CS769 Advanced NLP

Language Agents and LLM Programming

Junjie Hu

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Slides adapted from agent tutorial at EMNLP 2024

https://junjiehu.github.io/cs769-fall25/

Goal for Today

- **Part I: Introduction of Language Agents**
  - History, different AI agents
- **Part II: Foundations: Key Components of Language Agents**
  - Reasoning, Memory, Planning
- **Part III: LLM Coding Agents for Software Development**
  - Open Challenges

What is agent?

"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 Mordern Approach* (2020)

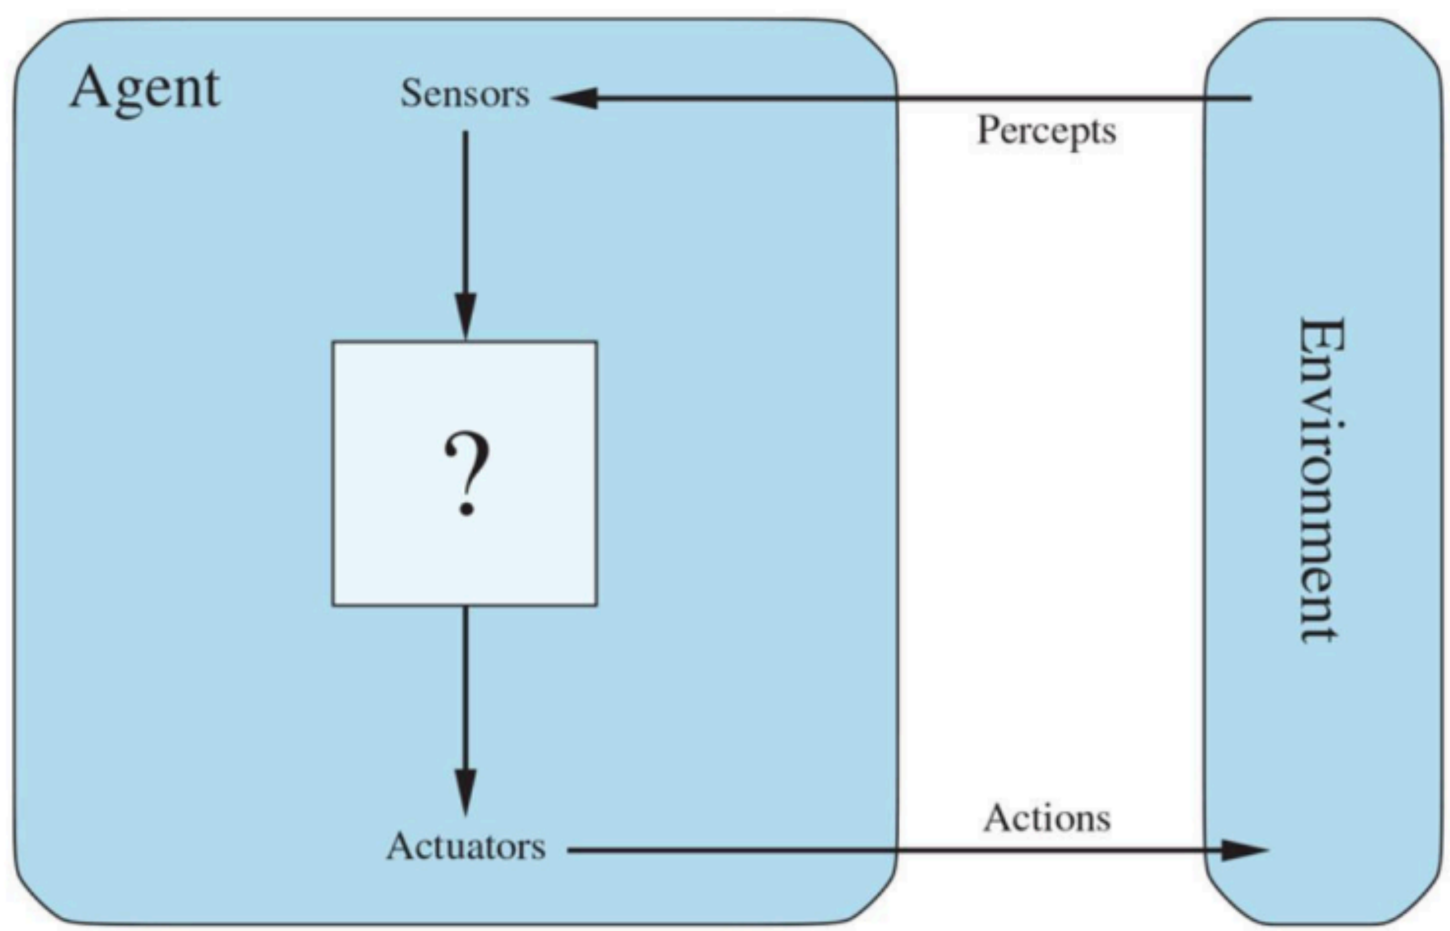

“Modern” agent = LLM + External Environment?

