- **Process:**
 - Use Likert scales or binary ratings.
 - Reliability is checked with inter-annotator agreement metrics (e.g., Cohen's Kappa)

Common Subjective Metrics (Assessed by Humans)

Evaluator

(Who)


Automated Metrics

Human

LLM-as-a-Judge

Humans are essential for evaluating:

- **S1: Fluency, Grammar, & Readability (27.7% papers)**
- **S2: Coherence & Logicity (12.8% papers)**
- **S3: Informativeness & Helpfulness (31.9% papers)**
- **S4: Relevance & Answerability (10.6% papers)**
- **S5: User Satisfaction, Enjoyment, & Future Use (6.4% papers)**
- **S6: Trustworthiness, Persuasiveness, & Explainability (8.5% papers)**
- **S7: Significant Gap:** Beyond-accuracy metrics like novelty and serendipity are rarely evaluated (8.5% papers)



LLM-based Evaluators (LLM-as-Judge)

Evaluator

(Who)

Automated Metrics

Humans

LLM-as-a-Judge

Definition: An emerging paradigm using LLMs to evaluate the outputs of other models.

Role: Act as a scalable, fast, and consistent alternative to human annotators.

Challenge: Reliability can vary, and results may not always align with human judgment.

Examples:

- **Rec-SAVER [tsai_2024]:** An LLM generates and then self-verifies reasoning references to create an evaluation benchmark automatically.
- **iEvalM [wang_2023]:** An LLM acts as both a user simulator to create interactions and an evaluator to score metrics like persuasiveness.
- **[sayana_2025]:** Uses Gemini Pro with ensemble rating (averaging over multiple runs) to score generated text on a 7-point Likert scale.

LLM-based Evaluators (LLM-as-Judge)

Evaluator

(Who)

Automated Metrics

Humans

LLM-as-a-Judge

Definition: An emerging paradigm using LLMs to evaluate the outputs of other models.

Role: Act as a scalable, fast, and consistent alternative to human annotators.

Challenge: Reliability can vary, and results may not always align with human judgment.

Examples:


- **Rec-SAVER [tsai_2024]:** An LLM generates and then self-verifies reasoning references to create an evaluation benchmark automatically.
- **iEvalM [wang_2023]:** An LLM acts as both a user simulator to create interactions and an evaluator to score metrics like persuasiveness.
- **[sayana_2025]:** Uses Gemini Pro with ensemble rating (averaging over multiple runs) to score generated text on a 7-point Likert scale.

Example: Rec-SAVER

Leveraging LLM Reasoning Enhances Personalized Recommender Systems

Alicia Y. Tsai^{*†1,3}, Adam Kraft^{*3}, Long Jin², Chenwei Cai²,
Anahita Hosseini³, Taibai Xu², Zemin Zhang², Lichan Hong³,
Ed H. Chi³, Xinyang Yi³

¹University of California, Berkeley ²Google ³Google DeepMind



- Rec-SAVER (Recommender Systems Automatic Verification and Evaluation of Reasoning) is introduced to automatically assess quality of reasoning outputs from LLMs without needing human raters or curated gold references.
- It does this by generating explanations + then performing self-verification, i.e. re-predicting user ratings based on those explanations and comparing with actual ratings. If the explanation supports a correct rating, it counts more favorably.
- It uses multiple automatic NLG / text similarity / coherence / faithfulness metrics (e.g. BLEU, ROUGE, METEOR, BERTScore) to evaluate different dimensions of reasoning: how coherent the reasoning is, how faithful (i.e. correct with respect to inputs), how insightful.
- They validate that Rec-SAVER's automatic judgments correlate well with human judgments on coherence and faithfulness, thus making it a reliable judge / benchmark.
- Using Rec-SAVER, they are able to compare models (zero-shot vs fine-tuned; smaller vs larger) not only on rating-prediction accuracy, but also on reasoning quality. They show that adding reasoning helps improve recommendation performance

Overview (For Whom)

Stakeholders

(For Whom)

Consumers

Item-Providers


Others

A crucial, yet often overlooked, question is: *For whom is the evaluation being conducted?*

We can categorize the focus of an evaluation based on its primary beneficiary:

1. The Consumer (End-User)
2. The Item Provider (e.g., sellers, content creators)
3. Other Stakeholders (e.g., the platform, society at large)

 **Ideal Scenario :**  A GenCRS that balances the **needs** of **all stakeholders** i.e. is **fair** to all of them



The typical multi-stakeholder environment in a hotel booking scenario**

Reality !!!

