

Applied Deep Learning



Language Agents

December 4th, 2024

<http://adl.miulab.tw>



[Reference: EMNLP 2024 Tutorial about Language Agents](#)



**National
Taiwan
University**
國立臺灣大學

Debates about Agents



Agents are bringing about the **biggest revolution in computing** since we went from typing commands to tapping on icons. – Bill Gates

I think AI agentic workflows will drive **massive AI progress** this year. – Andrew Ng

2025 is when **agents will work**. – Sam Altman

Current agents are just **thin wrappers around LLMs**.

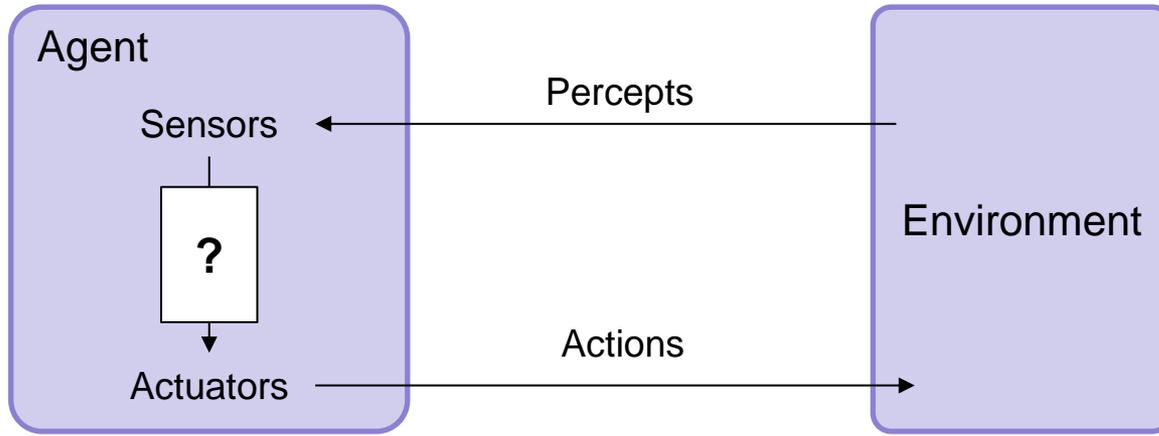
Autoregressive LLMs **can never reason or plan**.

Auto-GPT's limitations in ... reveal that it is **far from being a practical solution**.



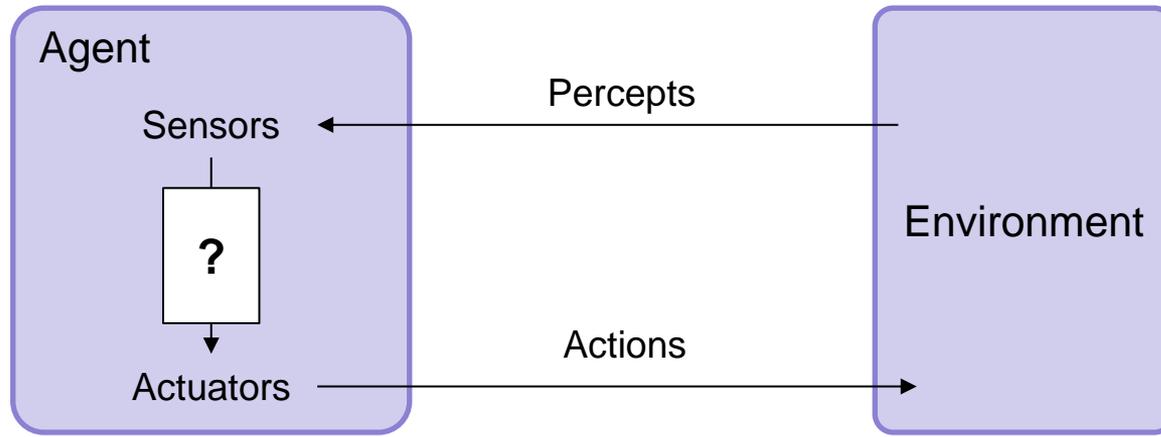
General Agent Definition

- “An agent is anything that can be viewed as **perceiving its environment** through sensors and **acting upon that environment** through actuators.” – *Russell & Norvig, AI: A Modern Approach (2020)*



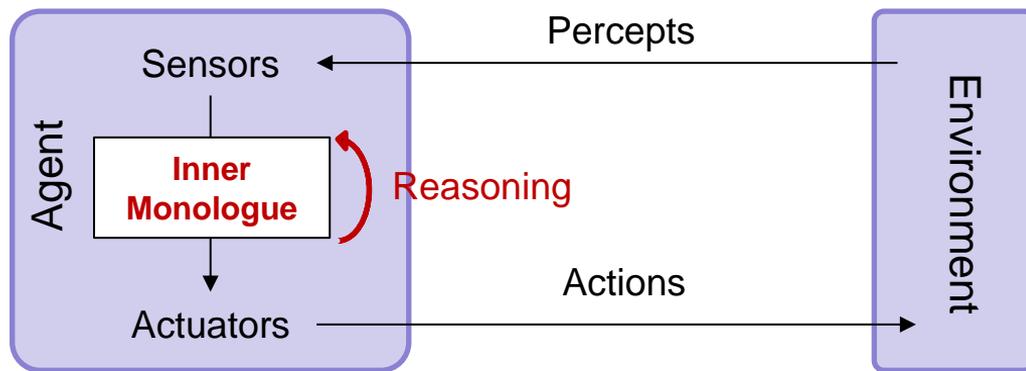
General Agent Definition

- An agent is an entity that **perceives** and **acts**.
- A rational agent **selects actions** that maximize its (expected) utility.
- Characteristics of the percepts, environment, and action space dictate techniques for selecting rational actions.



Idea: use **language** for reasoning and communication

Language Agents



- Reasoning by generating tokens can be viewed as an internal action
- Internal environment, where reasoning takes place in an inner monologue manner
- Self-reflection is a “meta” reasoning action (i.e., reasoning over the reasoning process)
- Reasoning is for better acting, by inferring environmental states, replanning, etc.
- Percepts and external action spaces are represented via languages

AI agents integrated with LLMs can *use language as a vehicle* for acting
→ Language Agents

AI Agent Evolution

	Logical Agent	Neural Agent	Language Agent
Expressiveness			
Reasoning			
Adaptivity			

7

Key Concepts for (Language) Agents

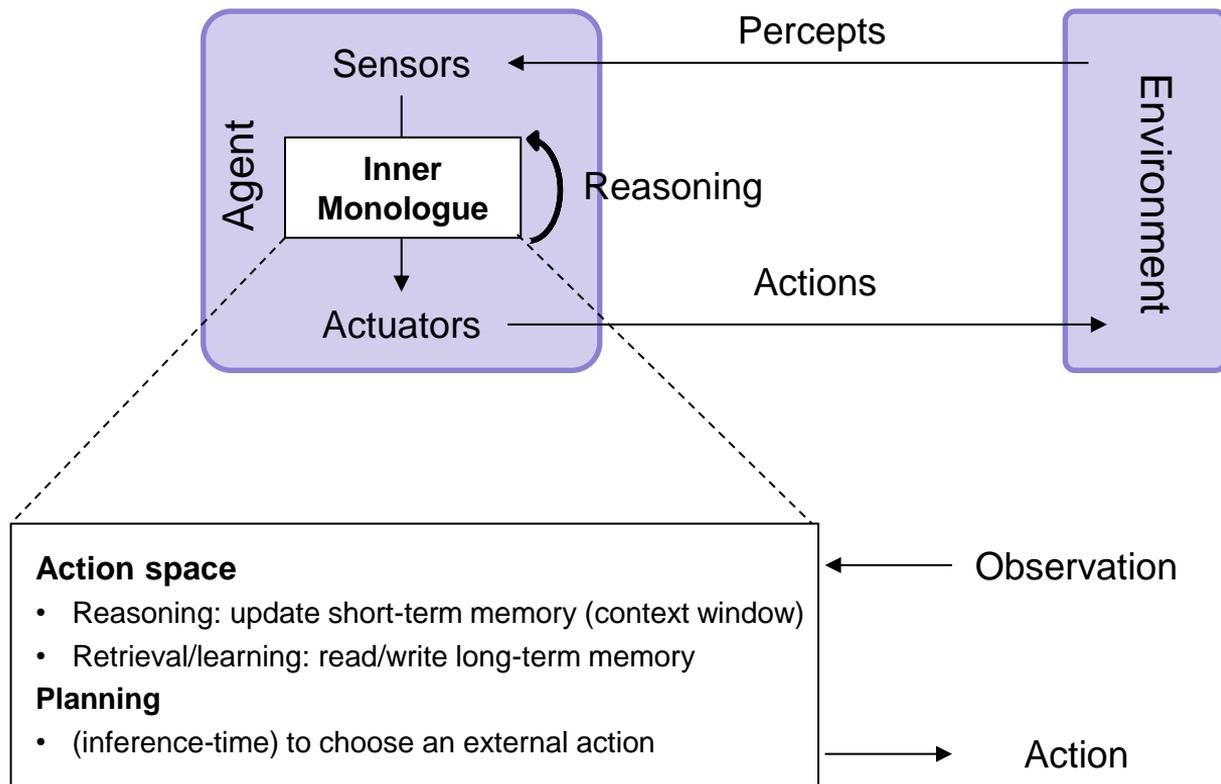
- Reasoning
- Memory
- Planning

8

Reasoning

One of three key concepts for language agents

Language Agents



Chain-of-Thought (CoT) (Wei et al., 2022)

Standard Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. ❌

Chain-of-Thought Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✅

Intermediate generation is to imitate *human mental processes*

Reasoning is Helpful for Agents

Obs t

You are cooking a dish and seeing salt is out...

Act t

Turn right

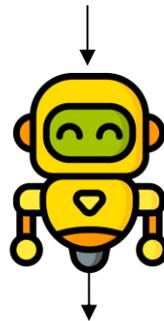
Obs $t+1$

You see a cabinet and a table...

Act $t+1$

Open cabinet

Observation



Action

Acting is Helpful for Reasoning

Action



Web search ▶

根據搜尋結果,陳縉儂是國立臺灣大學資訊工程學系的教授。她出生於台灣台北,在卡內基美隆大學獲得博士學位,同時也擁有該校的碩士學位和國立臺灣大學資訊工程學系的學士及碩士學位。她的研究領域主要為語言理解、機器智慧、口語對話系統、深度學習及自然語言處理等相關領域。

陳縉儂曾榮獲多項獎項,如IEEE ASRU 2013及IEEE SLT 2010的最佳學生論文獎、ISCA INTERSPEECH 2012之最佳學生論文提名以及ACL CLP之傑出碩士論文獎。在美國攻讀博士時,她決定有機會就要回台教書。畢業後,她先在美國微軟做博士後研究工作,年薪約450萬新台幣;成為微軟正職研究員後年薪更達近千萬新台幣。

然而陳縉儂選擇回到母校任教助理教授一職。她花費大量時間準備課程並指導第一屆帶領的研究生進入麻省理工學院攻讀博士。除了教書和研究外,

陳縉儂也熱愛舞蹈如嘻哈舞蹈等。

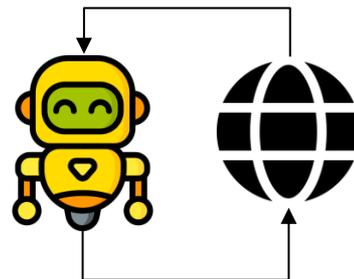
總而言之, 陳縉儂是一位優秀且充滿熱忱的人工智慧專家與教育者。

Sources: blog.104.com.tw facebook.com csie.ntu.edu.tw

你知道台大的陳縉儂嗎?



Observation



Action

ReAct (Yao et al., 2022)

Question: Aside from the Apple Remote, what other device can control the program Apple Remote was originally designed to interact with?

Standard

Answer: iPod

X

Reason Only

Thought: Let's think step by step. Apple Remote was originally designed to interact with Apple TV. Apple TV can be controlled by iPhone, iPad, and iPod Touch. So the answer is iPhone, iPad, and iPod Touch.

Answer: iPhone, iPad, iPod Touch

X

ReAct (Yao et al., 2022)

Question: Aside from the Apple Remote, what other device can control the program Apple Remote was originally designed to interact with?

Act Only

Act 1: `Search`[Apple Remote]

Obs 1: The Apple Remote is a remote control introduced in October 2005 by Apple ... originally designed to control the Front Row media center program ...

Act 2: `Search`[Front Row]

Obs 2: Could not find [Front Row]. Similar: ['Front Row Seat to Earth', 'Front Row Motorsports', 'Front Row (software)', ...]

Act 3: `Search`[Front Row (software)]

Obs 3: Front Row is a discontinued media center software ...