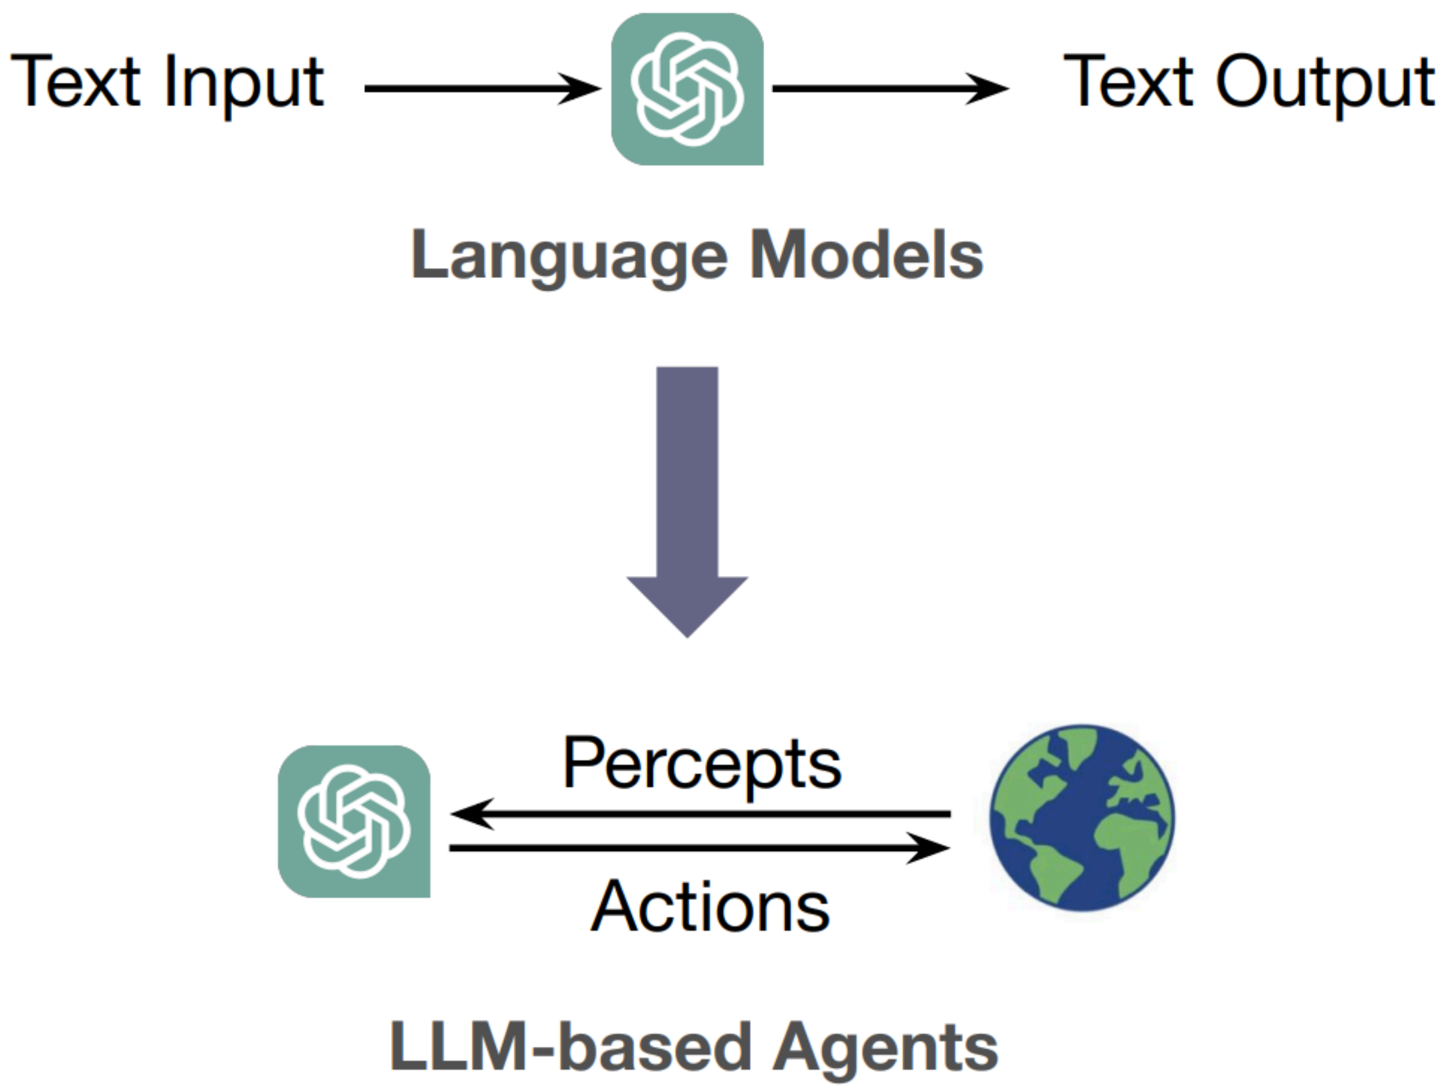

“Modern” agent = LLM + External Environment?