Stakeholders


(For Whom)

Consumers

Item-Providers

Others

- A staggering **93% of surveyed works** focus exclusively on the consumer.
 - This is understandable, as the main goal of a GenCRS is to satisfy the end-user.
- As a result, evaluations are dominated by user-centric metrics:
 - Recommendation Relevance
 - Dialogue Quality
 - User Satisfaction
- Evaluations considering other stakeholders are exceptionally rare.



The Gap: Item Providers & Society

Stakeholders

(For Whom)

Consumers

Item-Providers

Others

- Item Providers:
 - No studies explicitly focus on provider benefits (e.g., increased sales, fair exposure).
 - Some works indirectly address their interests:
 - Mitigating popularity bias to enhance fairness and visibility for less popular items [wang_2023].
 - Tackling user-item rating bias for a fairer assessment of items [kim_2024].



The Gap: Item Providers & Society

Stakeholders


(For Whom)

Consumers

Item-Providers

Others

- **Society at large:** Represents a major blind spot in current evaluation practices.
 - Environmental Cost: The significant carbon footprint of LLMs is **almost entirely ignored** and undocumented in the GenCRS literature.
- **Emerging Progress: Incorporating Sustainability (explicitly catering to Society as a stakeholder)**
 - Examples:
 - **System Design:** e.g., Collab-REC*.
 - **Data Generation:** e.g., SynthTRIPS**.
 - Promising shift towards multi-stakeholder evaluation frameworks.
- **Takeaway:** Expanding evaluation to include these broader impacts is an ethical imperative.



*Banerjee, Ashmi, et al. "Collab-REC: An LLM-based Agentic Framework for Balancing Recommendations in Tourism." arXiv preprint arXiv:2508.15030 (2025).

**Banerjee, Ashmi, et al. "SynthTRIPS: A Knowledge-Grounded Framework for Benchmark Data Generation for Personalized Tourism Recommenders." SIGIR 2025

The Gap: Item Providers & Society

Stakeholders


(For Whom)

Consumers

Item-Providers

Others

- **Society at large:** Represents a major blind spot in current evaluation practices.
 - Environmental Cost: The significant carbon footprint of LLMs is **almost entirely ignored** and undocumented in the GenCRS literature.
- **Emerging Progress: Incorporating Sustainability (explicitly catering to Society as a stakeholder)**
 - Examples:
 - **System Design:** e.g., Collab-REC*.
 - **Data Generation:** e.g., SynthTRIPS**.
 - Promising shift towards multi-stakeholder evaluation frameworks.
- **Takeaway:** Expanding evaluation to include these broader impacts is an ethical imperative.



*Banerjee, Ashmi, et al. "Collab-REC: An LLM-based Agentic Framework for Balancing Recommendations in Tourism." arXiv preprint arXiv:2508.15030 (2025).

**Banerjee, Ashmi, et al. "SynthTRIPS: A Knowledge-Grounded Framework for Benchmark Data Generation for Personalized Tourism Recommenders." SIGIR 2025

Example: Collab-REC

Collab-REC: An LLM-based Agentic Framework for Balancing Recommendations in Tourism

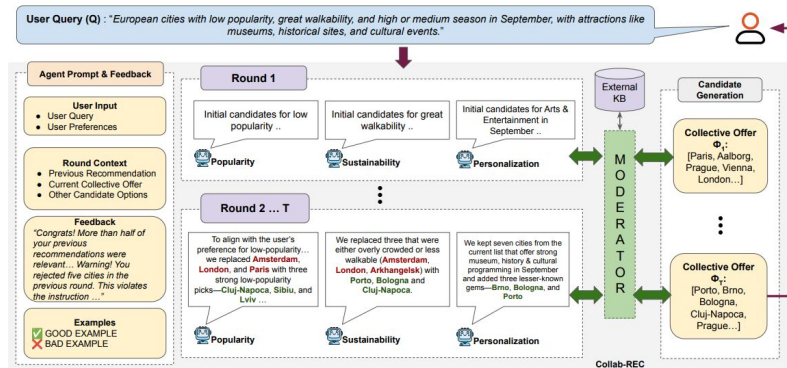
Ashmi Banerjee
ashmi.banerjee@tum.de
Technical University of Munich
Munich, Germany