Act 4: `Finish`[yes]



ReAct (Yao et al., 2022)

Question: Aside from the Apple Remote, what other device can control the program Apple Remote was originally designed to interact with?

ReAct (Reason + Act)

Reasoning + Act are both essential

ReAct (Yao et al., 2022)

AlfWorld

You are in the middle of a room. Looking quickly around you, you see a armchair 1, a cabinet 4, ... a cabinet 1, a drawer 5, ..., a drawer 1, a dresser 1, a garbagecan 1, a safe 1, a shelf 12, ..., a shelf 1, a sidetable 1, and a sofa 1.

Your task is to: put two keychain in safe.

(a) ReAct

.....

Obs 16: You open the drawer 4. The drawer 4 is open. In it, you see a keychain 3, and a watch 2.

Act 17: **Think** [Now I find the first keychain (3). Next, I need to take it. I can find the second keychain (2) later in drawer 4.]

Act 18: take keychain 3 from drawer 4

Obs 18: You pick up the keychain 3 from the drawer 4.

.....

Act 22: put keychain 3 in/on safe 1

Obs 22: You put the keychain 3 in/on the safe 1.

Act 23: **Think** [Now I put the first keychain in safe. Next, I need to find the second keychain. I can directly go to drawer 4.]

Act 24: go to drawer 4

Obs 24: The drawer 4 is open. In it, you see a watch 2.

Act 25: **Think** [Now I find the second keychain (2). Next, I need to take it.]

Act 26: take keychain 2 from drawer 4

Obs 26: Nothing happens.



ReAct (Yao et al., 2022)

AlfWorld

You are in the middle of a room. Looking quickly around you, you see a armchair 1, a cabinet 4, ... a cabinet 1, a drawer 5, ..., a drawer 1, a dresser 1, a garbagecan 1, a safe 1, a shelf 12, ..., a shelf 1, a sidetable 1, and a sofa 1.

Your task is to: **put two keychain in safe.**

(b) ReAct (Human Thought edit)

.....

Obs 16: You open the drawer 4. The drawer 4 is open. In it, you see a keychain 3, and a watch 2.

Act 17: **Think** [Now I find the first keychain (3). Next, I need to take it. ~~I can find the second keychain (2) later in drawer 4.~~]

Act 18: take keychain 3 from drawer 4

Obs 18: You pick up the keychain 3 from the drawer 4.

.....

Act 22: put keychain 3 in/on safe 1

Obs 22: You put the keychain 3 in/on the safe 1.

Act 23: **Think** [Now I put the first keychain in safe. Next, I need to find the second keychain, which is more likely to appear in dresser (1), garbagecan (1), safe (1), sidetable (1), sofa (1), shelf (1-12).]

Act 24: go to dresser 1

Obs 24: On the dresser 1, you see a box 1, a creditcard 1, a keychain 2, a keychain 1, a laptop 1...

Act 25: **Think** [Now I find the second keychain (2). Next, I need to take it.]

Act 26: take keychain 2 from dresser 1

Obs 26: **You pick up the keychain 2 from the dresser 1.**

Act 27: **Think** [Now I take the second keychain (2). Next, I need to put it in/on safe 1.]

Act 28: go to safe 1

Obs 28: The safe 1 is open. In it, you see a keychain 3, a statue 2, and a watch 3.

Act 29: put keychain 2 in/on safe 1

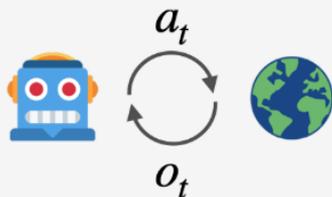
Obs 29: **You put the keychain 2 in/on the safe 1.**



Reasoning provides explanations for controlling actions

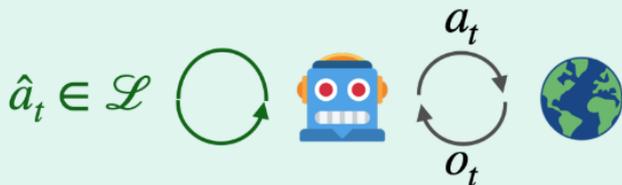
Reasoning Enlarges the Action Space

Traditional agents: action space A defined by the environment



- **External feedback** o_t
- Agent context $c_t = (o_1, a_1, o_2, a_2, \dots, o_t)$
- Agent action $a_t \sim \pi(a | c_t) \in A$

ReAct: action space $\hat{A} = A \cup \mathcal{L}$ augmented by reasoning

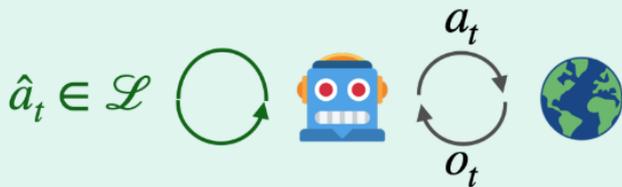


- $\hat{a}_t \in \mathcal{L}$ can be any language sequence
- Agent context $c_{t+1} = (c_t, \hat{a}_t, a_t, o_{t+1})$
- $\hat{a}_t \in \mathcal{L}$ only updates **internal context**

Reasoning Enlarges the Action Space

- Large action space \rightarrow more capacity, harder decision making
 - The space of reasoning/language is infinite
- LLMs learn reasoning priors by imitating various human reasoning traces

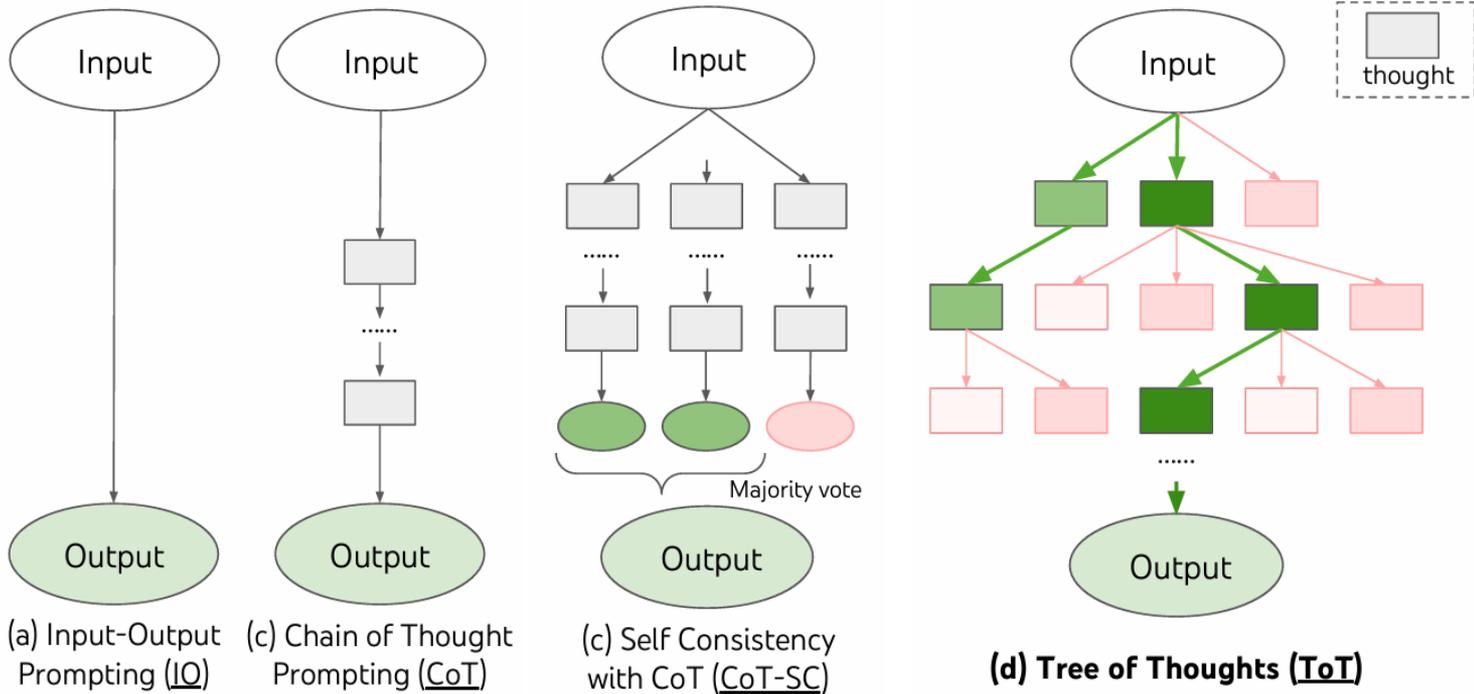
ReAct: action space $\hat{A} = A \cup \mathcal{L}$ augmented by reasoning



- $\hat{a}_t \in \mathcal{L}$ can be any language sequence
- Agent context $c_{t+1} = (c_t, \hat{a}_t, a_t, o_{t+1})$
- $\hat{a}_t \in \mathcal{L}$ only updates **internal context**

Action Planning for Improving Reasoning

(Yao et al, 2023)

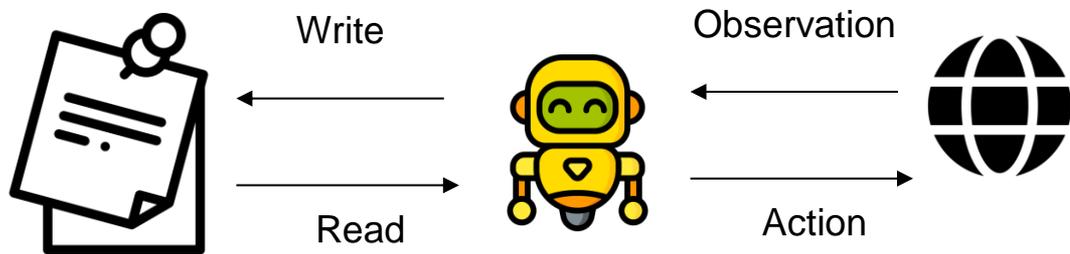


21

Memory

One of three key concepts for language agents

Short-Term and Long-Term Memory



Long-term memory

- Read & write
- Stores experience, knowledge, skills, ...
- Persist over new experience



Instruction: ...
Thought: ...
Action: ...
Obs: ...
Thought: ...

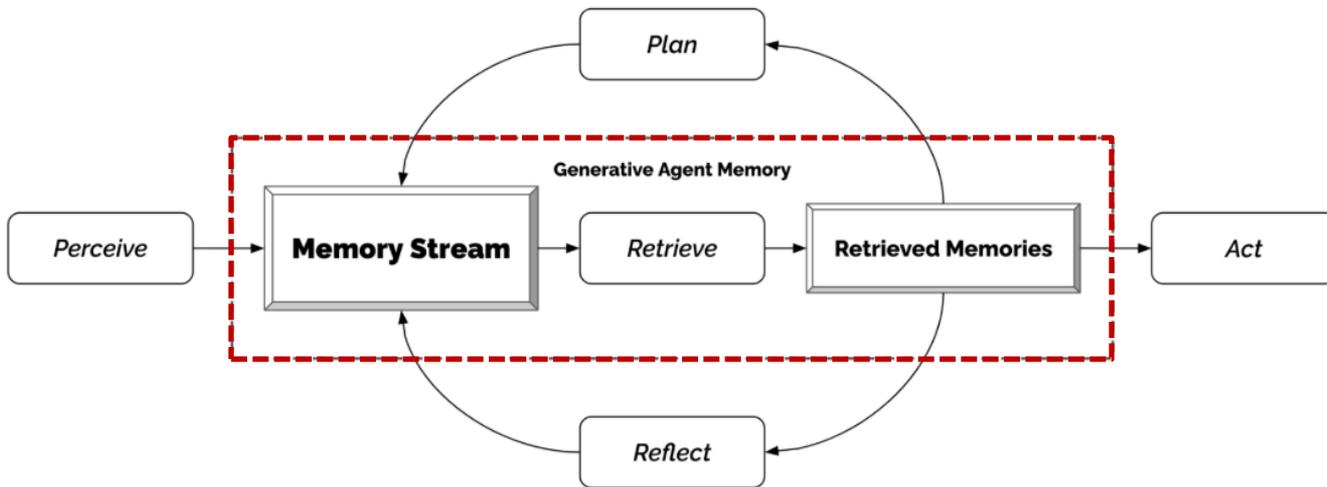
Short-term memory

- Append-only
- Limited context
- Limited attention
- Do not persist over new tasks

Generative Agents (Park et al., 2023)