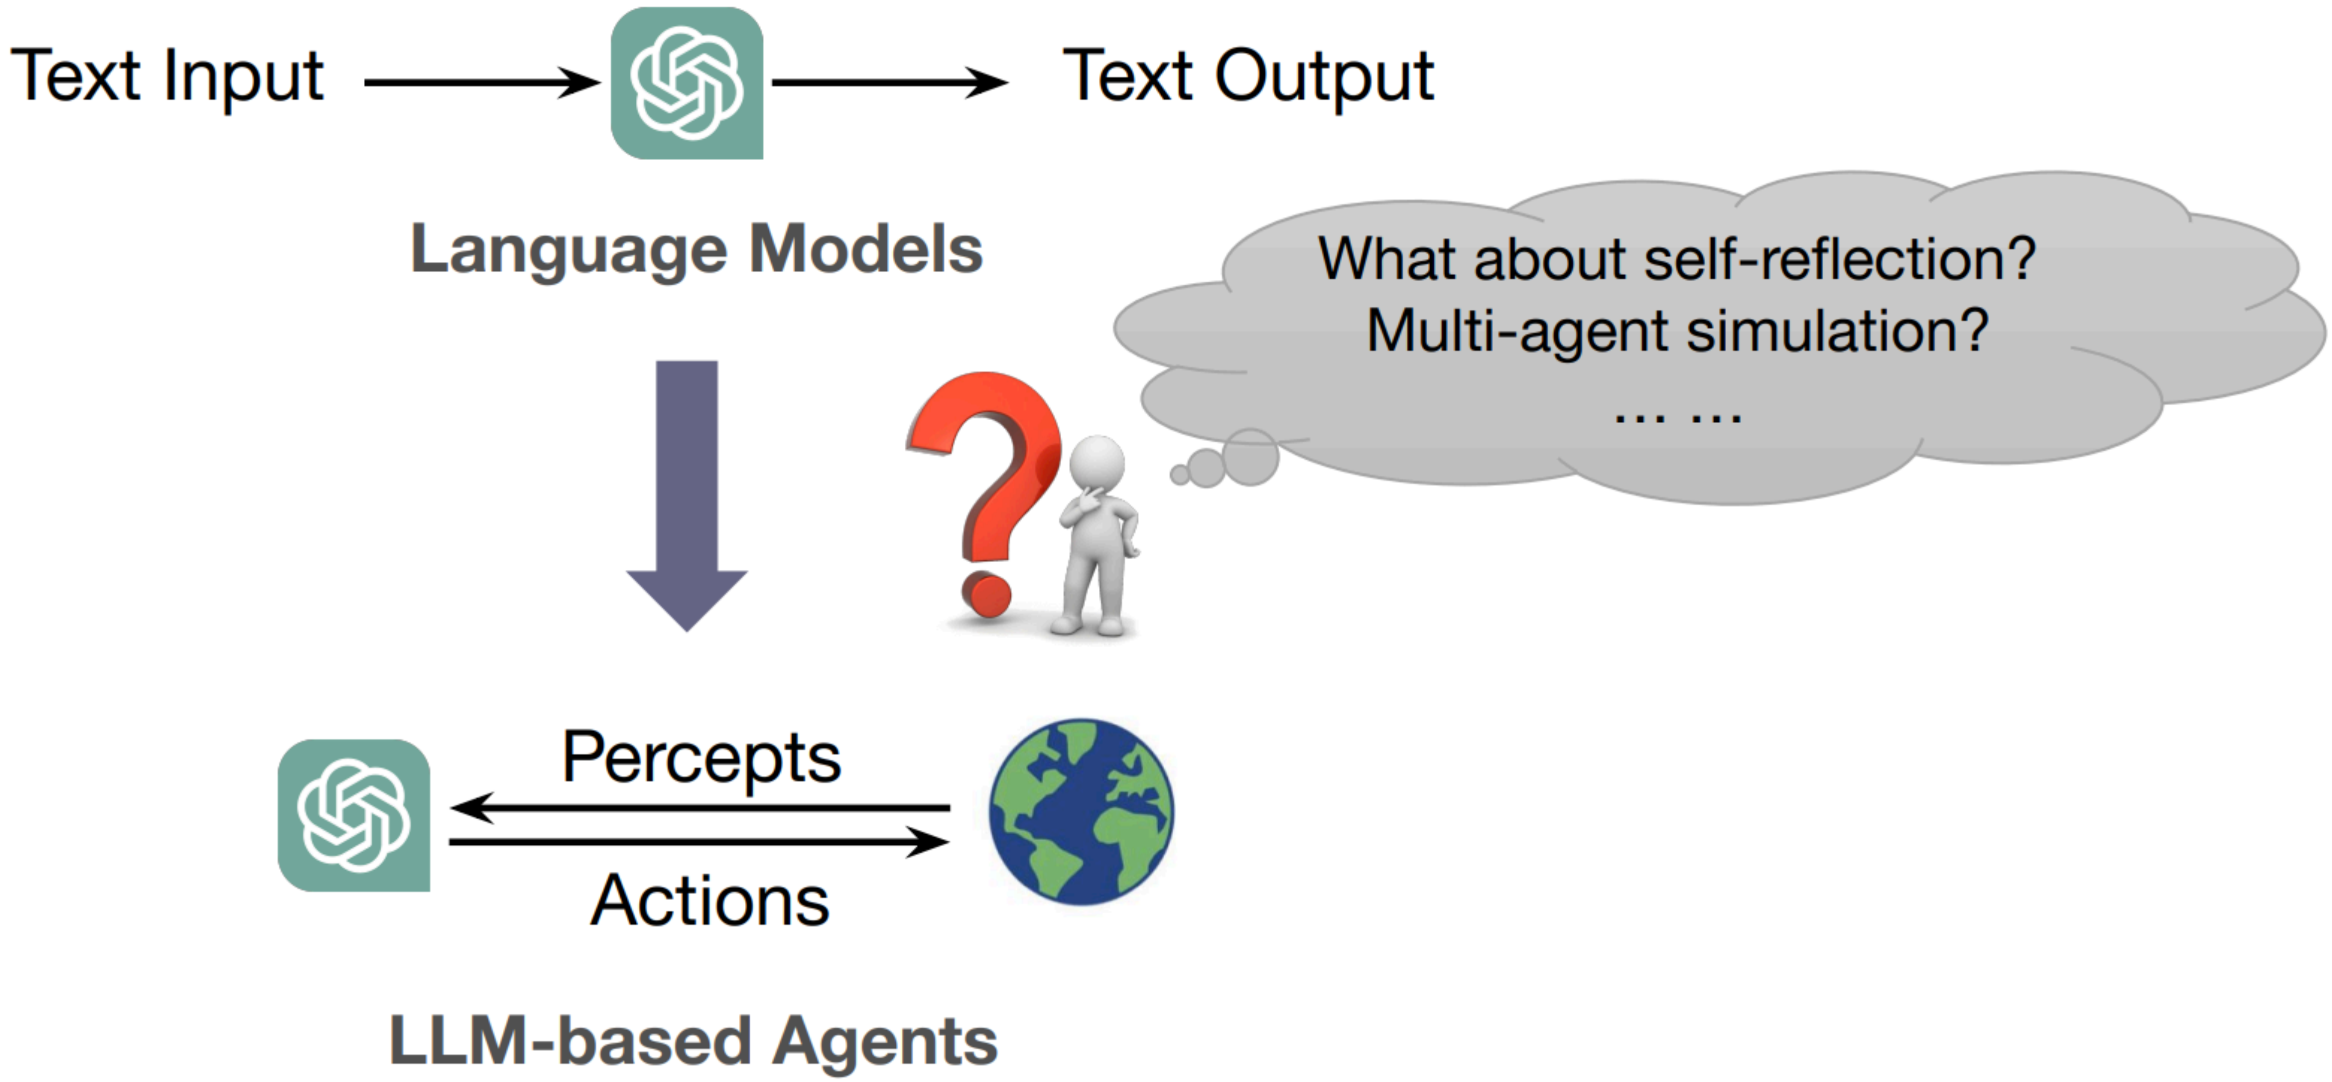

What's fundamentally different now?

- Contemporary AI agents, with integrated LLM(s), can use language as a vehicle for reasoning and communication
  - Instruction following, in-context learning, output customization
  - Reasoning (for better acting): state inferences, self-reflection, replanning, etc.