Fitri Nur Aisyah
fitri.aisyah@tum.de
Technical University of Munich
Munich, Germany

Adithi Satish
adithi.satish@tum.de
Technical University of Munich
Munich, Germany

Wolfgang Wörndl
woerndl@in.tum.de
Technical University of Munich
Munich, Germany

Yashar Deldjoo
yashar.deldjoo@poliba.it
Polytechnic University of Bari
Bari, Italy



- It includes a Sustainability Agent among its core multi-agent setup, whose role is to promote eco-centric criteria (e.g. walkability, air quality, seasonality) when proposing city recommendations.
- The framework uses multi-round negotiation, where the **Sustainability Agent's** suggestions are combined with those of a **Personalization Agent** and a **Popularity Agent**; a moderator enforces trade-offs so that sustainability isn't drowned out by more popular or purely preference-based choices.
- The moderator (non-LLM) integrates penalties for repeated or invalid proposals and scores candidates by factors including agent success, reliability, and hallucination penalty — this helps to ensure sustainability suggestions make it through even when they conflict with popularity bias.
- Empirical results show Collab-REC improves diversity of recommendations (lesser-known / less popular cities surfaced) and reduces popularity bias, thus contributing toward more socially sustainable tourism (e.g. avoiding over-tourism).

The Gap: Item Providers & Society

Stakeholders


(For Whom)

Consumers

Item-Providers

Others

- **Society at large:** Represents a major blind spot in current evaluation practices.
 - Environmental Cost: The significant carbon footprint of LLMs is **almost entirely ignored** and undocumented in the GenCRS literature.
- **Emerging Progress: Incorporating Sustainability (explicitly catering to Society as a stakeholder)**
 - Examples:
 - **System Design:** e.g., Collab-REC*.
 - **Data Generation:** e.g., SynthTRIPS**.
 - Promising shift towards multi-stakeholder evaluation frameworks.
- **Takeaway:** Expanding evaluation to include these broader impacts is an ethical imperative.



*Banerjee et al. "Collab-REC: An LLM-based Agentic Framework for Balancing Recommendations in Tourism." arXiv preprint arXiv:2508.15030 (2025).

**Banerjee et al. "SynthTRIPS: A Knowledge-Grounded Framework for Benchmark Data Generation for Personalized Tourism Recommenders." SIGIR 2025

Example: SynthTRIPS

SynthTRIPS: A Knowledge-Grounded Framework for Benchmark Query Generation for Personalized Tourism Recommenders


Ashmi Banerjee
ashmi.banerjee@tum.de
Technical University of Munich
Munich, Germany

Adithi Satish
adithi.satish@tum.de
Technical University of Munich
Munich, Germany

Fitri Nur Aisyah
fitri.aisyah@tum.de
Technical University of Munich
Munich, Germany

Wolfgang Wörndl
woerndl@in.tum.de
Technical University of Munich
Munich, Germany

Yashar Deldjoo
yashar.deldjoo@poliba.it
Polytechnic University of Bari
Bari, Italy



- Generates a list of personalized, synthetic queries for European city trip recommendations. Also, caters for sustainable trips.
- Example:
 - Persona: *"A wanderlust-filled trader who appreciates and sells the artisan's creations in different corners of the world"*
 - Filters: **Popularity = Low; Interest = Nightlife Spot**
 - Persona-Specific Query: *"Unique nightlife and cultural experiences in off-the-beaten-path European cities for a budget-conscious traveler interested in local artisans."*

Evaluation of GenCRS: Challenges

Overemphasis on Ground-Truth Matching

- Traditional evaluations focus on ground-truth matching, overlooking the interactive and evolving nature of CRS from the user's perspective


Limitations of LLM-Based Evaluators

- **Limited Human Representation:** Constrained by prompt templates and datasets; may miss real user complexity and dynamics.
- **Bias Propagation:** LLM evaluators can embed and amplify biases, reducing fairness
- **Ethical and Privacy Concerns:** Handling conversational data raises significant ethical and privacy risks.