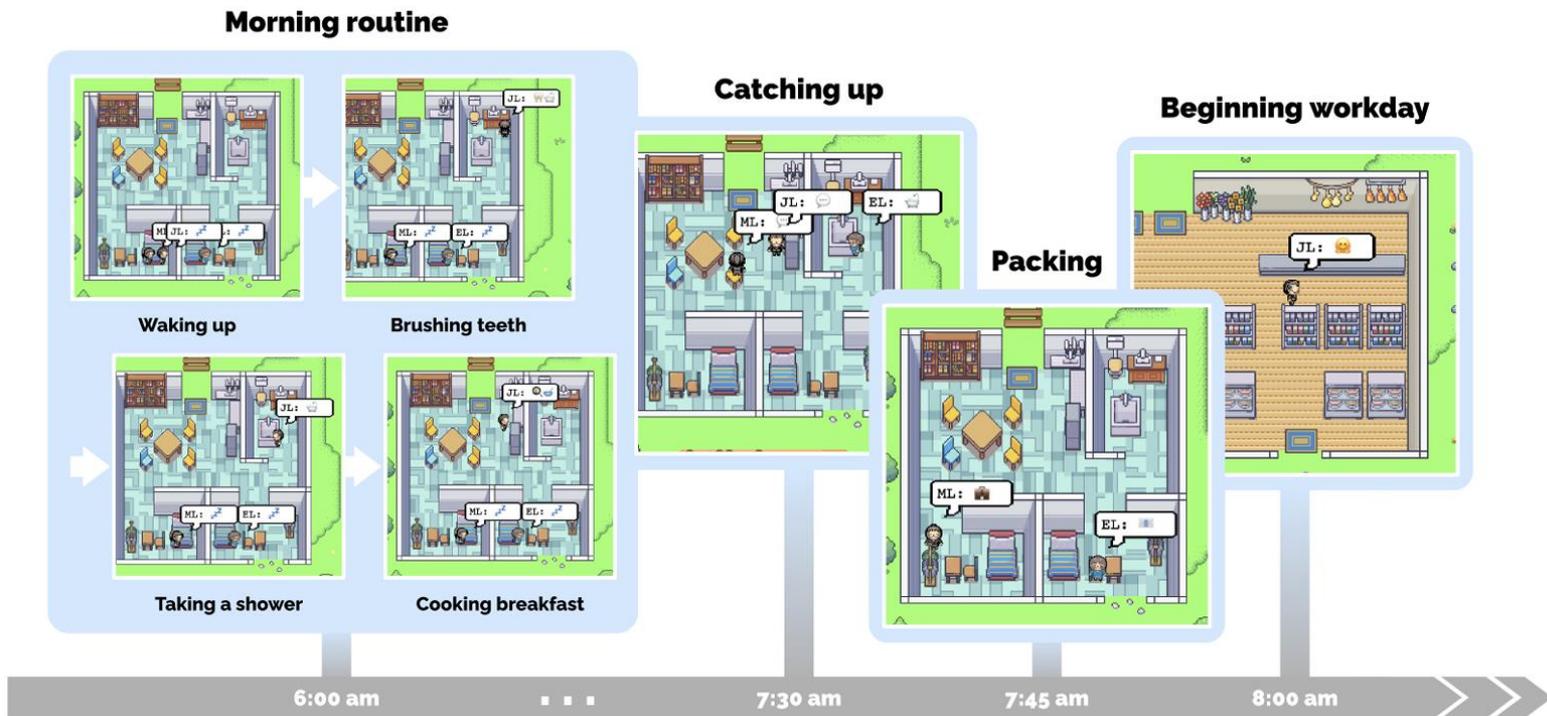
Memory requirement

- Context window cannot **hold** all the event streams
- It is difficult to **attend** to relevant events



Generative Agents (Park et al., 2023)

- Step 1: simulating a series of events for *episodic memory*



Generative Agents (Park et al., 2023)

- Step 2: **memory retrieval** considers *recency*, *importance*, and *relevance*

Memory Stream

```

2023-02-13 22:48:20: desk is idle
2023-02-13 22:48:20: bed is idle
2023-02-13 22:48:10: closet is idle
2023-02-13 22:48:10: refrigerator is idle
2023-02-13 22:48:10: Isabella Rodriguez is stretching
2023-02-13 22:33:30: shelf is idle
2023-02-13 22:33:30: desk is neat and organized
2023-02-13 22:33:10: Isabella Rodriguez is writing in her journal
2023-02-13 22:18:10: desk is idle
2023-02-13 22:18:10: Isabella Rodriguez is taking a break
2023-02-13 21:49:00: bed is idle
2023-02-13 21:48:50: Isabella Rodriguez is cleaning up the kitchen
2023-02-13 21:48:50: refrigerator is idle
2023-02-13 21:48:50: bed is being used
2023-02-13 21:48:10: shelf is idle
2023-02-13 21:48:10: Isabella Rodriguez is watching a movie
2023-02-13 21:19:10: shelf is organized and tidy
2023-02-13 21:18:10: desk is idle
2023-02-13 21:18:10: Isabella Rodriguez is reading a book
2023-02-13 21:03:40: bed is idle
2023-02-13 21:03:30: refrigerator is idle
2023-02-13 21:03:30: desk is in use with a laptop and some papers on it
...

```

Q. What are you looking forward to the most right now?

Isabella Rodriguez is excited to be planning a Valentine's Day party at Hobbs Cafe on February 14th from 5pm and is eager to invite everyone to attend the party.

retrieval	recency	importance	relevance
2.34	0.91	0.63	0.80

ordering decorations for the party

2.21	0.87	0.63	0.71
------	------	------	------

researching ideas for the party

2.20	0.85	0.73	0.62
------	------	------	------

...



I'm looking forward to the Valentine's Day party that I'm planning at Hobbs Cafe!



Not only relevance should be considered.

Social Simulation Agents (Zhang et al., 2024)

- Idea: different **moods** leads to different **responses**.



Hey, I passed the bar exam!!!

Context-emotion

I feel **Joyful, Impressed** because my friend passed the bar exam

I'm so glad to know that! Let's have a party to celebrate!!



Self-emotion

I feel **Proud, Excited** because I've been promoted.
Expressed emotion: Excited

Oh, I'm so happy for you and wish you best of luck.



Self-emotion

I feel **Disappointed, Jealous** because I also took the exam, but failed.
Expressed emotion: Disappointed

Social Simulation Agents (Zhang et al., 2024)

Self-emotion simulation through emotion-related events

Instruction:

Generate a series of events ...

Profile:

Sofia is a software engineer working at DiamondSoft. She graduate from ...



Start to the day
(coffee spill on the
desk).

7:00



Get an email about a
critical feature request
from the product team.

9:00



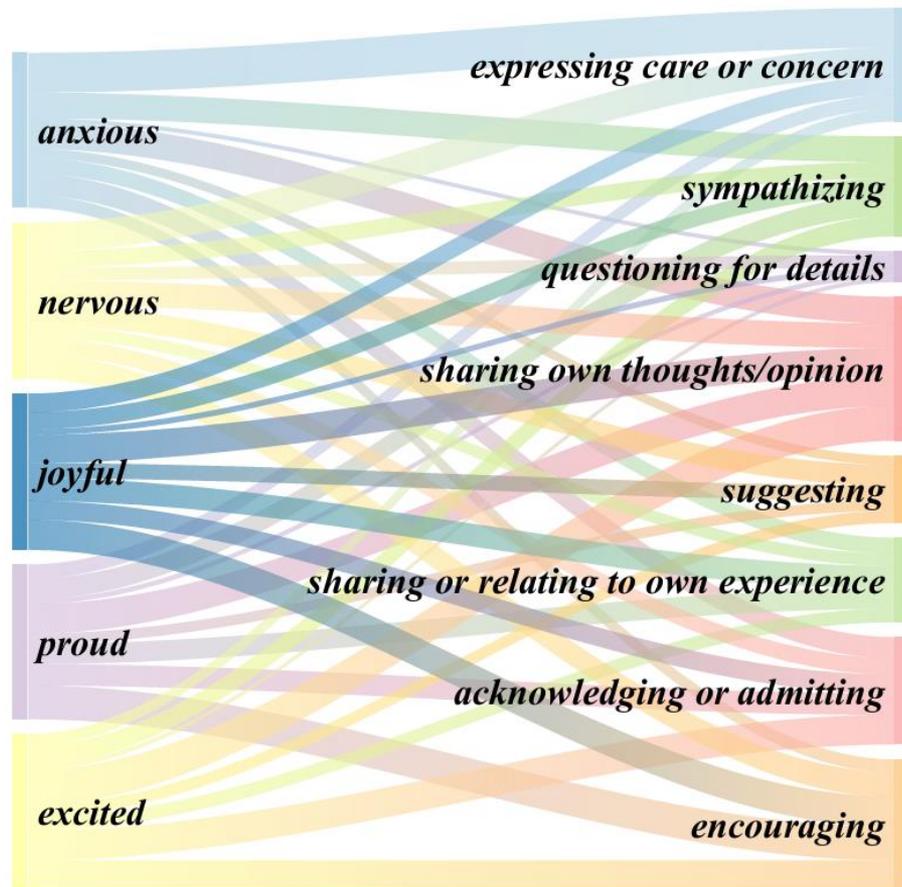
Successful debugging
a complex issue from
the previous day.

11:30



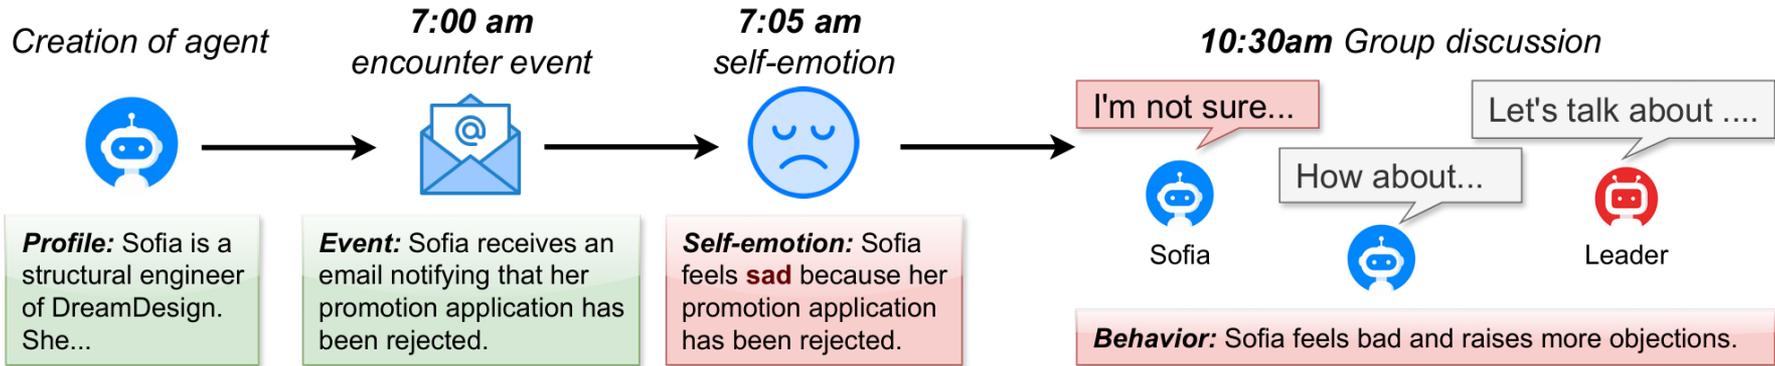
Social Simulation Agents (Zhang et al., 2024)

- Flow between most frequent *self-emotion* and the *dialogue strategies*



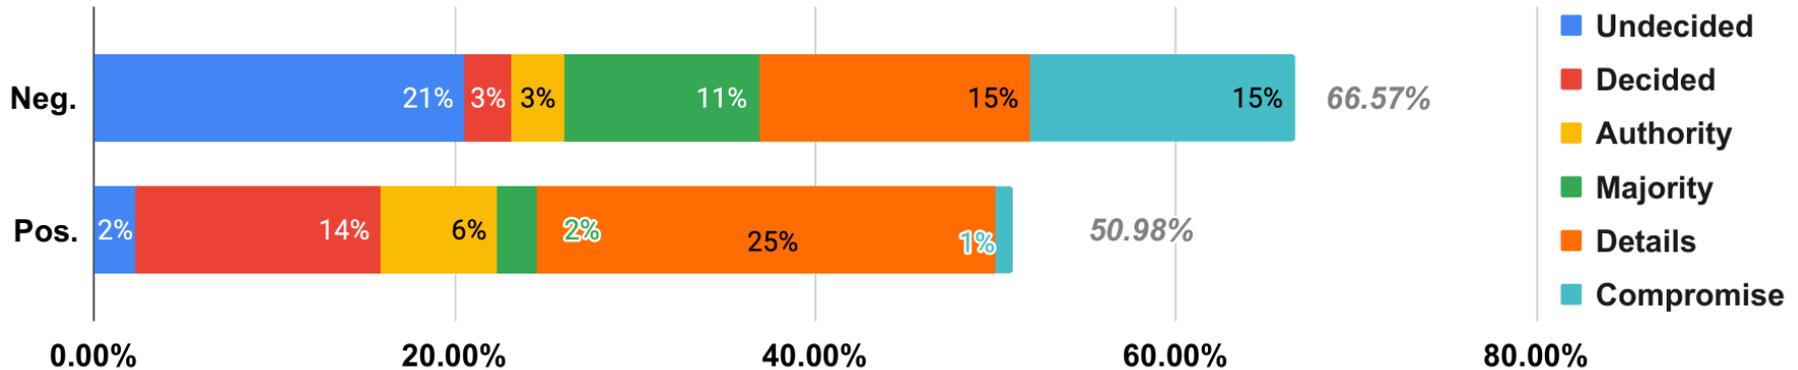
Social Simulation Agents (Zhang et al., 2024)

- Simulate group discussion in different timing / moods
 - Each agent has an *individual profile*
 - Agent's self-emotion *changes the decision*
 - Neg. → objections; Pos. → agreements



Social Simulation Agents (Zhang et al., 2024)

- Simulate group discussion in different timing / moods
 - Each agent has an *individual profile*
 - Agent's self-emotion *changes the decision*
 - Neg. → objections; Pos. → agreements



Positive self-emotion tends to lead more peaceful decisions in group discussions

31

Planning

One of three key concepts for language agents

Planning Definition

- Given a goal G , decides on a sequence of actions $(a_0, a_1, a_2, \dots, a_n)$ that will lead to a state that passes the goal test $g(\cdot)$.

