



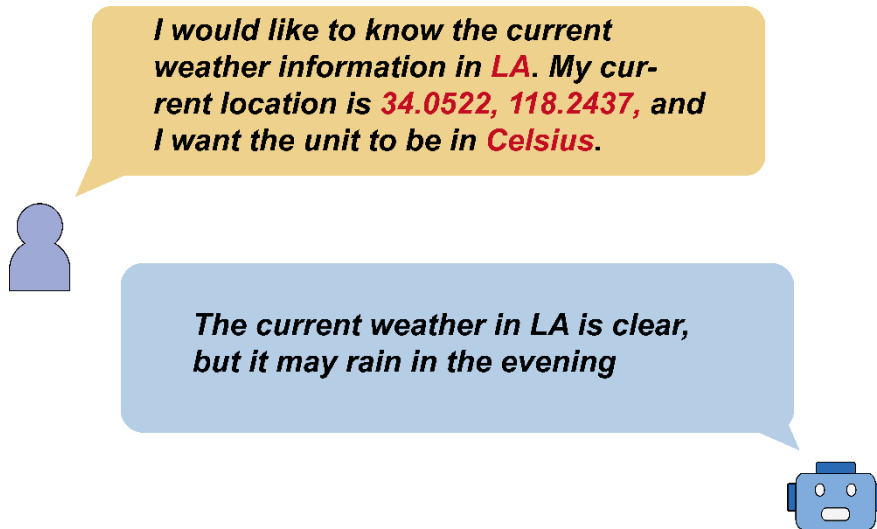
ToolDial: Multi-turn Dialogue Generation Method for Tool-Augmented Language Models

Jeonghoon Shim, Gyuhyeon Seo, Cheongsu Lim, Yohan Jo

Motivation



● Settings of current research on TALM



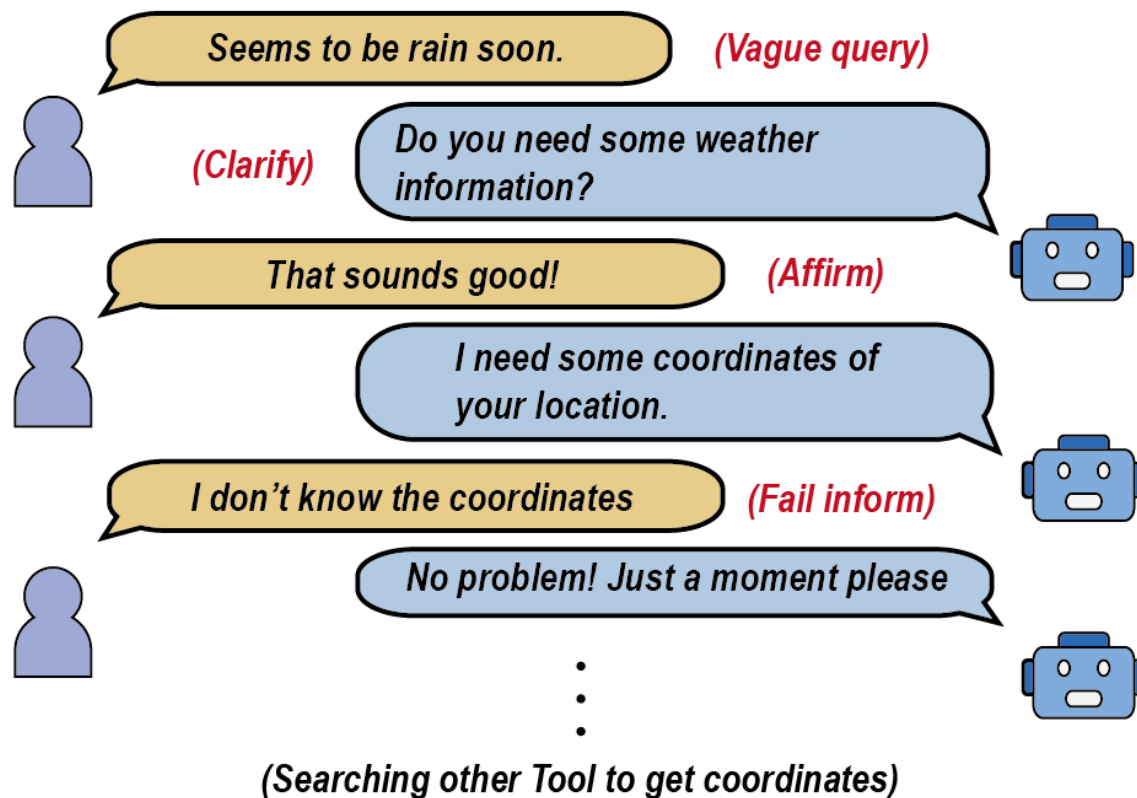
- Only the conversation between the user and TALM's single turn is assumed.
- All the information necessary to make an API call based on the user's query is available.

Is this setting truly aligned with the “real world”?