Evaluation: Trends

PRE-2023: THE STATIC PARADIGM




Pre-2023: The Static Paradigm

- Evaluation was predominantly static and relied on ground-truth-based metrics.
- **Key Metrics:** Recall@k, BLEU, Distinct-n.
- **Key Limitations****
 - Failed to capture the interactive and subjective nature of dialogue (e.g., conversation quality, user engagement).
 - Ineffective for open-ended, free-text responses where no single "correct" answer exists.
 - Could not properly assess issues like hallucination.

**Jannach, Dietmar, et al. "A survey on conversational recommender systems." *ACM Computing Surveys (CSUR)* 54.5 (2021): 1-36.

Evaluation: Trends

PRE-2023: THE STATIC PARADIGM



POST-2023: HOLISTIC & INTERACTIVE




Post-2023: The Shift to Holistic, Interactive Evaluation with a user-centric perspective

- **The Rise of "LLM-as-a-Judge":**
 - LLMs are increasingly used as evaluators to assess subjective qualities like coherence, helpfulness, and naturalness
- **New Evaluation Dimensions Have Emerged:**
 - Factual Accuracy & Hallucination Detection
 - Faithfulness to user instructions
 - Groundedness on external knowledge
- **New Practical Concerns:**
 - System latency has become a key metric, as complex agent setups can impact user satisfaction.

Evaluation: Summary

- Evaluation is evolving from a **static to a multidimensional and interactive paradigm**.
- **The Core Shift in Focus:**
 - **FROM:** Matching a single 'gold-standard' response...
 - **TO:** Measuring how helpful, human-like, and contextually appropriate a system feels.
- **The Dual Role of LLMs:**
 - LLMs now function as both the recommender agent being tested and the automated evaluator assessing performance.
- This new paradigm enables **more scalable and nuanced assessments** of the overall user experience.





Yashar Deldjoo

Open Challenges & Future Directions

Agenda

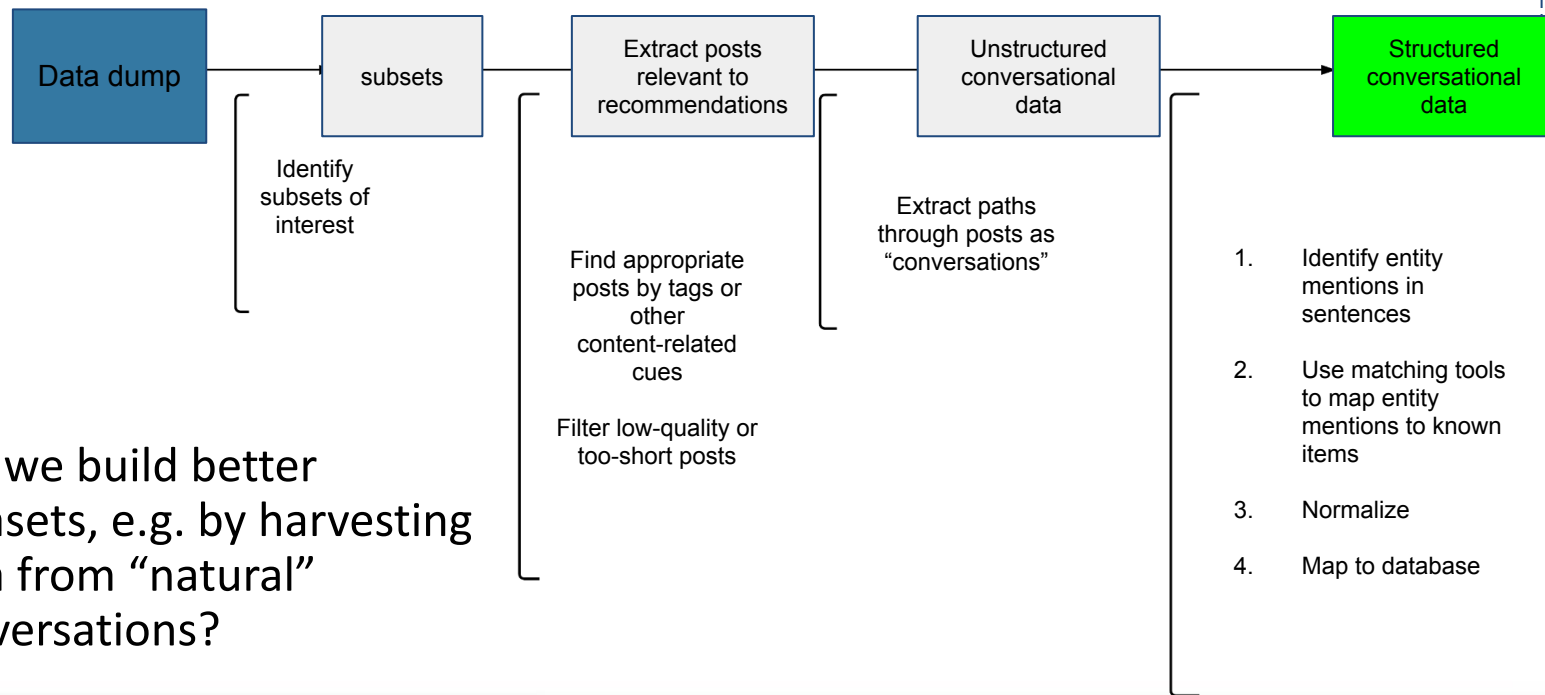
- Introduction
- Core Systems & Components
- Foundation Model Integration & Generative Paradigms
- Knowledge and Data Foundation
- Simulation
- Evaluation
- **Open Challenges & Future Directions**