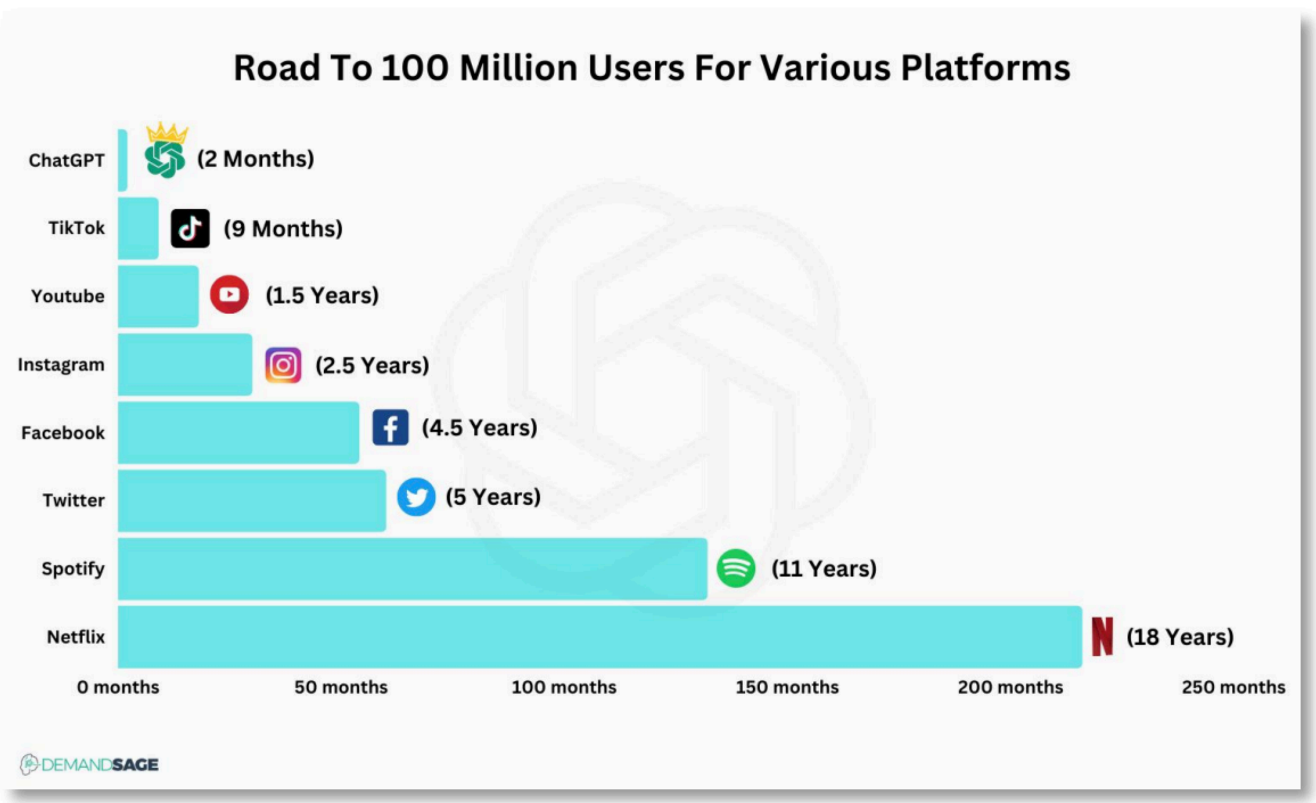

https://www.demandsage.com/chatgpt-statistics/

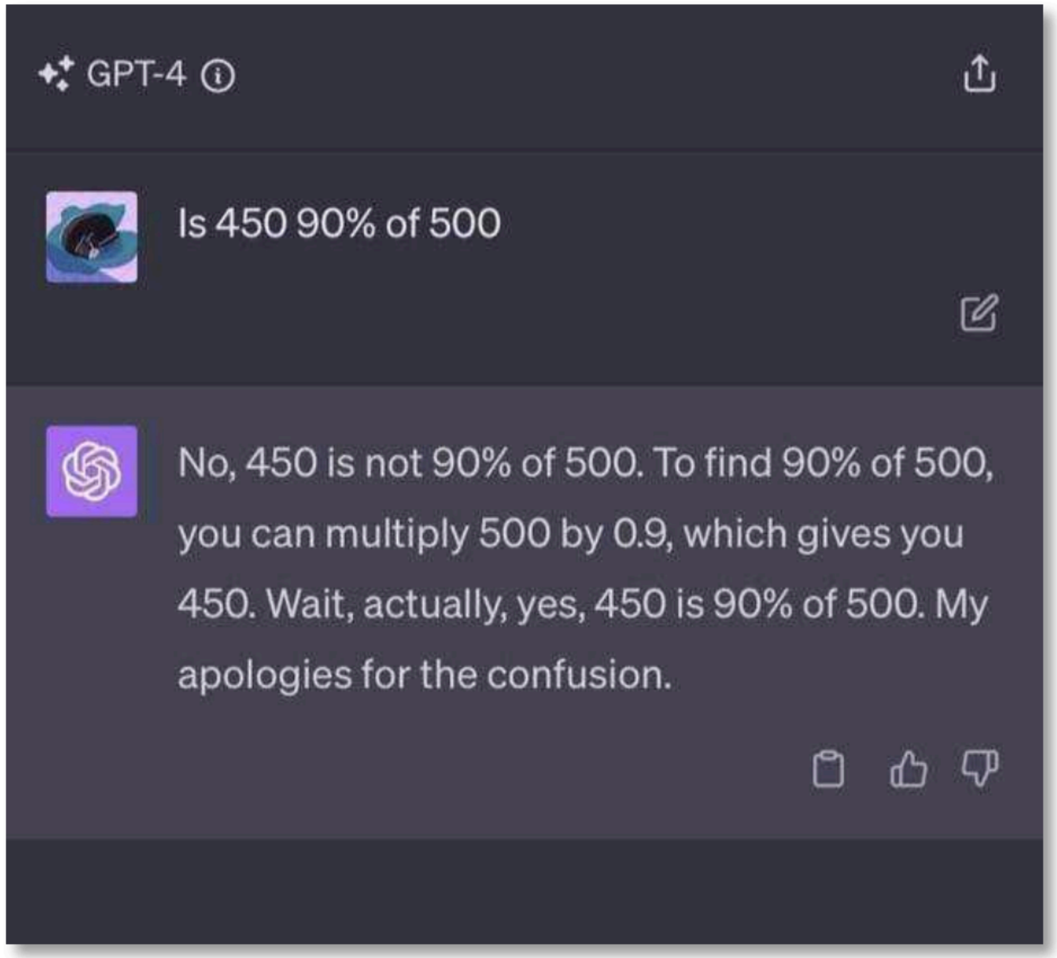

https://www.reddit.com/r/ChatGPT/comments/16jvl4x/wait_actually_yes/

Language agents: a new type of AI agents

- These contemporary AI agents capable of using language for reasoning and communication are best called "**language agents**." They are qualitatively a different type of AI agents with **language** being their most distinct trait.
- What about **multimodal agents**?
  - While there's perception of other input modalities, language is still doing the heavy lifting (i.e., reasoning and communication)
- What about simply LLM agents?
  - The key is using language for reasoning and communication, but that doesn't have to come from an LLM;