Motivation

● In real world

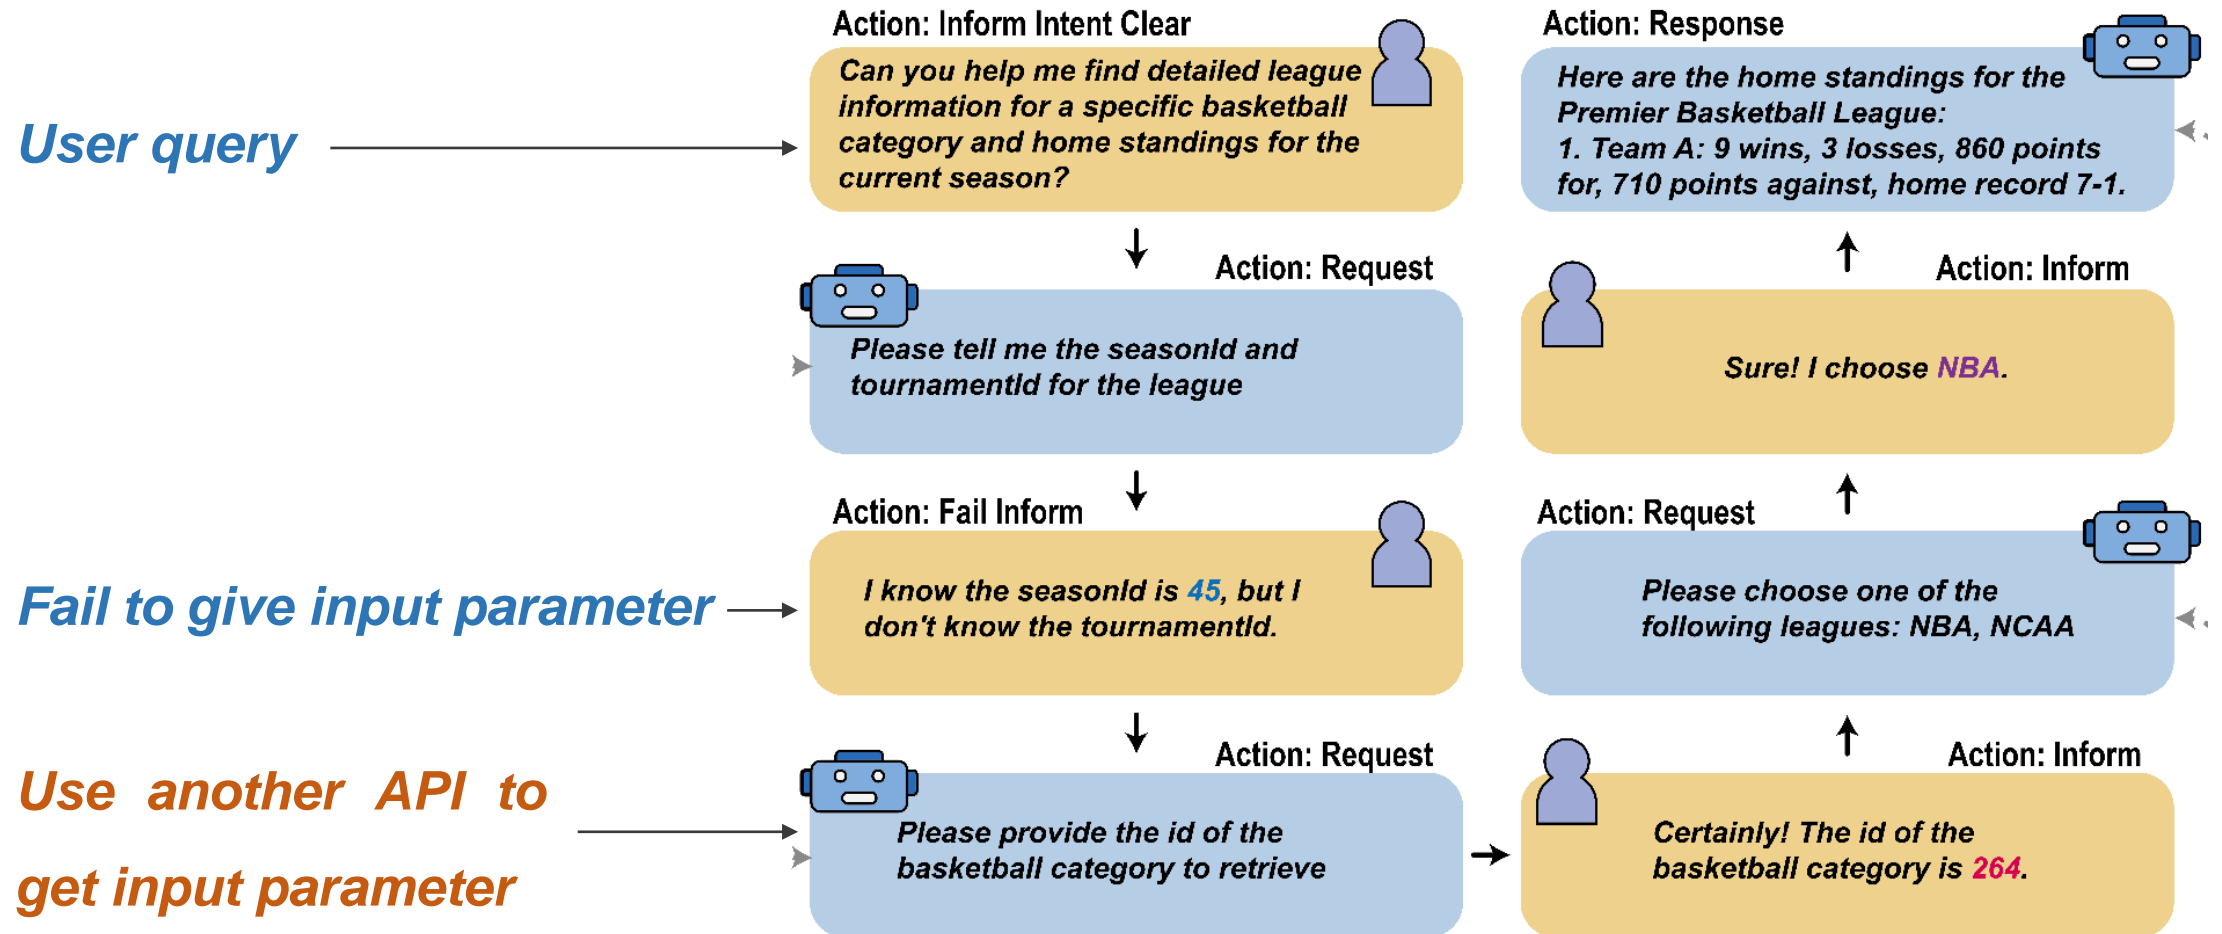


- Situation is more complex than previous settings.
- Multi-turn conversation is needed to fulfill the user's query.