Datasets

- How can datasets be built that are more ***natural***? E.g. actually how humans would interact when making movie recommendations, versus current, more synthetic, settings?
- Other efforts (e.g. INSPIRED) aim for a more natural setting, but are also very small
- Need datasets that are **bigger** and **more realistic**
- Our previous efforts (e.g. to synthesize conversational datasets from product review text) were much larger but of low quality

Dataset Construction Pipeline

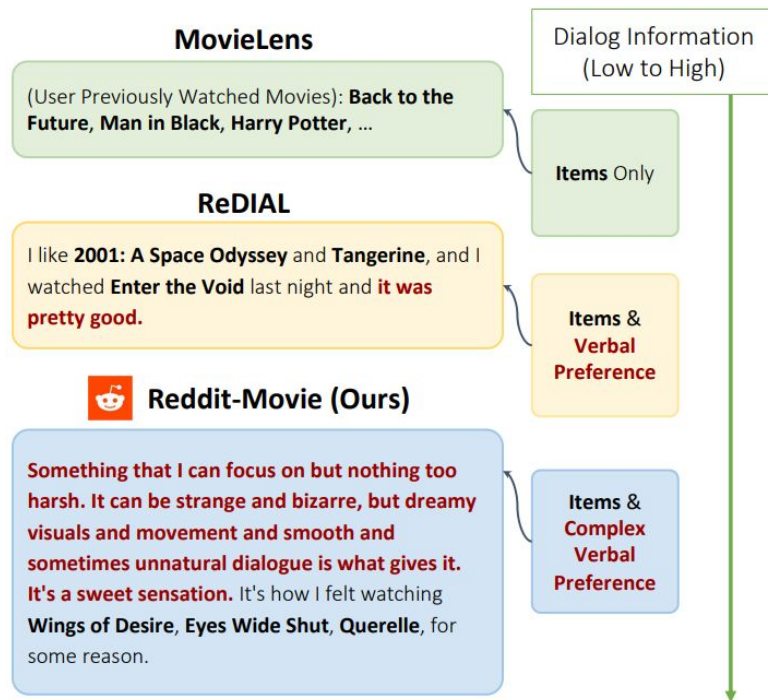
For Models ...