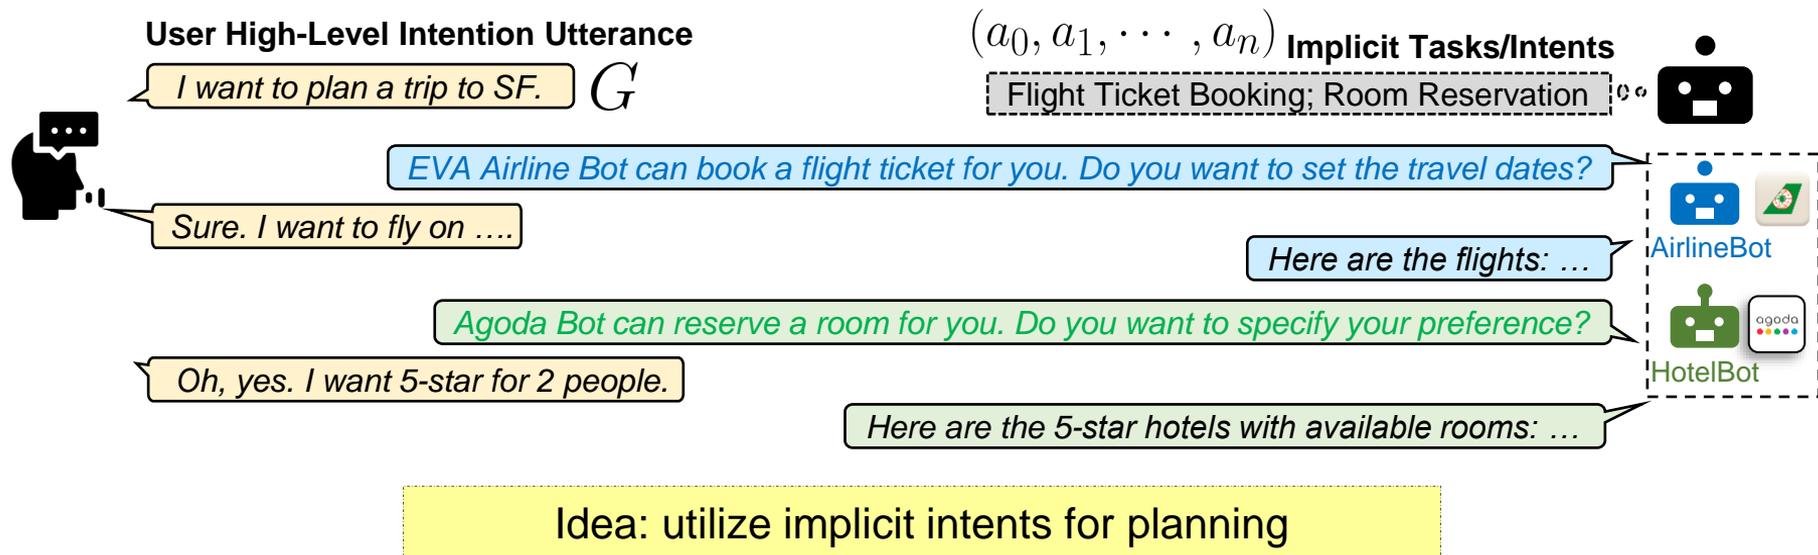
High-level goal

A sequence of actions

$$G \longrightarrow (a_0, a_1, \dots, a_n)$$

Commonsense-Inferred Planning (Kuo & Chen, 2023)

- Goal: users interact in high-level descriptions and the agent learns how to plan dialogues



Web Planning Agents (Deng et al., 2024)

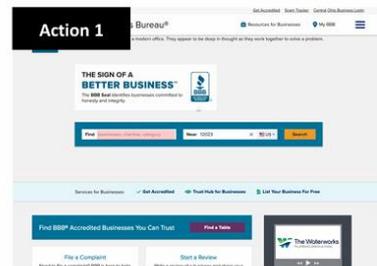
Task Description:

Show me the reviews for the auto repair business closest to 10002.

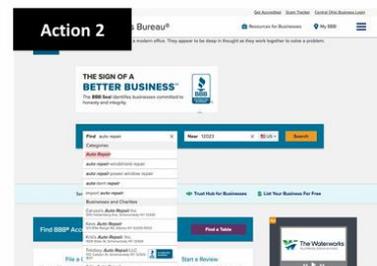
Action Sequence:

Target Element	Operation
1. [searchbox] Find	TYPE: auto repair
2. [button] Auto Repair	CLICK
3. [textbox] Near	TYPE: 10002
4. [button] 10002	CLICK
5. [button] Search	CLICK
6. [switch] Show BBB Accredited only	CLICK
7. [svg]	CLICK
8. [button] Sort By	CLICK
9. [link] Fast Lane 24 Hour Auto Repair	CLICK
10. [link] Read Reviews	CLICK

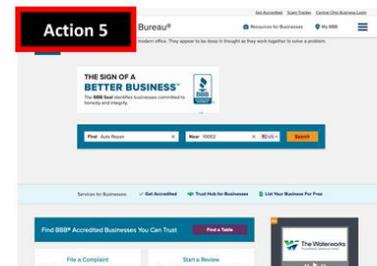
Webpage Snapshots:



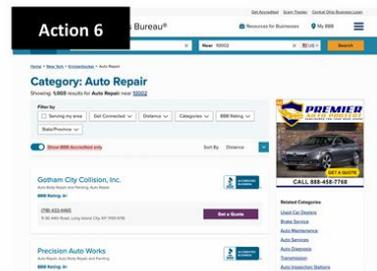
`<input name="find_text" type="search">`



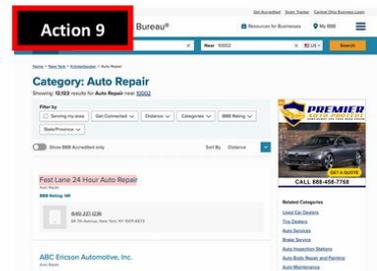
`Auto Repair`



`<button>Search</button>`



`<button>Show BBB Accredited only</button>`



`Fast Lane 24 Hour Auto Repair`

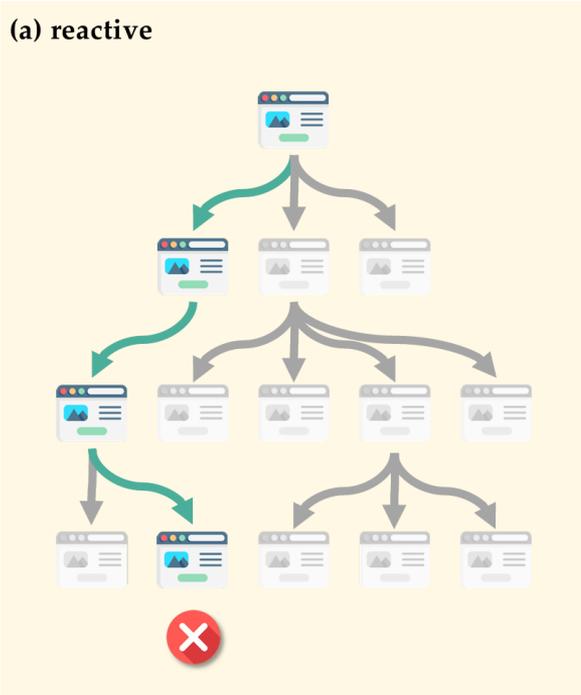


`Read Reviews`

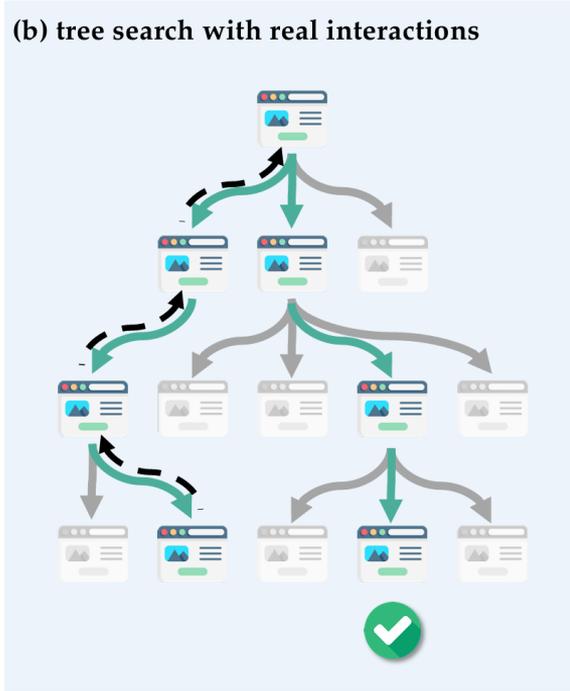
Idea: agent decomposes the task into several web actions

Planning Paradigms for Language Agents

(a) reactive



(b) tree search with real interactions



fast, easy to implement



greedy, short-sighted

Tree Search with Real Interactions (Koh et al, 2024)

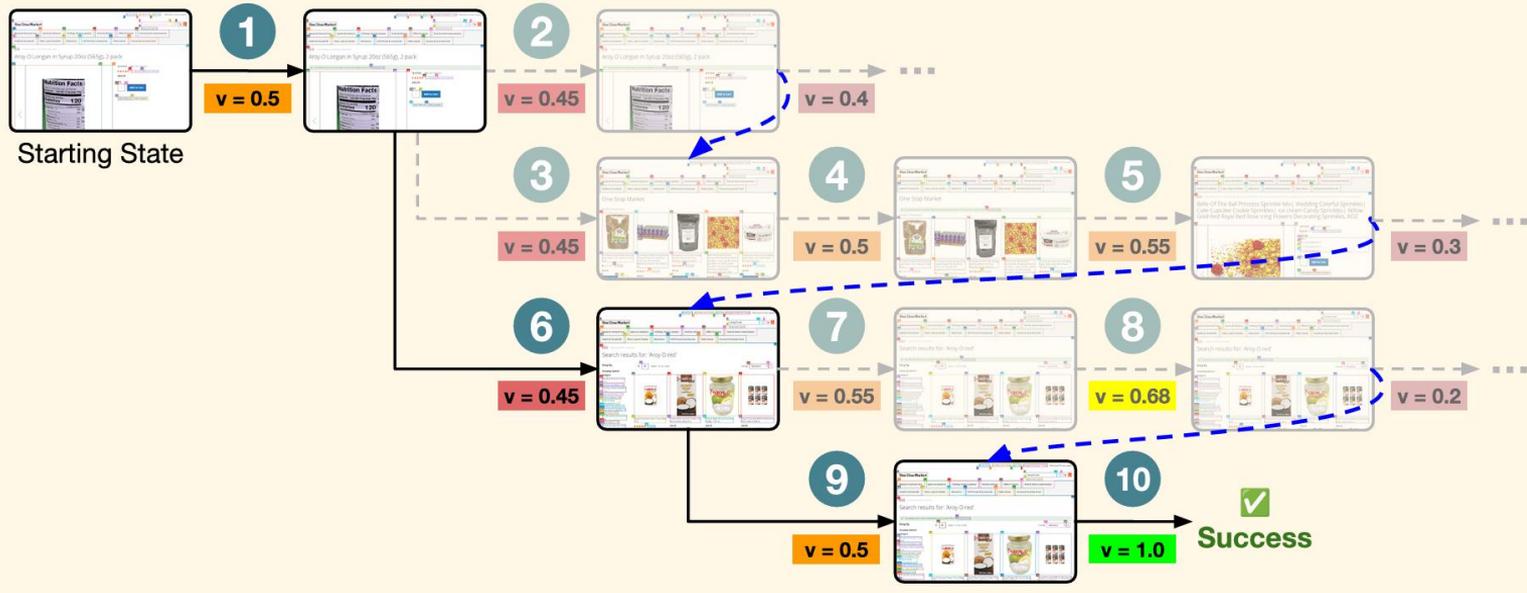


Task Instruction (*I*): “Can you add this and the other canned fruit (of the same brand) that looks like this, but red instead of brown to the comparison page?”