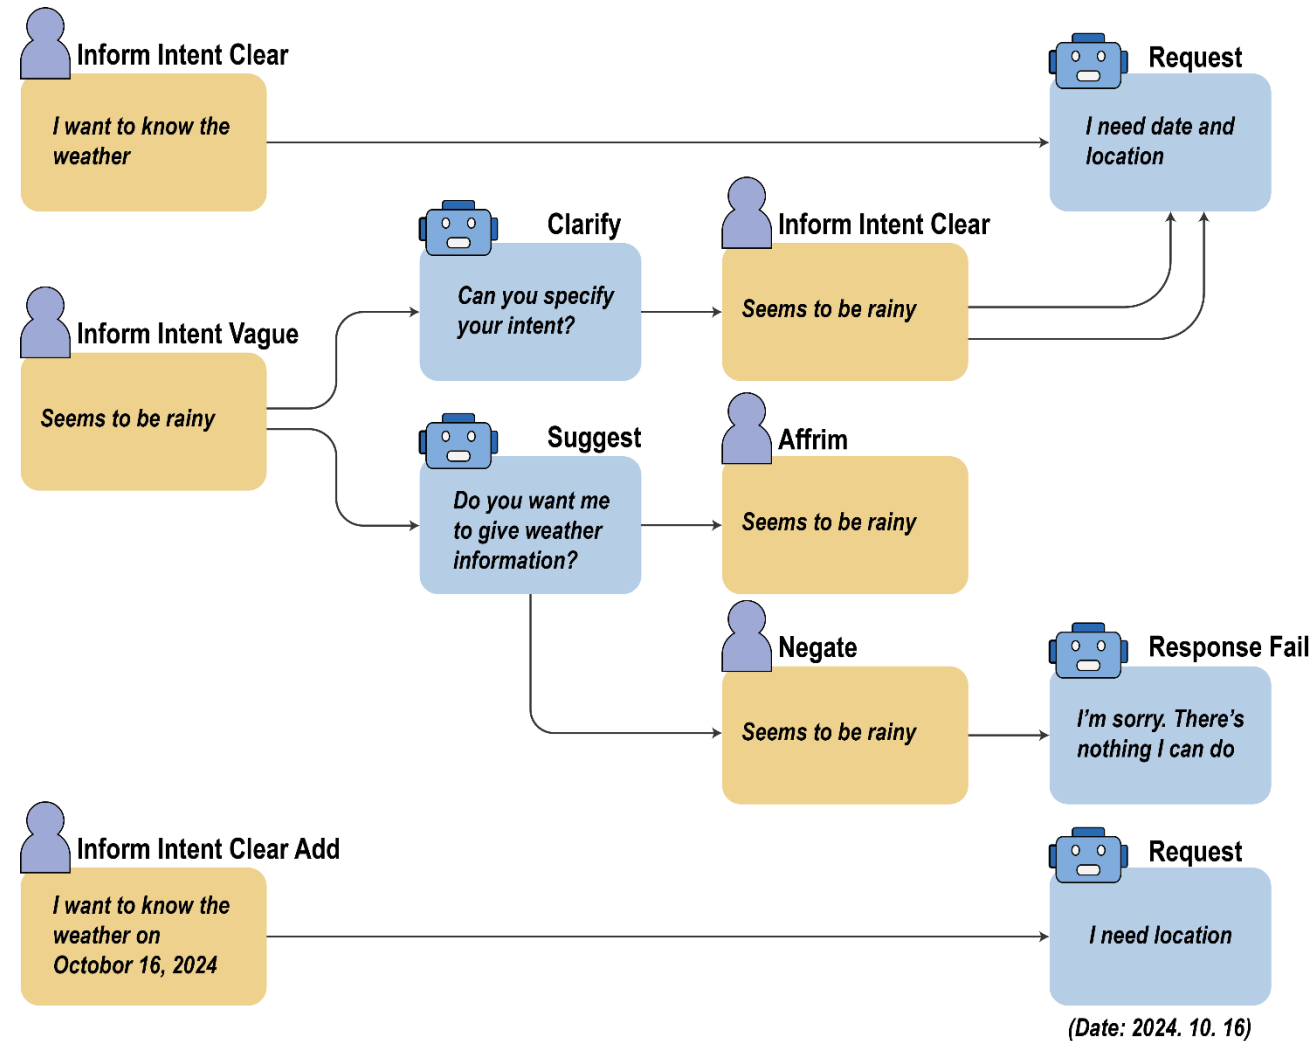
Motivation



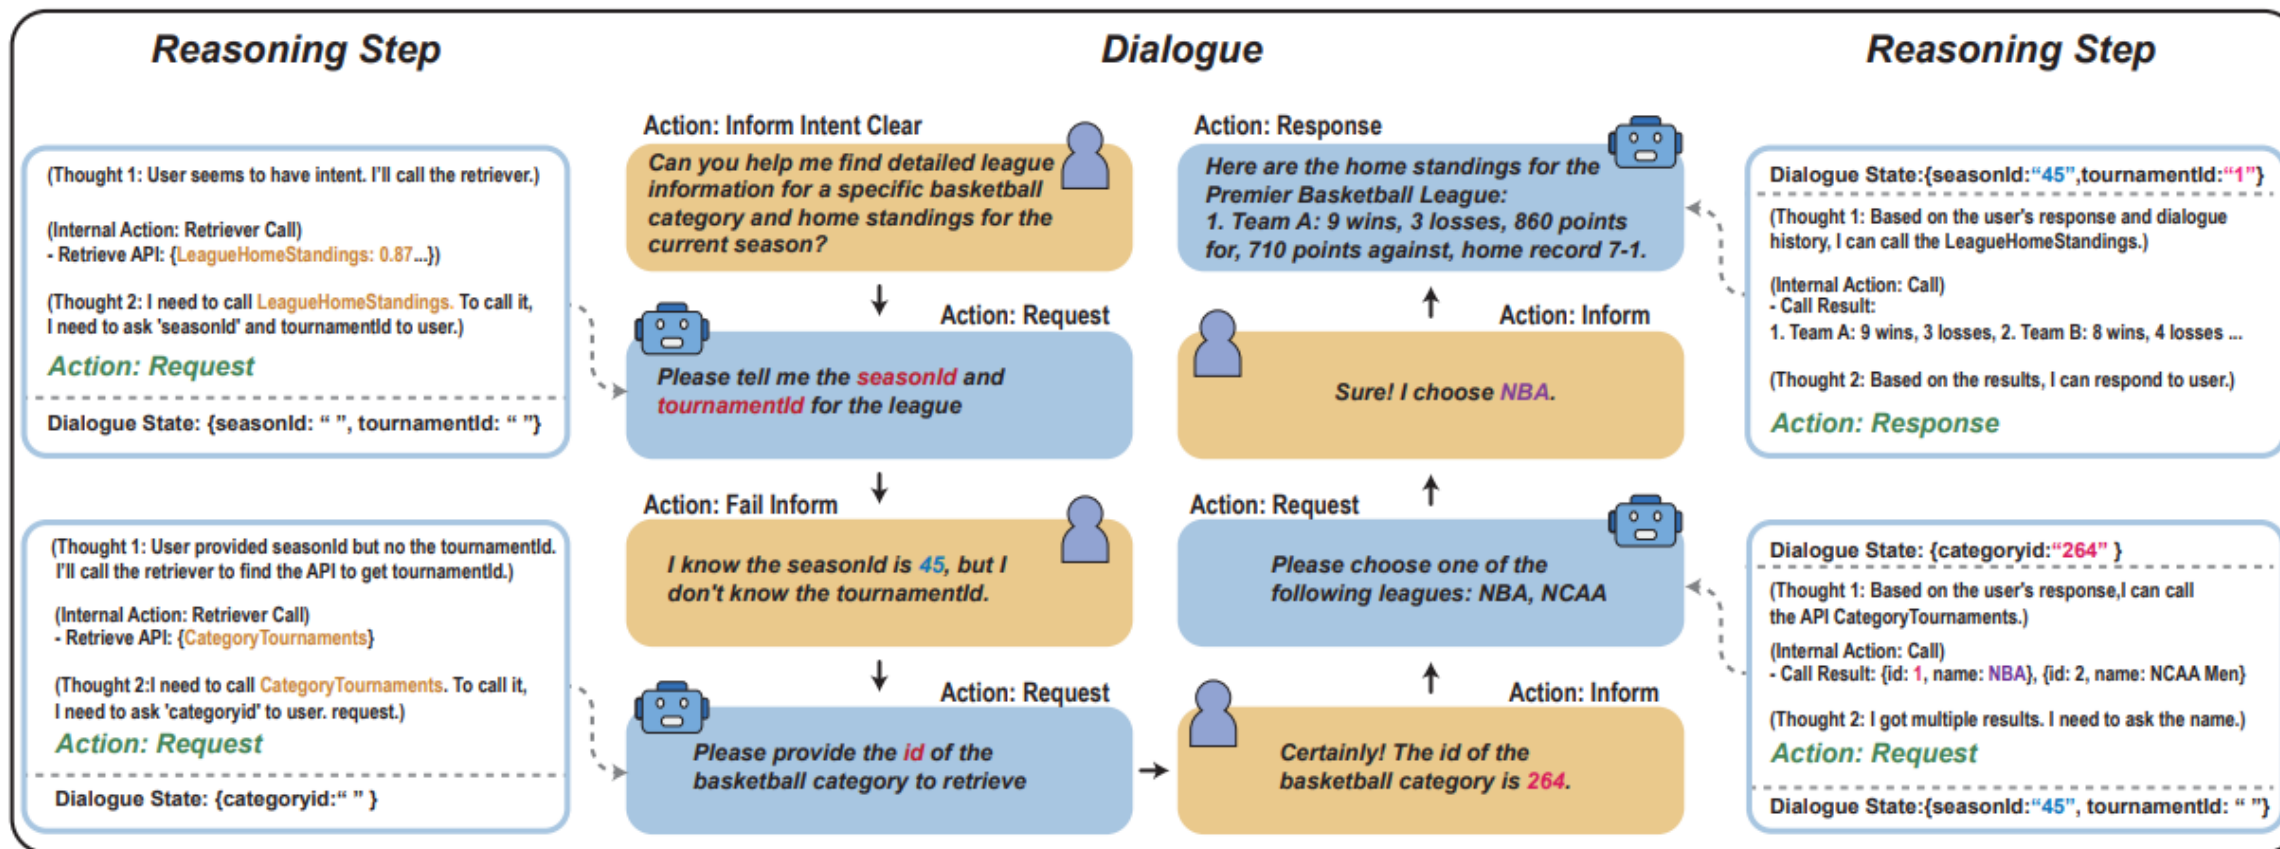
Seoul National University
Graduate School of Data Science