  - Maybe in a few years, we will move beyond LLMs, but the need for universal language understanding and production in agents will remain

Reconciling with classical view of language agents

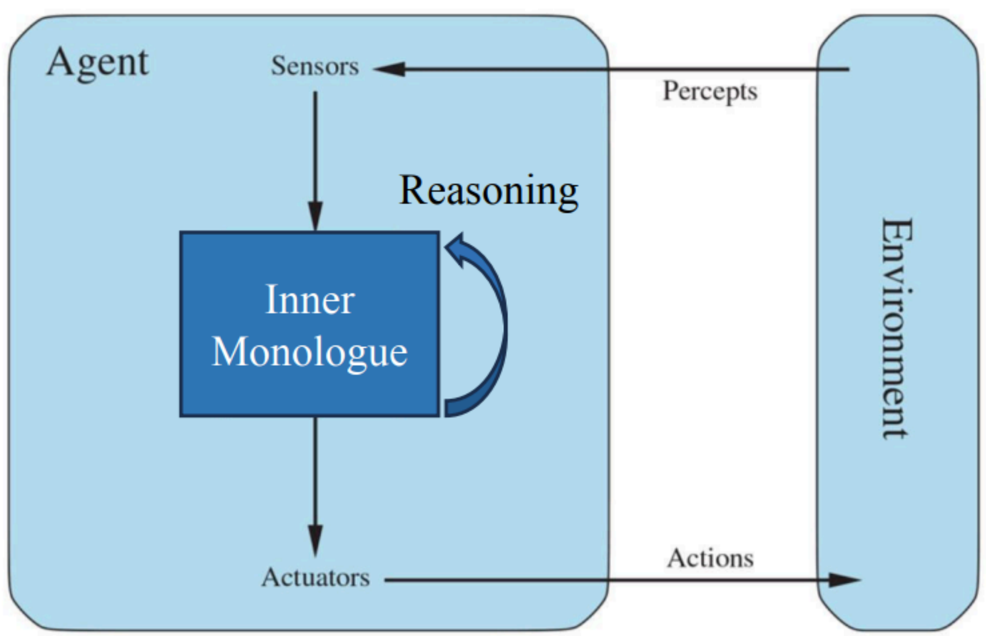

Adapted from Russell & Norvig (2020)

- **Reasoning by generating tokens** is a new type of action (vs. actions in external environments)
- **Internal environment**, where reasoning takes place in an inner monologue fashion
- **Self-reflection** is a 'meta' reasoning action (i.e., reasoning over the reasoning process), akin to metacognitive functions
- **Reasoning is for better acting**, by inferring environmental states, retrospection, dynamic replanning, etc.
- **Percept** and **external action spaces** are substantially expanded, thanks to using language for communication and multimodal perception