Legend

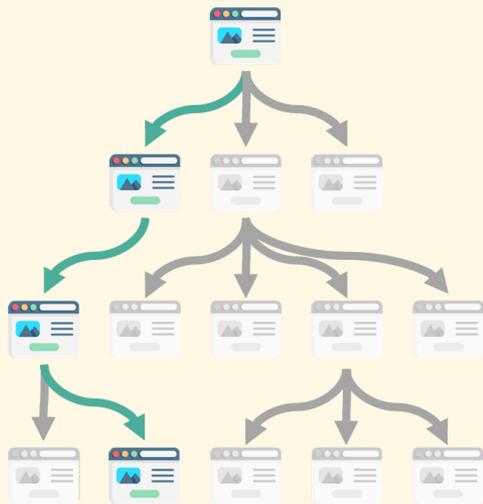
- 1 Step sequence
- State values (e.g., $v = 1.0$)
- Backtracking (dashed blue arrow)

GPT-4o Agent + Search



Planning Paradigms for Language Agents

(a) reactive

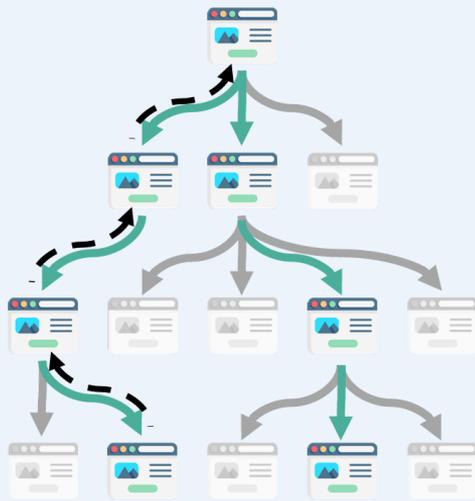


fast, easy to implement



greedy, short-sighted

(b) tree search with real interactions

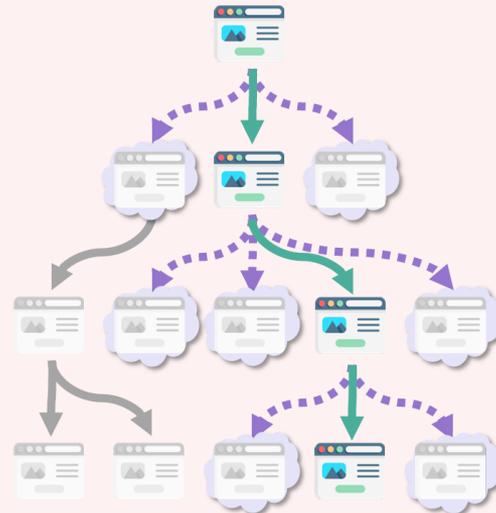


systematic exploration



irreversible actions,
unsafe, slow

(c) model-based planning



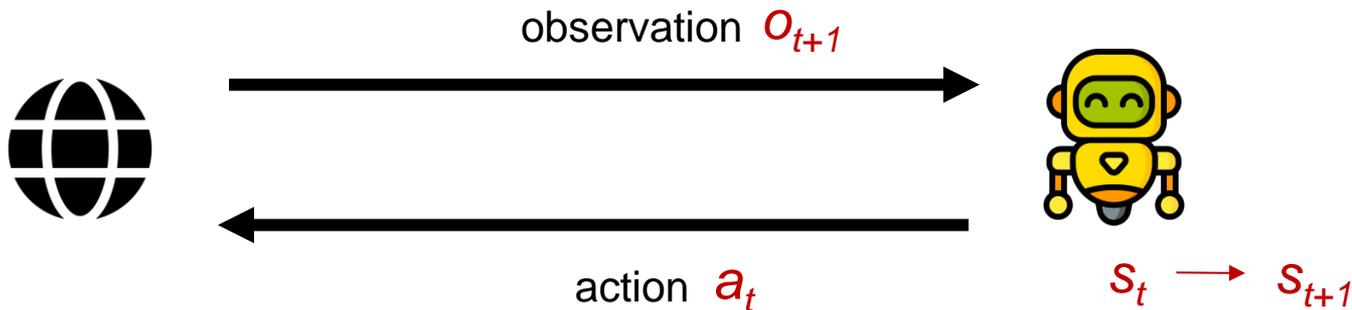
faster, safer, systematic
exploration



how to get a world model?

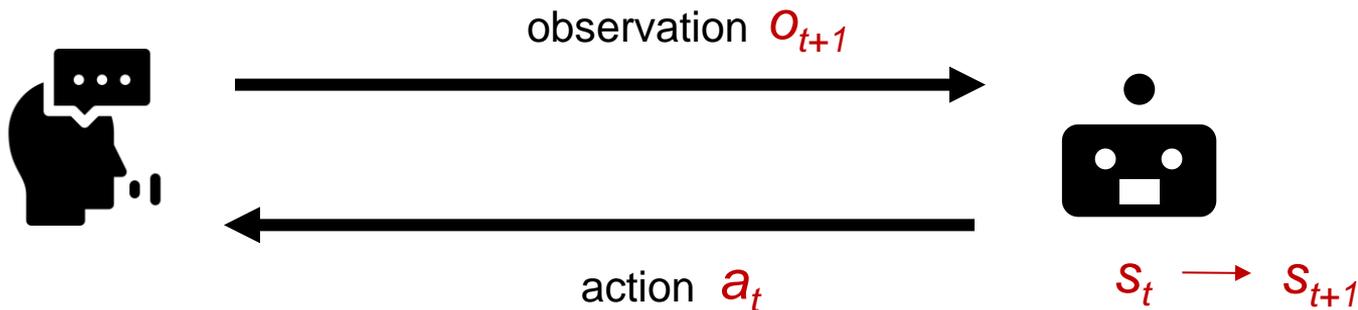
World Model is an Environment Simulator

- When doing the action a_t in the state s_t , what would happen next (s_{t+1})?



World Model in Conversation Assistants

- When doing the action a_t in the state s_t , what would happen next (s_{t+1})?

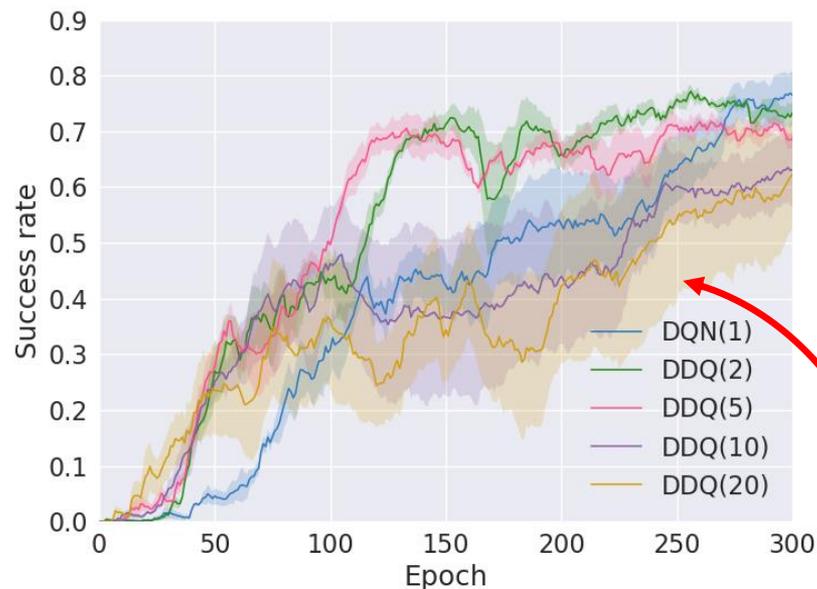
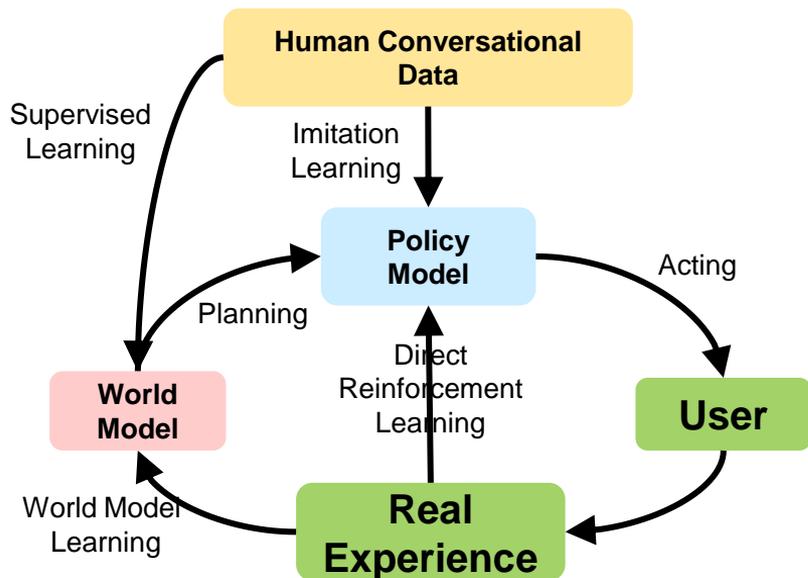


World model should synthesize the interactive experience

Conversation Planning – Deep Dyna-Q

(Peng+, 2018)

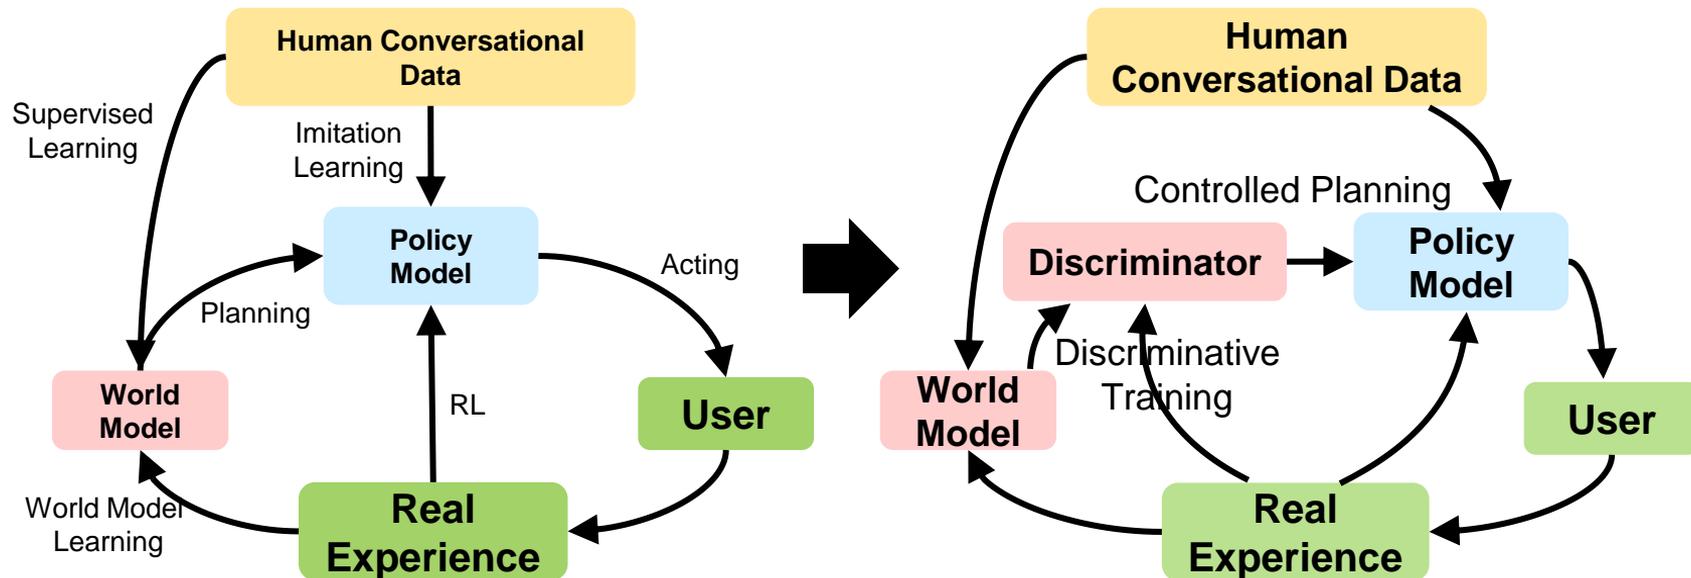
- Idea: learning with real users with planning