Motivation



● Overall Data



● Dialogue Statistics

Resource	ToolDial	ToolBench	API-Bank	ToolAlpaca
Real-world API?	✓	✓	✓	✗
Multi-turn Scenario?	✓	✗	✓	✗
Multi-tool Scenario?	✓	✓	✓	✗
Multi-step Reasoning?	✓	✓	✓	✗
Situation Complexity?	✓	✗	✗	✗
Number of Actions	16	3	7	3
Number of Dialogues	11,111	188,304	6,860	4,889
Avg. Turn per Dialogue	8.95	2	2.84	2

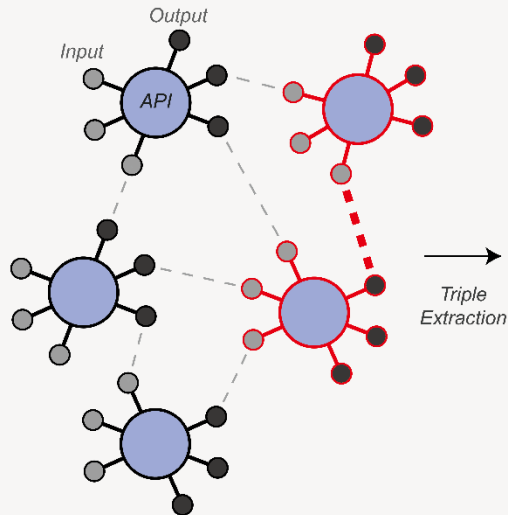
- Compared to previous benchmarks, ToolDial has much longer dialogue turns and incorporates more complex situations.

Method

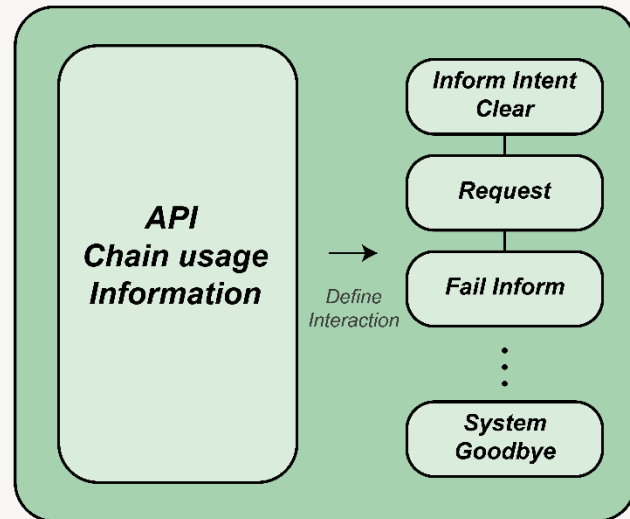


Seoul National University
Graduate School of Data Science

Graph Construction



Action Sequence Construction



Scenario Generation

User turn
- Action: Inform intent clear
- User ask to system with ...

System turn
- Action: Request
- System asks information...

User turn
- Action: Fail Inform
- User can't provide the...

...

System turn
- Action: System Goodbye
- System says good bye



<Dialogue Data>

Inform Intent Clear

User: I want the weather information.

Request

System: I need the coordinates!

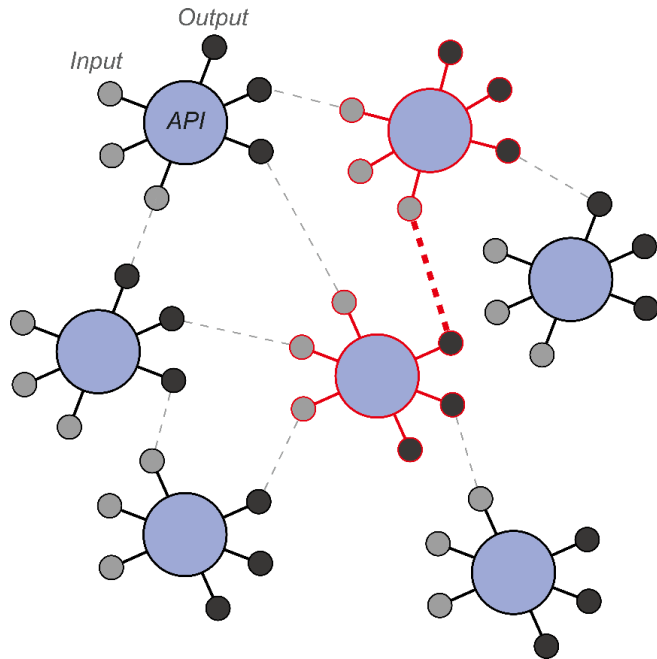
Fail Inform

User: I don't know the coordinates.