Reddit-Movie Dataset

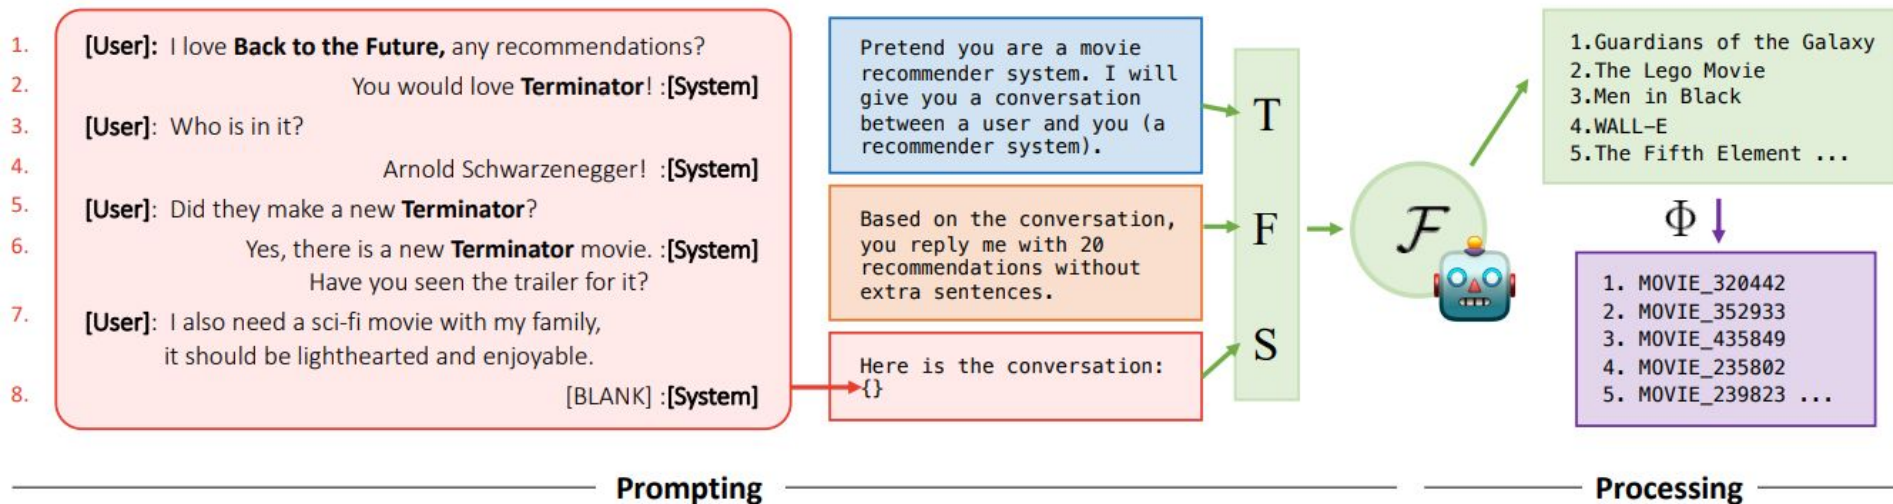
- **634,392** movie recommendation conversations, featuring 1.7M dialog turns
- ~11k users, ~24k items
- (compare to e.g. ReDial, featuring ~10k conversations, ~139k turns, ~800 users)

Much bigger than existing datasets; conversations are shorter; they have much more *context*; and (for better or worse) have much more varying structure



What do these new datasets reveal?


We use a simple prompting setup to compare LLMs:



What do these new datasets reveal?

Some observations about model performance:





- Existing models engage in *shortcut learning* by focusing on repeated items (i.e., items already mentioned in a dialog but not as recommendations)
- LLMs outperform existing fine-tuned models; GPT-4 outperforms other LLMs**
- LLMs generate some out-of-dataset items, but not many hallucinated recommendations (<5%); can be dealt with by string matching



What do these new datasets reveal?

Some observations about model performance:

- **Significant “popularity bias”** (and other bias) issues
- Recommendation performance is highly sensitive to geographical region (presumably just due to ground truth frequency)




What do these new datasets reveal?

Some observations about model performance:

- Users mention both *previous items* (collaborative information) and *context* (semantic information) in their queries
- By ablating one or the other (basically, just deleting either items or other text from the query), and measuring performance, we find that (pre-trained) conversational models rely much more on *semantics* than *collaborative* signals
- Suggests that there's still a lot of room for improvement in terms of **leveraging collaborative knowledge** (i.e., recommender systems stuff!) in conversational models

More Datasets (Complex Contexts Created by ChatGPT)

- One thing our datasets reveal is that real-world (or at least online) conversations generally consist of long-context “queries” followed by relatively shallow conversations
- We can construct such datasets synthetically...
- Basically, we can ask an LLM to construct a contextual query using product reviews (fairly easy); we can then evaluate conversational recommenders based on their ability to find the “right” product given one of these contextual queries



qid	query	item_id	user_id	ori_rating	ori_review
0	I need filters that effectively trap dust and improve the air quality in my home. It's surprising how much dust they can collect in just a few months.	B0C5QYHTJ	AGRE0263GTRNY03K4CIQV2DTZLSQ	5	These filters work I collected a lot of dust in the house. I replaced after 3 months the dust and make my home place.
1	I need to find a protein that is...	B0C7D3VLXW	AFLIZT24MDW4XG4HBYKOI3BZGDHQ	5	Love We love this protein it for a few years now
2	I need a pillow that helps keep m...	B0C3QRMPVN	APCSK3W3GI7PGT4655HHKZ2CFPMA	5	CPAP help I use this pillow. It greatly helps

More Datasets ...