Evolution of AI Agents

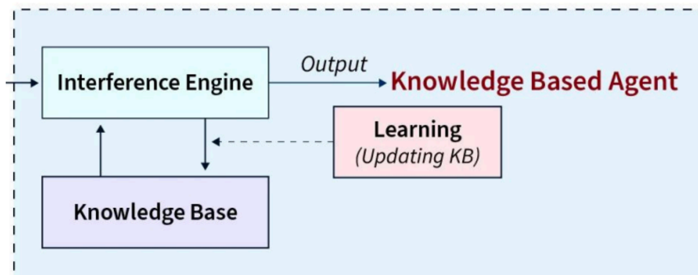

Logical Agent

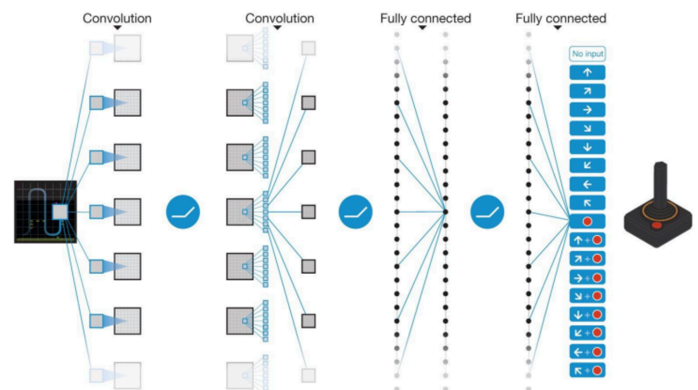

Neural Agent

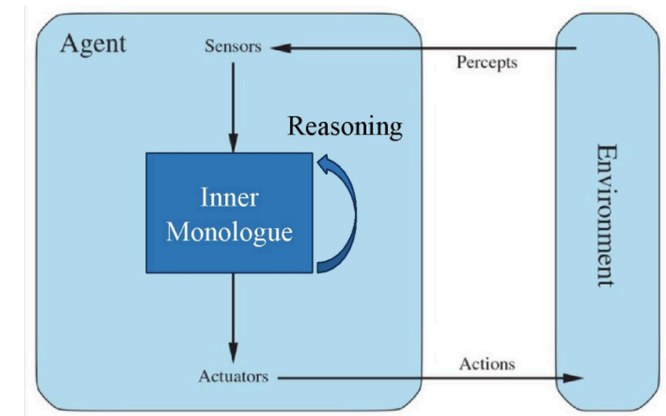

Language Agent

Expressiveness			
Reasoning			
Adaptivity			

Image sources: https://www.scaler.com/topics/artificial-intelligence-tutorial/knowledge-based-agent/
Mnih et al., “Human-level control through deep reinforcement learning.” Nature (2015)

Evolution of AI Agents

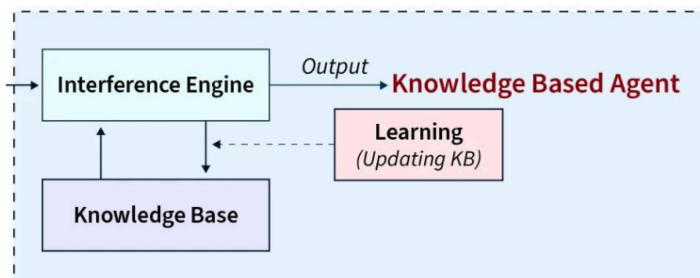

Logical Agent

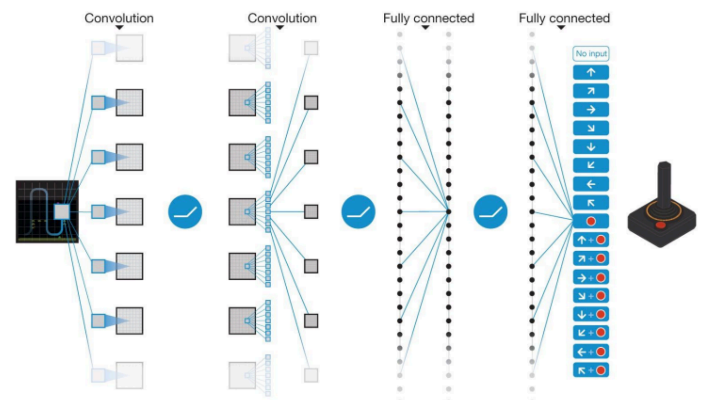

Neural Agent

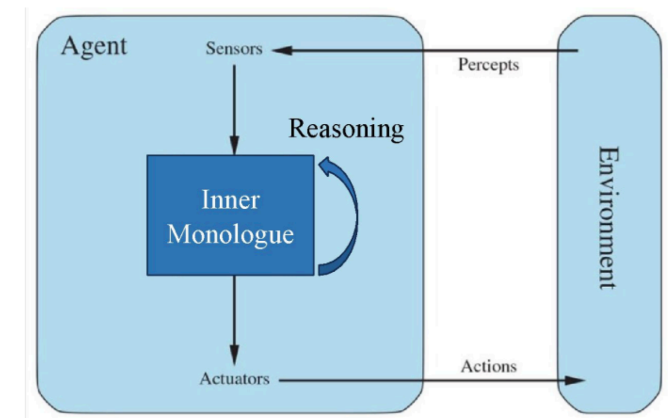

Language Agent

Expressiveness	Low bounded by the logical language		
Reasoning	Logical inferences sound, explicit, rigid		
Adaptivity	Low bounded by knowledge curation		