System Goodbye

System: Thank you! Goodbye!

● Graph Construction



API docs

Output

Name: "coordinates"

Description: ~~

API docs

Input parameter

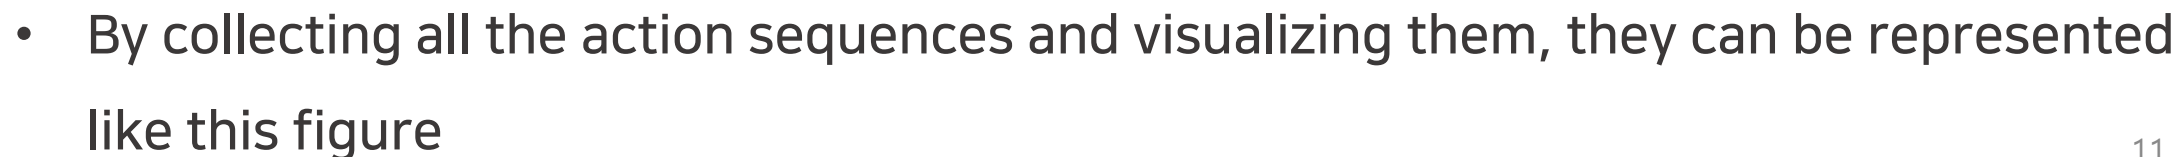
Name: "coordinates"

Description: ~~

- Longest Common Subsequence similarity between name
- Embedding similarity between description



- Action Sequence Construction



● Scenario Instruction Generation

<Scenario Instruction>

<Dialogue state>

- Inform intent clear: the user utters a pre-constructed query related with API LeagueHome-Standings and CategoryTournament.
- (Retriever call) → request: the system to ask the user for **seasonId** and **tournamentId**. {seasonId: " ", tournamentId: " "}
- Fail inform: the user responds with seasonId **45** but fails to provide tournamentId. {seasonId: "45", tournamentId: " "}
- (Retriever call) → request: the system prompts the user for **id**. {id: " "}
- Inform: the user responds with the requested information.
- (call) → request: the system asks the user for the **name** variable, to select one **id** from multiple results.
- Inform: Instruction: the user responds with **NBA**

● Dialogue Data Generation

<Scenario Instruction>

<Dialogue state>

- Inform intent clear: the user utters a pre-constructed query related with API LeagueHome-Standings and CategoryTournament.
- (Retriever call) → request: the system to ask the user for **seasonId** and **tournamentId**. {seasonId: " ", tournamentId: " "}
- Fail inform: the user responds with seasonId **45** but fails to provide tournamentId. {seasonId: "45", tournamentId: " "}
- (Retriever call) → request: the system prompts the user for **id**. {id: " "}
- Inform: the user responds with the requested information.
- (call) → request: the system asks the user for the **name** variable, to select one **id** from multiple results.
- Inform: Instruction: the user responds with **NBA**