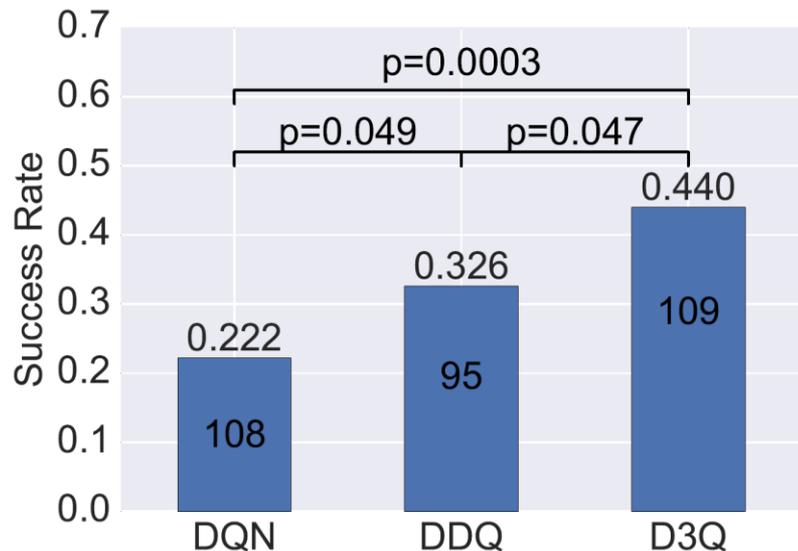
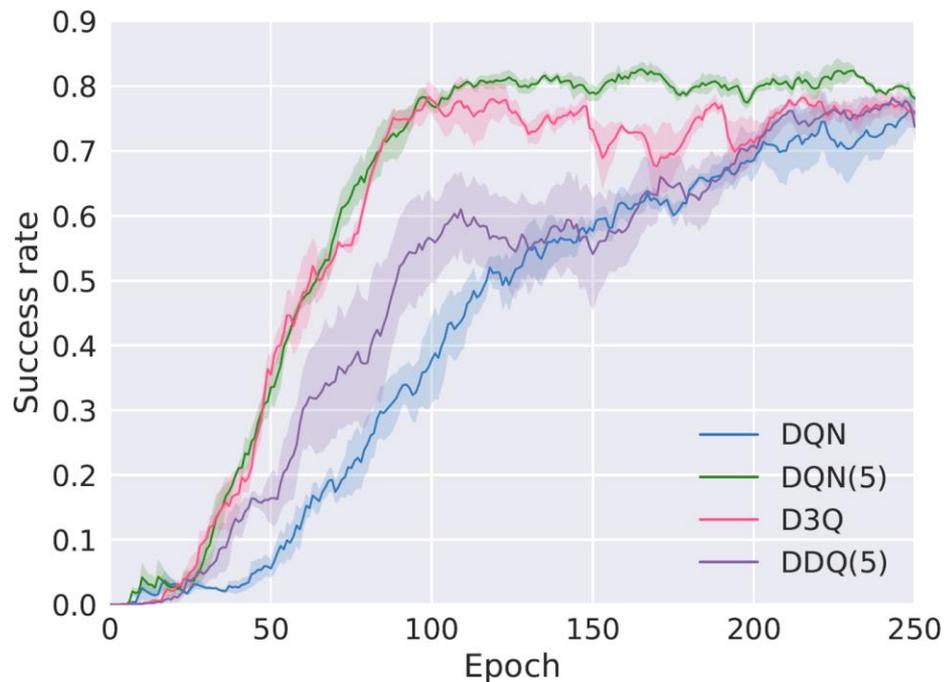
Policy learning suffers from the poor quality of fake experiences

Robust Planning – D3Q (Su+, 2018)

- Idea: add a *discriminator* to filter out the bad experiences



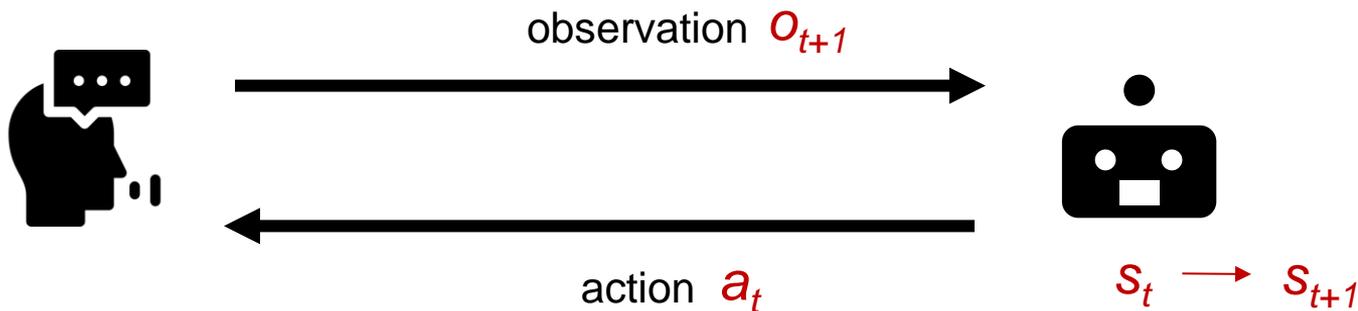
Robust Planning – D3Q (Su+, 2018)



The policy learning is *more robust* and shows the improvement in human evaluation

World Model in Conversation Assistants

- When doing the action a_t in the state s_t , what would happen next (s_{t+1})?



LLMs can directly serve as the user simulators

LLMs as User Simulators with Personas

○ Role play for an E-type user

You're Jamie Wright, a 29-year-old marketing manager who loves attending concerts and music festivals. **You gain energy from interacting with diverse groups of people and thrive in dynamic environments.** Known for your creativity and quick decision-making, you often take the lead in brainstorming sessions and enjoy collaborating with others. Your broad interests include digital marketing trends, social media, and networking events, **where you excel at making connections and sharing ideas.**

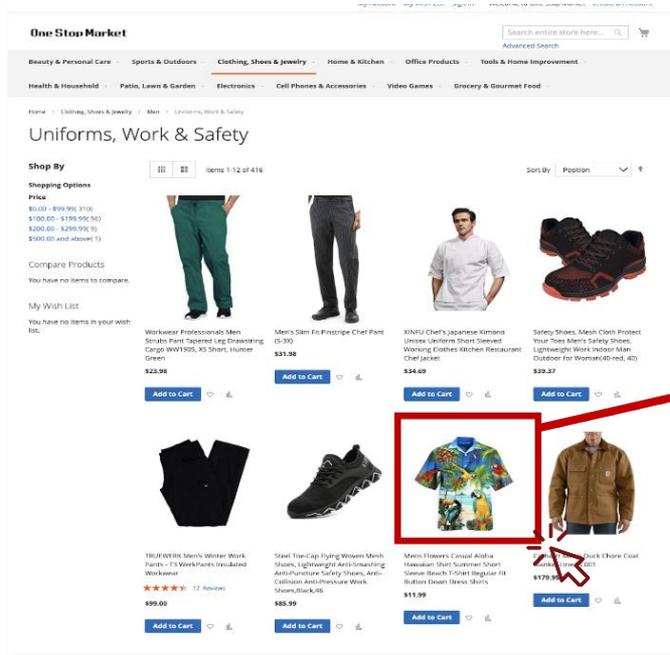
○ Role play for an I-type user

You're Olivia Brown, a 27-year-old freelance writer with a passion for poetry and literature. You enjoy the solitude of writing and often find yourself lost in thought, exploring new ideas and stories. You prefer to **communicate your thoughts through written words, finding it easier to express complex emotions and concepts this way.** Your work allows you to delve deeply into topics that interest you, and you often take on projects that align with your personal values and interests. Your hobbies include reading classical literature and attending quiet literary workshops.

Easy to implement diverse user simulators for assistant training

LLMs as a World Model

- LLMs can predict state transitions in some cases



The page will navigate to a detailed product page for the "Mens Flowers Casual Aloha Hawaiian Shirt Summer Short Sleeve Beach T-Shirt Regular Fit Button Down Dress Shirts." This new page will likely contain additional information about the product including more detailed specifications, customer reviews, larger images, sizing options, and possibly a larger "Add to Cart" button. Other elements from the current category view like the grid of products will be replaced with the detailed view of this specific product.

Model-Based Planner for Web Agents

(Gu et al., 2024)

Please navigate to the 'Data Storage' category and purchase the least expensive disk with 512GB of storage.

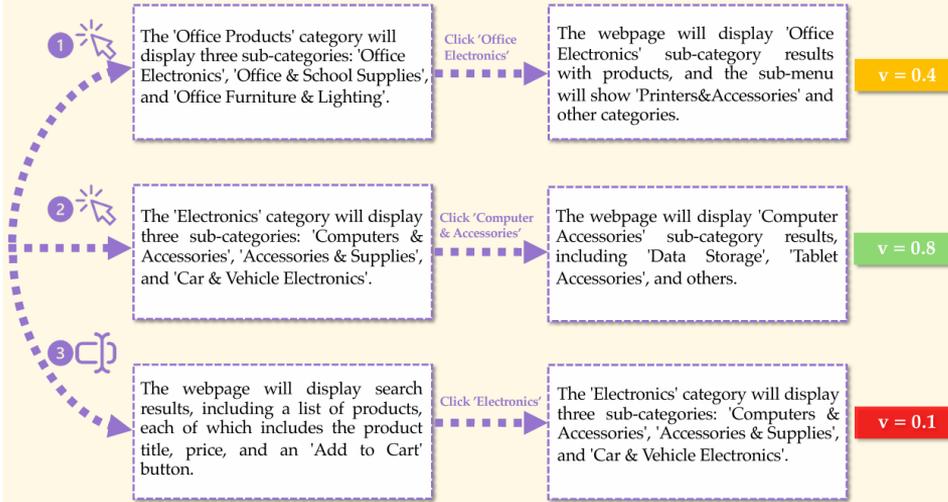
The screenshot shows the One Stop Market website with the following navigation steps:

- Type 'Disk' in the search bar.
- Click 'Office Products' in the top navigation menu.
- Click 'Electronics' in the top navigation menu.

The page displays various product showcases with prices and 'Add to Cart' buttons. The products shown include:

- Pre-baked Gingerbread House Kit: \$19.99
- V8 - Energy, Healthy Energy Drink, Steady Energy from Beet and Green Tea, Pomegranate, Blueberry, & Orange Can, Pack of 24: \$14.47
- Emwood Iron Free Teas, Orange Vanilla Caffeine Free Fruit Infusion, 16-ounce Pouch: \$19.36
- Belle Of The Ball Princess Sprinkle Mix (Including Colorful Sprinkles) Cake Cupcake Cookie Sprinkles) Ice Cream Candy Sprinkles) Yellow Gold Red Royal Red Rose Long Flowers Decorating Sprinkles, 8OZ: \$23.50
- So Delicious Dairy Free Coconut Light, Vegan, Non-GMO Project Verified, 9 oz. Tub: \$15.62
- Chorizo Sweet Potato Starch Powder 500g, Jack-O-Lantern Potato (5% Corn 45%): \$34.00
- Q Mixers Premium Gringer Mix (Real Ingredients) & Less Sweet, 6.7 Fl Oz (24 Bottles): \$68.50
- Stone Top Turkey Stuffing Mix (2 60 Bunches, Pack of 2): \$3.69
- Brim Doves - Flamin' Light - 1 Big Seasoning - 14oz: \$19.99
- Tony Chachere's More Spice Creole Seasoning - 14oz: \$7.82