- We've mostly looked at movies so far, but where else do people have similar conversations?
- Can look for reddit posts that have Amazon product links as endpoints

category	subreddits	# of conversations	total
broad	advice; random_acts_of_amazon	89,871; 117,305	207,176
electronics	buildapc; pcmasterrace; suggesta-	71,547; 23,602;	131,256
	laptop; buildapcmonitors; mechan-	14,333; 11,425;	
	calkeyboards	10,349	
tech support	techsupport; homenetworking	14,333; 10,532	24,865
fashion	malefashionadvice; watches	11,963; 12,111	24,074
books	books; booksuggestions	9,897; 10,038	19,935
audio-visual	vinyl; vinyldeals; hometheater;	10,722; 12,160;	46,347
	headphones	12,626; 10,839	
DIY	homeimprovement	11,995	11,995

amazon cat.	# of convs.	amazon category	# of convs
books	162,047	amazon home	48,240
buy a kindle	38,012	tools & home impr.	31,672
computers	28,796	toys & games	27,746
all electronics	24,443	sports & outdoors	24,154
amazon fashion	23,638	industrial & scientific	18,771
automotive	18,554	health & personal care	15,550
all beauty	14,235	digital music	14,100
movies & tv	12,518	cell phones & access.	12,492
grocery	10,759	pet supplies	9,055
video games	8746	office products	8,311

More Datasets ...

- Using these, we can easily build datasets containing:
 - Real-world recommendation-oriented conversations
 - Signals from collaborative filtering (mainly to harvest pre-trained item representations based on “denser” data than what is available in conversational datasets)
 - Item metadata etc.
- This can be done elsewhere (much as we’ve done for movies), but for Amazon, the process is trivial as the item IDs are already in reddit conversations

External Item Representations

- **Note:** very similar to the “pre-LLM” state-of-the-art: i.e., conversational components and recommendation components are joined together (which makes sense!)
- Also, not *quite* a conversational model, but rather a “contextual” recommender
- Hints at possible new paradigms, e.g. where users interact with a system by editing a complex query

Lots More ...


Mostly forms of “system building” or RAG strategies:

- **Collaborative retrieval:** Can we retrieve related items (or interactions) to use as prompts to construct evidence for or against particular recommendations? (KDD’24, Wu++)
- **Retrieve related training samples:** Similar to nearest-neighbor language models (RecSys’24, Xie++)
- **Retrieve related knowledge:** E.g. from an external corpus with domain information about items being recommended

Evaluation

Can we do better than held-out item prediction?

- Users may interact with conversational recommenders precisely because they struggle to articulate their preferences, or because they need to be persuaded to select a particular item;
- User studies are expensive, and generally non-reproducible
- Outside of industrial settings, user studies generally don't involve 'real' users
- User studies may be suitable for 'general knowledge' items and domains, but are unsuitable in cases where users requiring specific knowledge or expertise may be difficult to recruit



Open Challenges

- Challenges in providing diverse or novel recommendations during conversations
- Go beyond passive preference elicitation based on parsing input
- Persistent semantic gap between recommendation and response generation
- Scalable grounding methods and knowledge updates
- Lack of longitudinal benchmarks to investigate evolving preferences
- Bias, fairness and ethical considerations
- Matching system capabilities and user expecting to improve user experience and trust
- Scalability and robustness to enable feasible real-world applications
- Effective integration of multi-modal data

Summary

- Conversational recommendation represents a promising frontier in building recommender systems that are more “human-like”
- This line of research has been somewhat blown open by the excellent performance of general-purpose language models
- There’s still plenty to do (even if, arguably, less of it is about modeling...)
- Many “traditional” questions about recommender systems (evaluation, fairness, etc.) have new life in light of conversational paradigms

Link: <https://recsys-lab.at/gen-conv-recsys-tutorial>