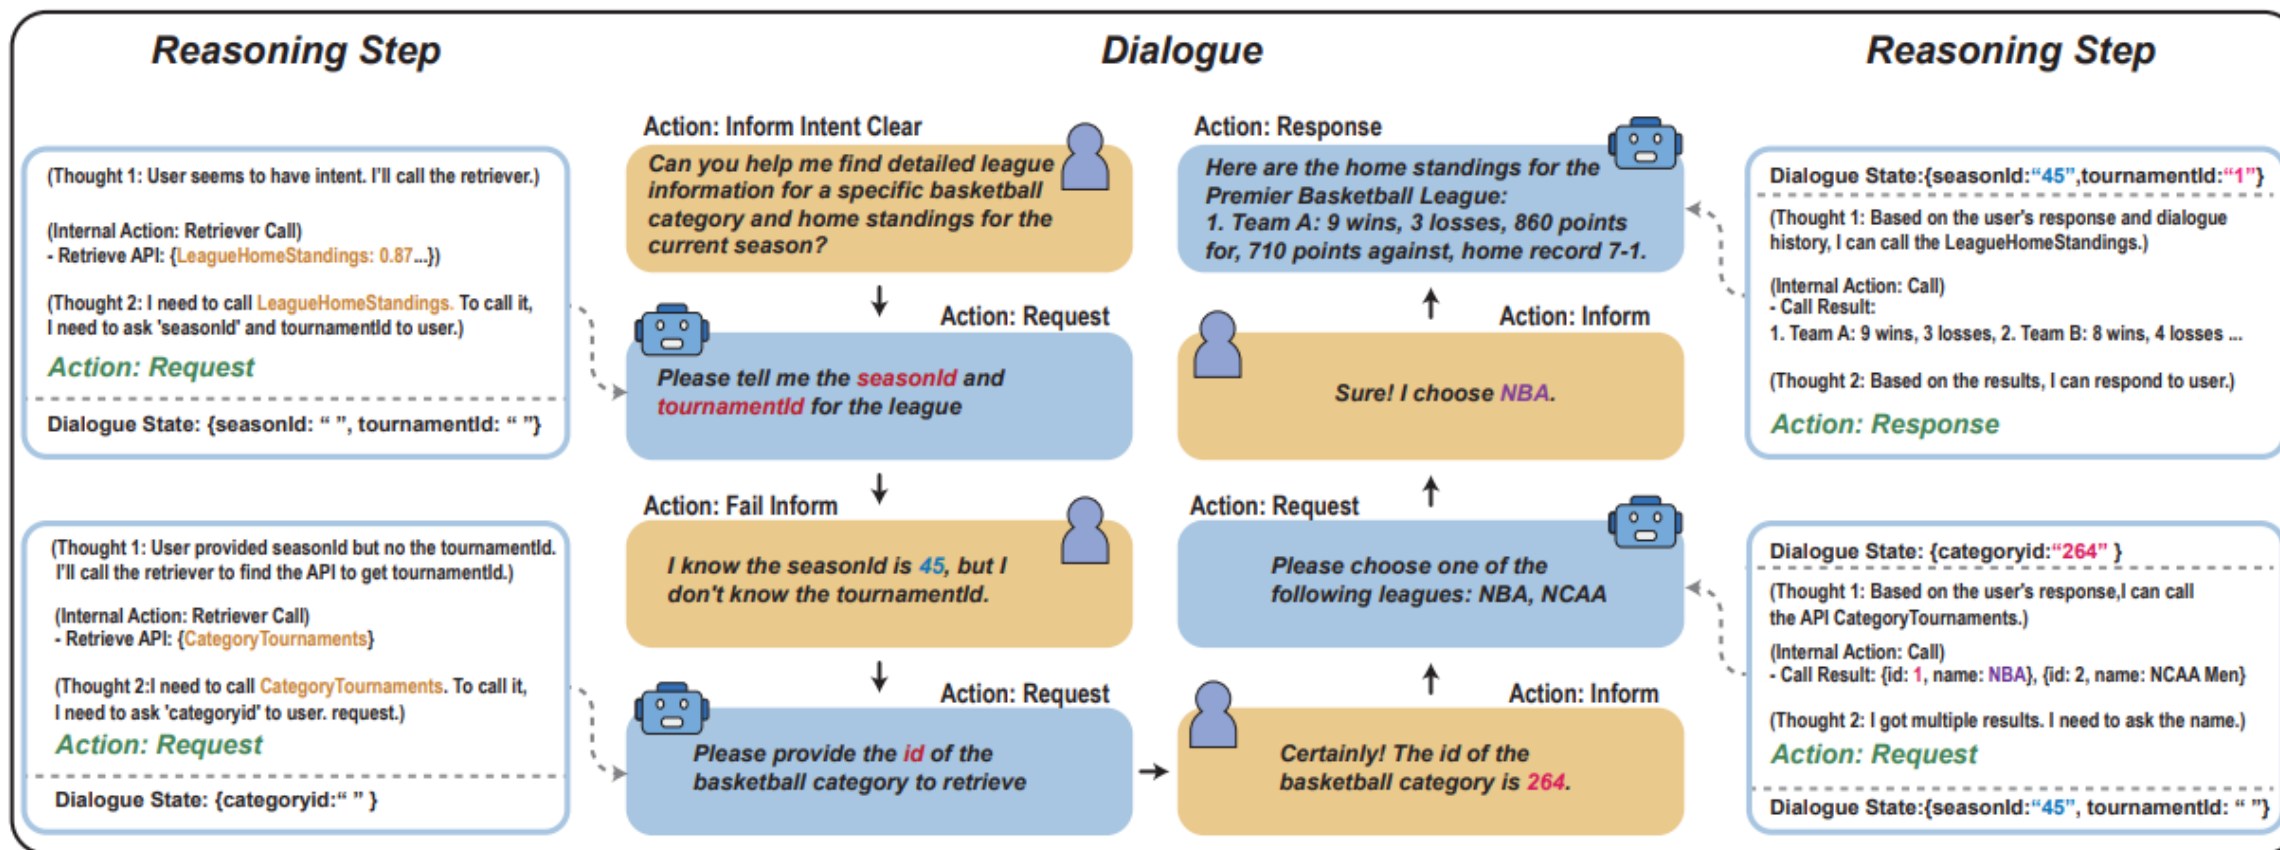


Directly prompt to ChatGPT

Method



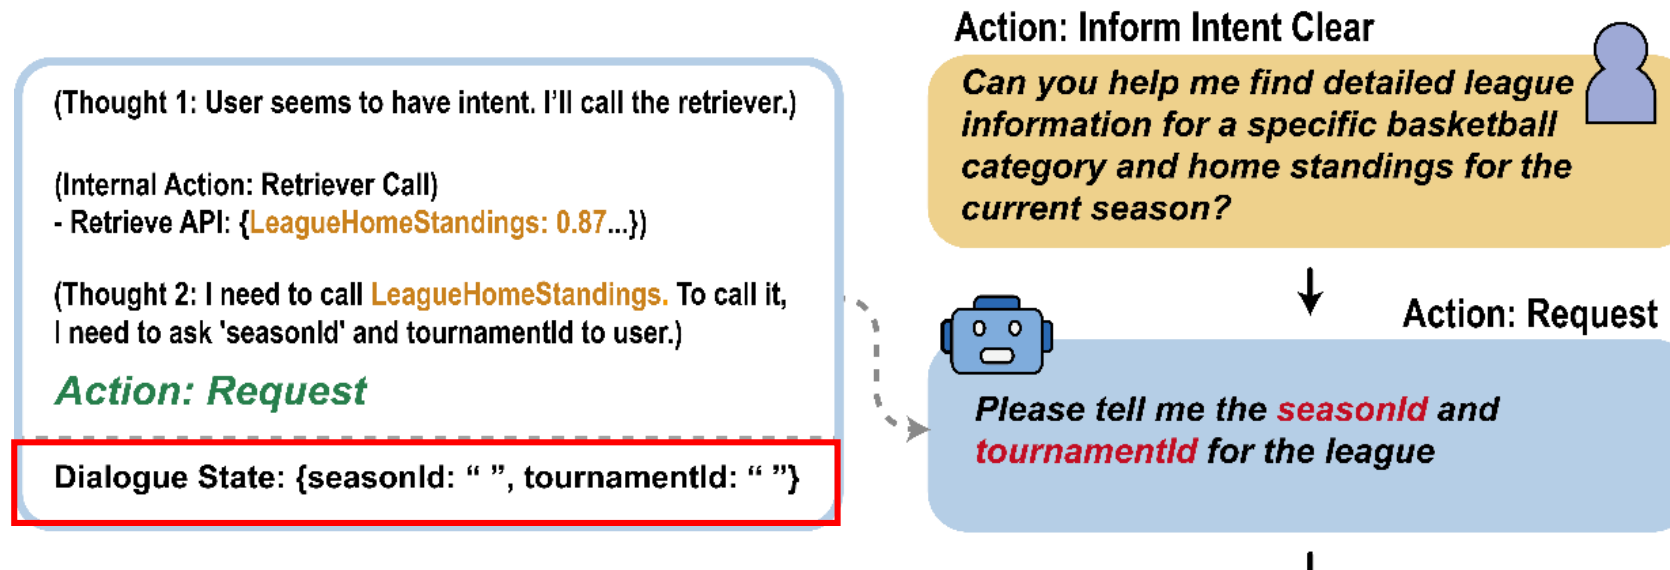
● Dialogue Data Generation



Experiments



● Dialogue state tracking



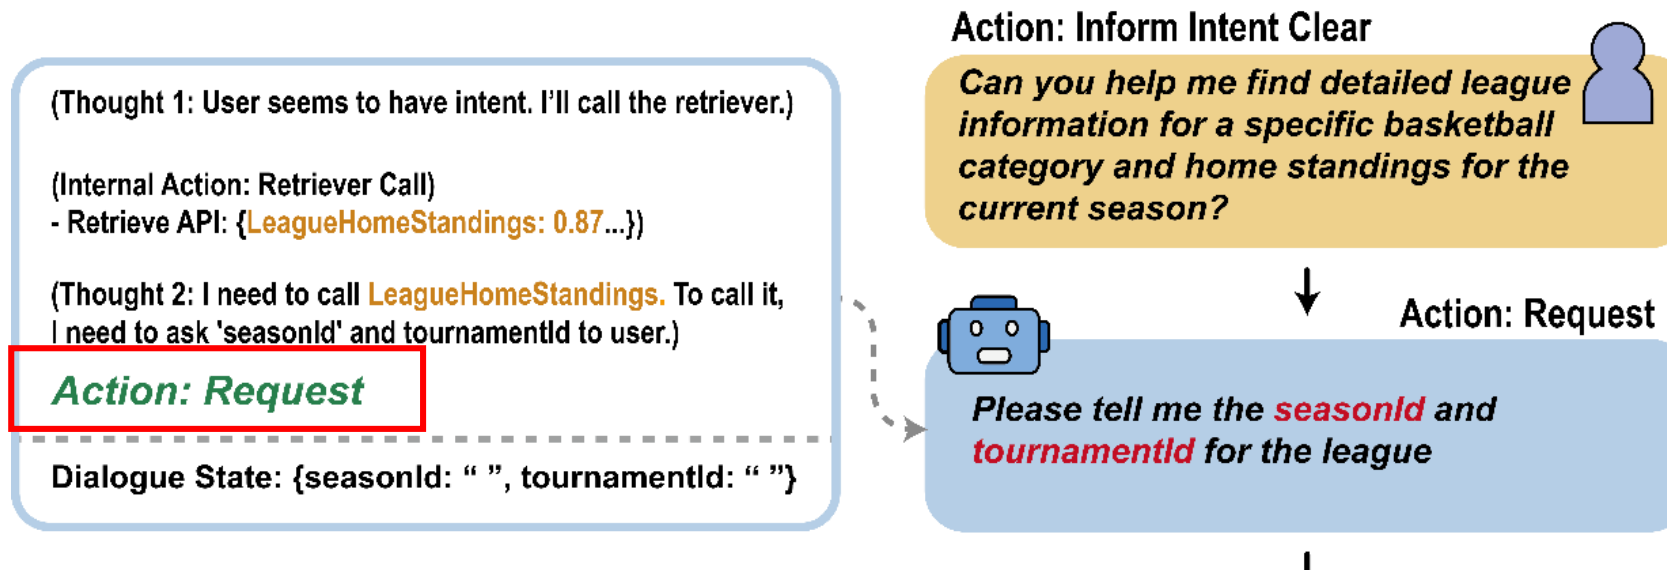
- A task that tracks the input parameter information of the API currently being executed from the dialogue history.
- Specific instruction is given on both train and inference

Experiments