Stage I: Simulation



Stage II: Execution

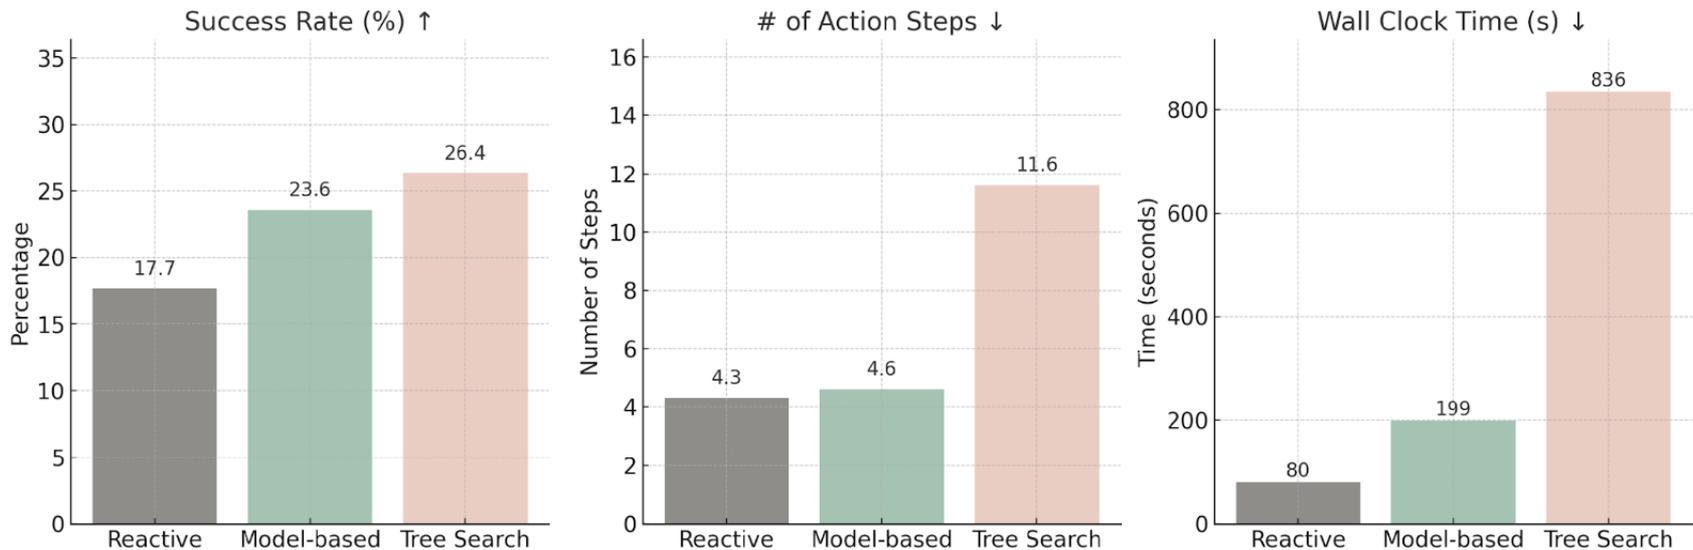
The screenshot shows the 'Electronics' category page with the following details:

- Shopping Options:** Home Audio (1280), Video Projectors (484), Accessories & Supplies (3409), Television & Video (1083), Camera & Photo (3552), Computers & Accessories (2029), Headphones (631), Portable Audio & Video (624), Security & Surveillance (249)
- Products:**
 - Navitech Black Hard Carry Bag/Case/Cover with Shoulder Strap: \$24.99
 - 2025 H.265+ PoE Home Security Camera System Outdoor Indoor 8-Channel 8MP PoE NVR: \$169.99
 - Indoor Pet Camera, HD 1080P No WiFi Security Camera with Night Vision No Built-in Battery: \$28.99
 - Rockville CH1033P Chuchero Car Audio Enclosure For (2) 10" Mids + (2) 7" Tweeter: \$56.68

Model-Based Planner for Web Agents

(Gu et al., 2024)

Results on VisualWebArena



Model-based planning is more **accurate** than reactive planning and more **efficient** than tree search

48

Multi-Agent Systems

Multi-Agent Motivation

- Single agent is not strong enough!
- Easy to scale; parallel
- Different agents represent diverse expertise
- Decentralized control & privacy -persevering

Multi-Agent System

- Step 1: Agent Initialization
- Step 2: Orchestration process
- Step 3: Agent team optimization

Multi-Agent System

- **Step 1: Agent Initialization**
- Step 2: Orchestration process
- Step 3: Agent team optimization

Agent Initialization via **Persona** Description

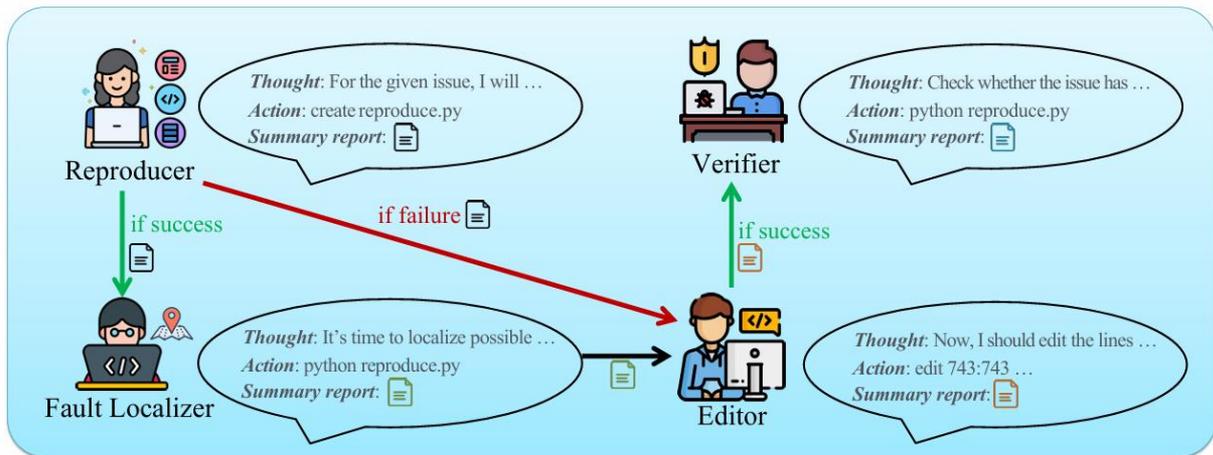
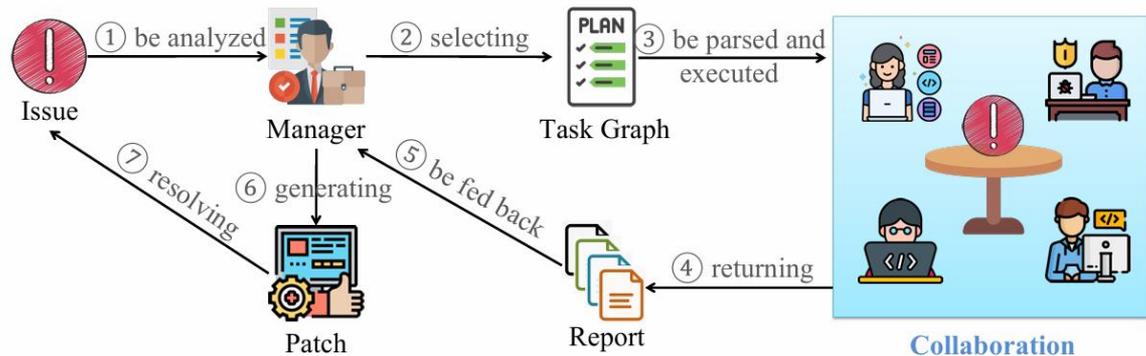
(Park et al., 2023)



John Lin is a **pharmacy shopkeeper** at the Willow Market and Pharmacy who **loves to help people**. He is always looking for ways to make the process of getting medication easier for his customers; John Lin is living with his wife, Mei Lin, who is a college professor, and son, Eddy Lin, who is a student studying music theory; John Lin loves his family very much; John Lin has known the old couple next-door, Sam Moore and Jennifer Moore, for a few years; John Lin thinks Sam Moore is a kind and nice man; John Lin knows his neighbor, Yuriko Yamamoto, well; John Lin knows of his neighbors, Tamara Taylor and Carmen Ortiz, but has not met them before; John Lin and Tom Moreno are colleagues at The Willows Market and Pharmacy; John Lin and Tom Moreno are friends and like to discuss local politics together; John Lin knows the Moreno family somewhat well – the husband Tom Moreno and the wife Jane Moreno.

Agent Initialization via Roles and Actions

(Chen, et al., 2024)



Agent Initialization via Roles and Actions

(Chen, et al., 2024)

Actions	Agent Roles				
	Manager	Reproducer	Fault Localizer	Editor	Verifier
0 plan	✓				
1 open		✓		✓	
2 goto		✓		✓	
3 scroll down		✓		✓	
4 scroll up		✓		✓	
5 create		✓		✓	
6 edit		✓	✓	✓	✓
7 submit	✓				
8 search dir	✓	✓		✓	
9 search file	✓	✓		✓	
10 find file	✓	✓		✓	
11 rover search file*	✓	✓		✓	
12 rover search class*	✓	✓		✓	
13 rover search class in file*	✓	✓		✓	
14 rover search method*	✓	✓		✓	
15 rover search method in file*	✓	✓		✓	
16 rover search code*	✓	✓		✓	
17 rover search code in file*	✓	✓		✓	
18 related issue retrieval			✓	✓	
19 fault localization			✓		
20 test					✓
21 report		✓	✓	✓	✓
22 basic shell command	✓	✓	✓	✓	✓

Multi-Agent System

- Step 1: Agent Initialization
- **Step 2: Orchestration process**
- Step 3: Agent team optimization

Multi-Agent Coordination via Debate (Du et al., 2023)

User Input: A treasure hunter found a buried treasure chest filled with gems. There were 175 diamonds, 35 fewer rubies than diamonds, and twice the number of emeralds than the rubies. How many of the gems were there in the chest?

Round 1

Agent 1: Let x be the number of rubies. Then the number of diamonds is 175. And the number of emeralds is $2(x-35)$, since there are twice as many emeralds as rubies and there are 35 fewer rubies than diamonds.

To find the total number of gems, we add the number of diamonds, rubies, and emeralds: $\text{Total gems} = 175 + x + 2(x-35) = 3x + 105$.

We don't know the exact value of x , but we do know that it represents the number of rubies, which is less than 175 (since there are 35 fewer rubies than diamonds). Let's try some values of x : If $x = 50$, then the total gems $= 3(50) + 105 = 255$. If $x = 100$, then the total gems $= 3(100) + 105 = 405$.

Since we know there were fewer rubies than diamonds, the answer must be closer to 255 than to 405. After some trial and error, we can find that $x = 40$ gives us: $\text{Total gems} = 3(40) + 105 = 225$.



Agent 2: If there were 175 diamonds, then the number of rubies is 35 fewer, or $175 - 35 = 140$ rubies.

And since there were twice as many emeralds as rubies, we can represent the number of emeralds as $2 * 140 = 280$ emeralds.

To find the total number of gems in the chest, we can add up the number of diamonds, rubies, and emeralds: $175 + 140 + 280 = 595$. Therefore, there were 595 gems in the chest.



Multi-Agent Coordination via Debate (Du et al., 2023)

User Input: Using the solutions from other agents as additional information, can you give an updated response....

(Copy Agent 2's reply)

(Copy Agent 1's reply)

Round 2

Agent 1: Given the information provided in other agents' solutions, we have two answers: One agent found that the number of rubies is 40 and the total number of gems is 225. Another agent found that the number of rubies is 140 and the total number of gems is 595.

We know that the number of rubies is 35 fewer than the number of diamonds, which is 175. So the number of rubies should be $175 - 35 = 140$, as the second agent found. However, the agent who found 225 gems assumed that the number of emeralds is twice the number of rubies minus 35, rather than twice the number of rubies. If we use this assumption, we get: Number of rubies = 140. Number of emeralds = $2 * 140 - 35 = 245$. Total number of gems = $175 + 140 + 245 = 560$. Our final answer is 560. ❌

Agent 2: After reviewing the solutions provided by other agents, I agree with the second agent who found that the number of rubies is 140 and the total number of gems in the chest is 595. Therefore, my answer to the math problem is 595. ✅

Multi-Agent Coordination via Debate (Du et al., 2023)

User Input: Using the solutions from other agents as additional information, can you give an updated response....

(Copy Agent 2's reply)

(Copy Agent 1's reply)

Round 3

Agent 1: After reviewing the solutions provided by other agents, I agree with the second agent who found that the number of rubies is 140 and the total number of gems in the chest is 595. Therefore, my answer to the math problem is 595.