Seoul National University
Graduate School of Data Science

● Action Prediction

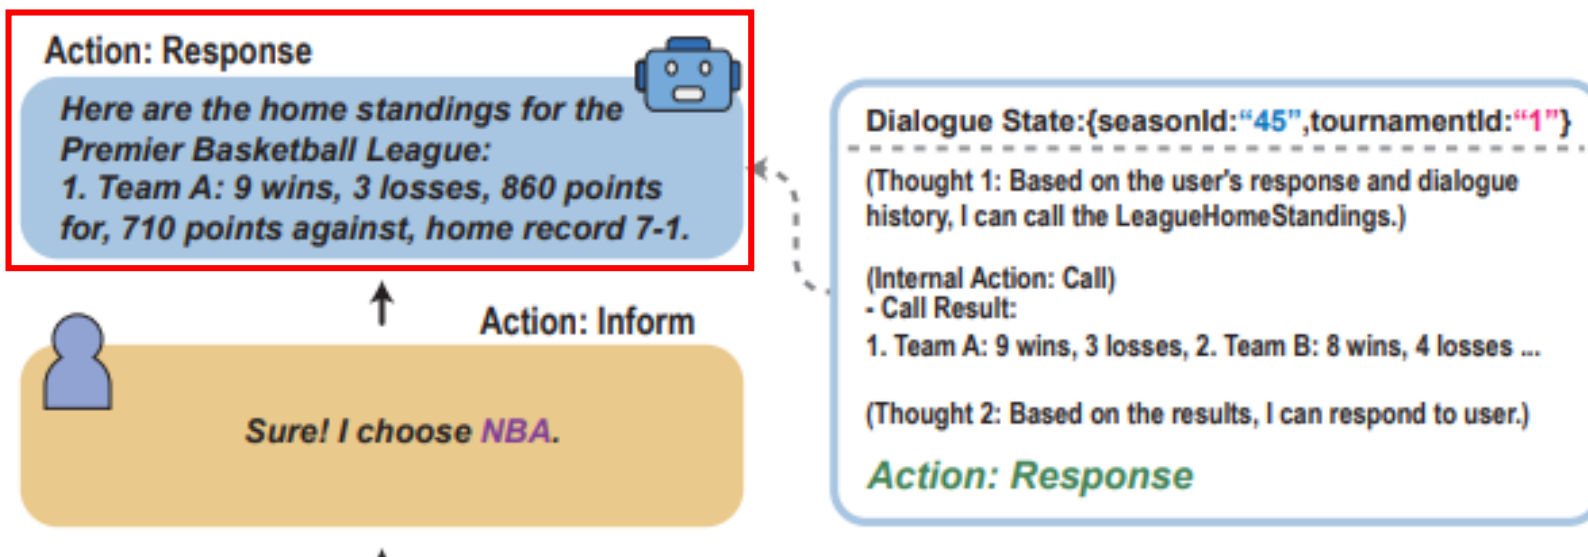


- Predict the next action that TALM will take based on the dialogue history.
- Specific instruction is given on both train and inference

Experiments



● Faithfulness



- We evaluate whether the system's response aligns with the API call results.
- The evaluation is performed using G-Eval [1].