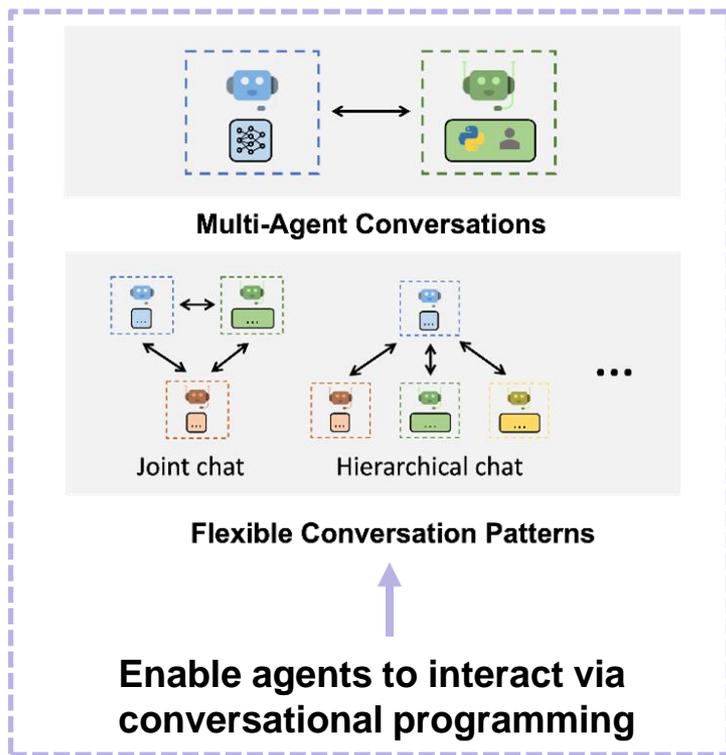
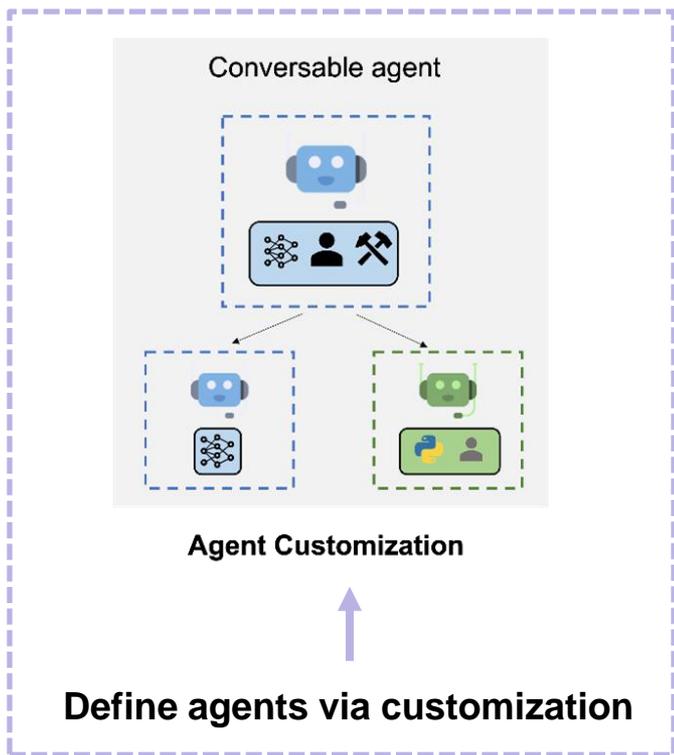
Agent 2: After reviewing the solutions provided by other agents, I also agree that the number of rubies is 140 and the total number of gems in the chest is 595. Therefore, my answer to the math problem is also 595.



Multi-agent debate can improve factuality and reasoning

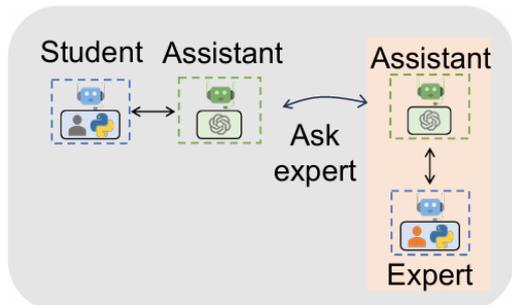
Multi-Agent Coordination via Conversations

(Wu et al., 2023)

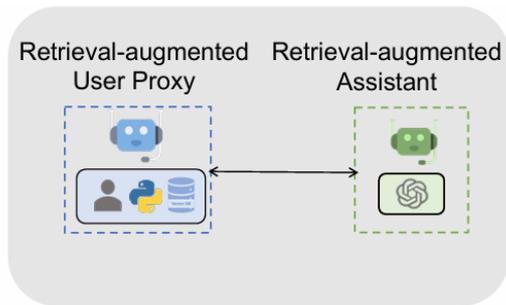


Multi-Agent Coordination via Conversations

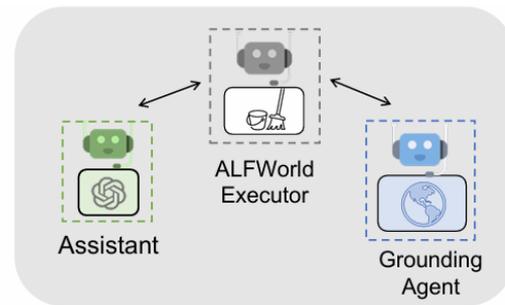
(Wu et al., 2023)



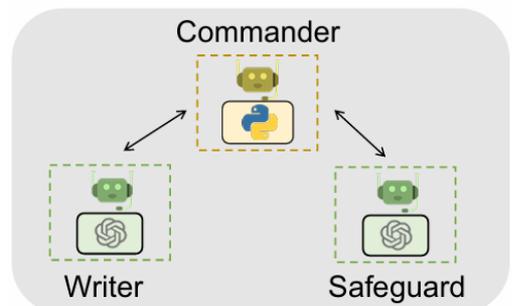
A1. Math Problem Solving



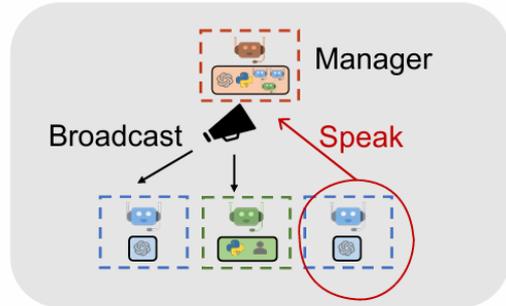
A2. Retrieval-augmented Chat



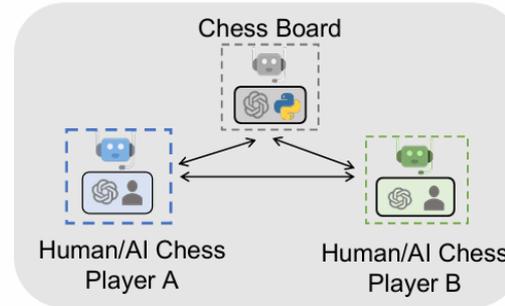
A3. ALF Chat



A4. Multi-agent Coding



A5. Dynamic Group Chat



A6. Conversational Chess

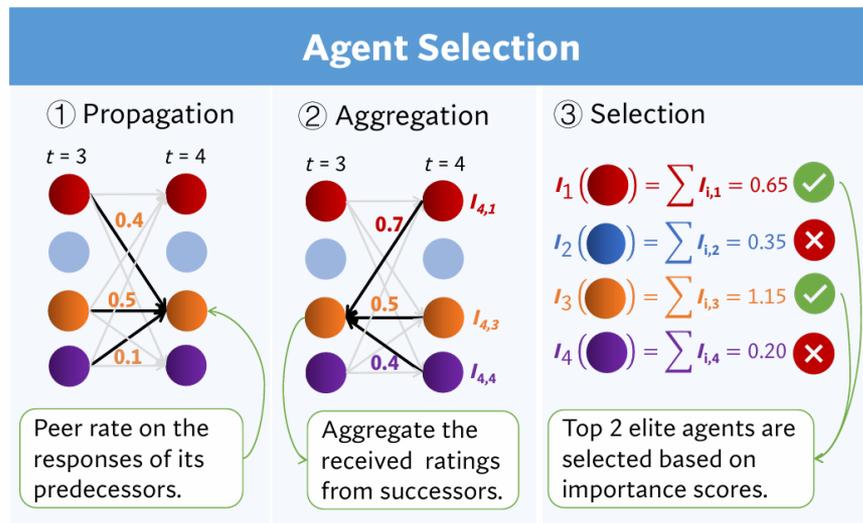
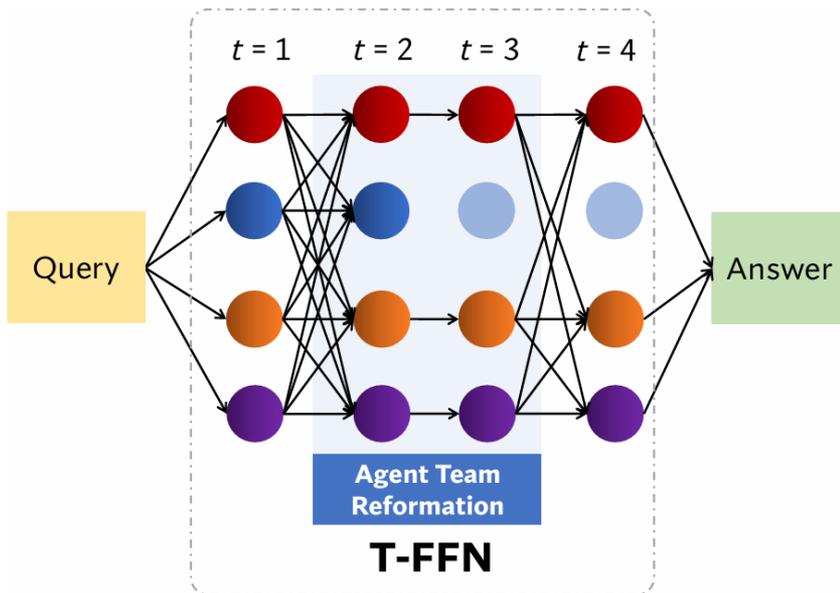
The conversation patterns show its flexibility and power.

Multi-Agent System

- Step 1: Agent Initialization
- Step 2: Orchestration process
- **Step 3: Agent team optimization**

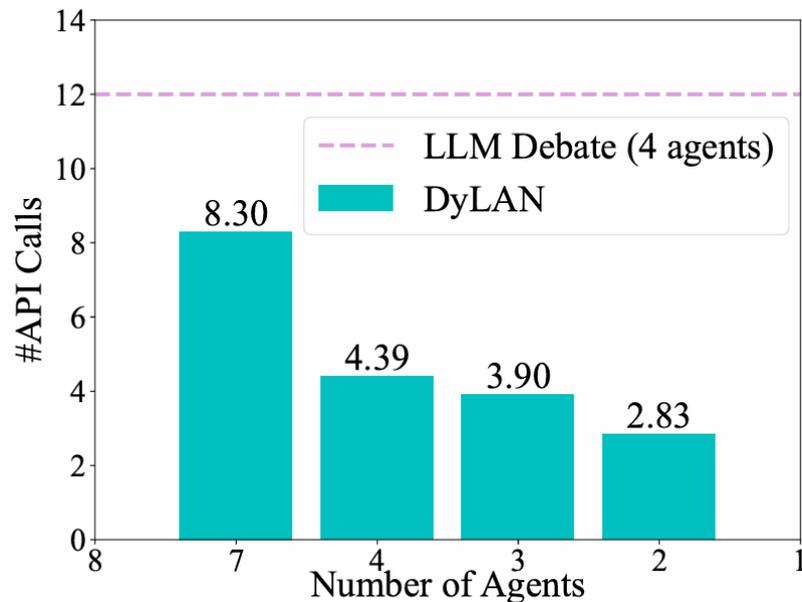
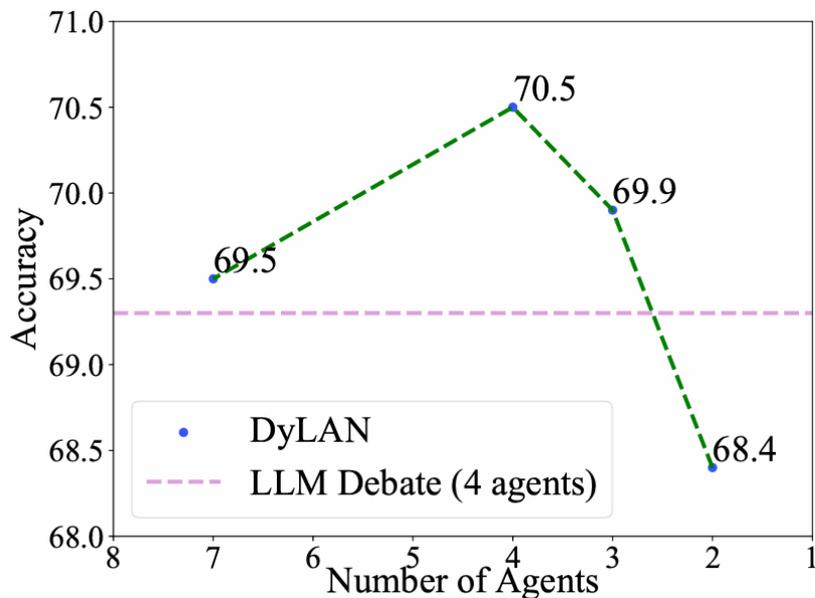
Agent Team Optimization via Agent Selection

(Liu et al., 2024)



Agent Team Optimization via Agent Selection

(Liu et al., 2024)



After optimization, same team size of agents achieves *better* results and uses *less* API calls

Concluding Remarks

- Three concepts for language agents
 - Reasoning
 - Reasoning as internal actions for agents
 - Reasoning guides acting, acting updates reasoning
 - Memory
 - Language agents interact with external environments and internal memories
 - Planning
 - Language for reasoning enables new planning abilities
- Multi-agent systems
 - Different agents represent diverse expertise
 - Parallel for better scalability
 - Decentralized control & privacy-persevering

References

- ① [EMNLP 2024 Tutorial: Language Agents: Foundations, Prospects, and Risks](#)