Experiments



Seoul National University
Graduate School of Data Science

● Results

Dialogue State Tracking

Model	Dialogue State Tracking		Action Prediction		Faithfulness
	w GT	w/o GT	w GT	w/o GT	w/o GT
GPT-3.5-turbo	38.8	33.1	53.5	54.1	95.4
GPT-4o-mini	58.8	67.7	63.7	60.2	96.6
GPT-4-turbo	77.5	68.6	64.2	61.5	97.1
GPT-4o	81.4	67.8	57.6	63.7	96.7
CodeLlama-7b-Instruct-hf	47.2	28.9	35.7	30.0	81.7
Qwen2.5-Coder-7B-Instruct	48.9	34.2	55.8	46.8	93.9
Llama3-8B-Instruct	53.4	24.5	37.7	35.5	91.5
TD-Llama	92.7	72.2	77.5	91.0	88.4

Experiments



Seoul National University
Graduate School of Data Science

● Results

Action Prediction

Model	Dialogue State Tracking		Action Prediction		Faithfulness
	w GT	w/o GT	w GT	w/o GT	w/o GT
GPT-3.5-turbo	38.8	33.1	53.5	54.1	95.4
GPT-4o-mini	58.8	67.7	63.7	60.2	96.6
GPT-4-turbo	77.5	68.6	64.2	61.5	97.1
GPT-4o	81.4	67.8	57.6	63.7	96.7
CodeLlama-7b-Instruct-hf	47.2	28.9	35.7	30.0	81.7
Qwen2.5-Coder-7B-Instruct	48.9	34.2	55.8	46.8	93.9
Llama3-8B-Instruct	53.4	24.5	37.7	35.5	91.5
TD-Llama	92.7	72.2	77.5	91.0	88.4

Experiments



Seoul National University
Graduate School of Data Science

● Results

Faithfulness

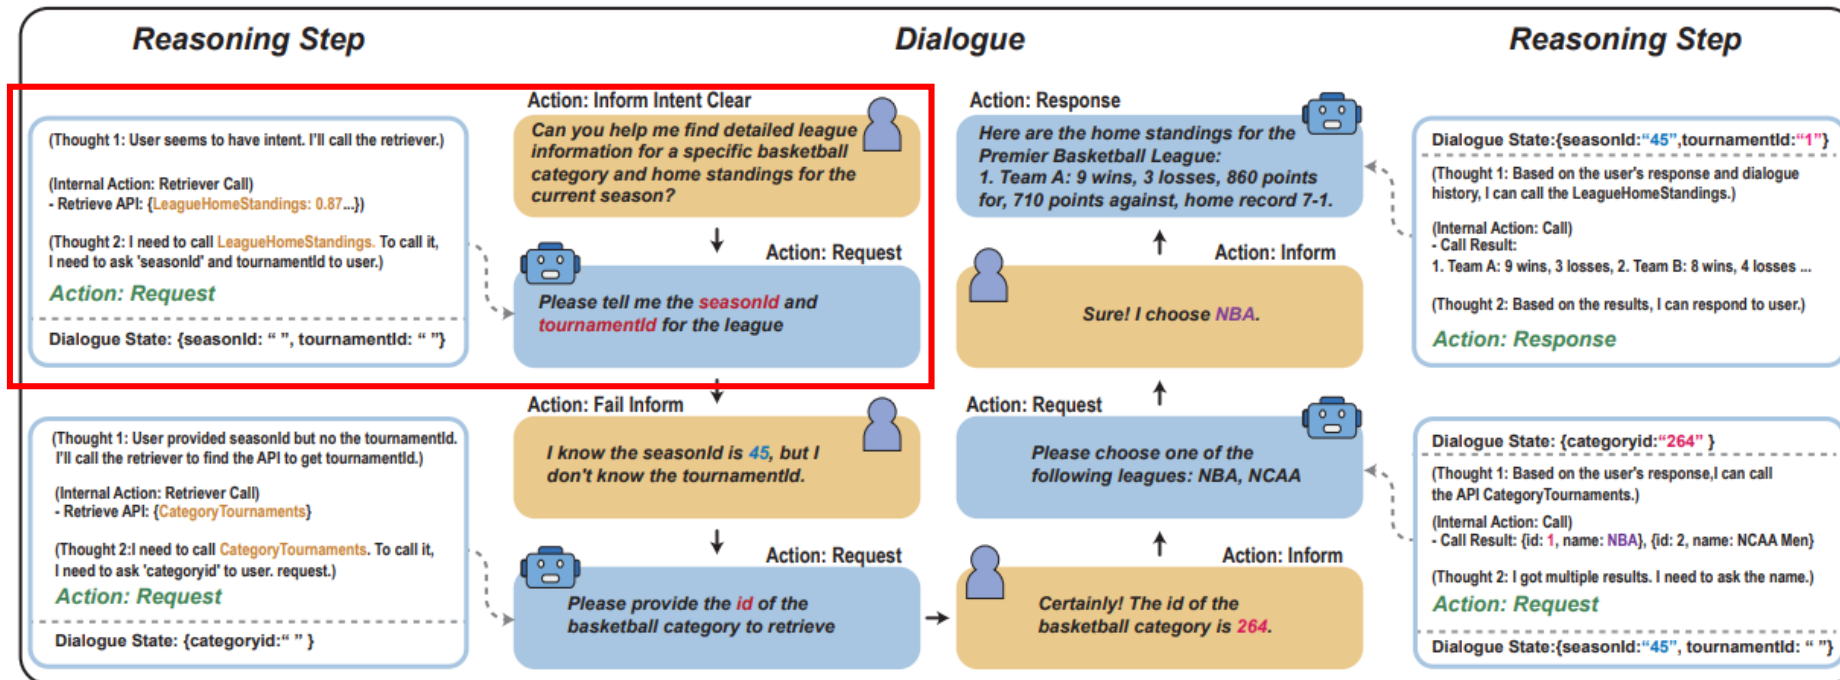
Model	Dialogue State Tracking		Action Prediction		Faithfulness
	w GT	w/o GT	w GT	w/o GT	w/o GT
GPT-3.5-turbo	38.8	33.1	53.5	54.1	95.4
GPT-4o-mini	58.8	67.7	63.7	60.2	96.6
GPT-4-turbo	77.5	68.6	64.2	61.5	97.1
GPT-4o	81.4	67.8	57.6	63.7	96.7
CodeLlama-7b-Instruct-hf	47.2	28.9	35.7	30.0	81.7
Qwen2.5-Coder-7B-Instruct	48.9	34.2	55.8	46.8	93.9
Llama3-8B-Instruct	53.4	24.5	37.7	35.5	91.5
TD-Llama	92.7	72.2	77.5	91.0	88.4

Experiments



Seoul National University
Graduate School of Data Science

● Overall Performance



- Evaluation of the ability to generate the next utterance after receiving the user's input, through appropriate reasoning steps.
- Task: Generating proper action and dialogue state for each reasoning step

- Overall Performance

- 77.1 % accuracy for 5,213 turn
- An accuracy of 28.3% for the 1,166 dialogues in the test set, indicating the proportion of dialogues correctly generated from start to finish.