Image sources: https://www.scaler.com/topics/artificial-intelligence-tutorial/knowledge-based-agent/
Mnih et al., "Human-level control through deep reinforcement learning." Nature (2015)

Evolution of AI Agents

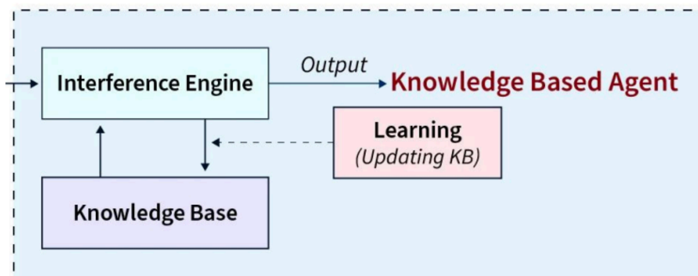

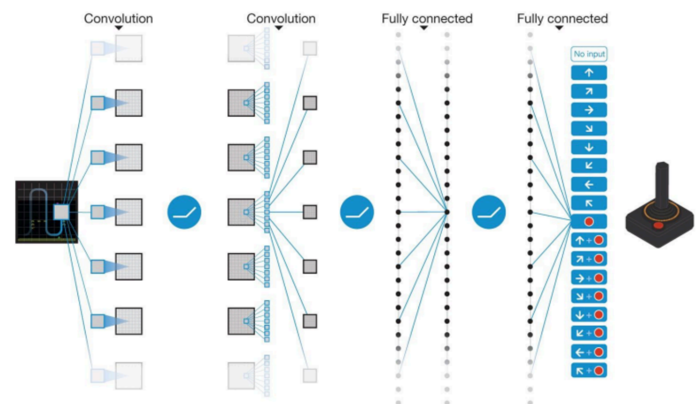

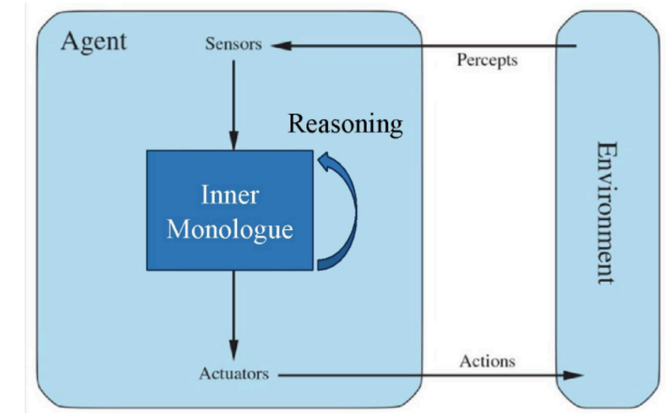

	Logical Agent	Neural Agent	Language Agent
Expressiveness	Low bounded by the logical language	Medium anything a (small) NN can encode	
Reasoning	Logical inferences sound, explicit, rigid	Parametric inferences stochastic, implicit, rigid	
Adaptivity	Low bounded by knowledge curation	Medium data-driven but sample inefficient	

Image sources: https://www.scaler.com/topics/artificial-intelligence-tutorial/knowledge-based-agent/
Mnih et al., "Human-level control through deep reinforcement learning." Nature (2015)

Evolution of AI Agents

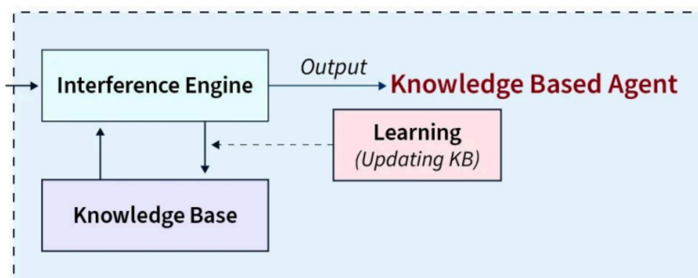

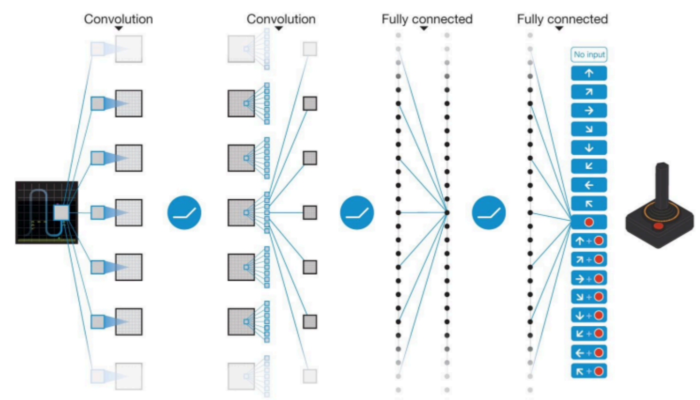

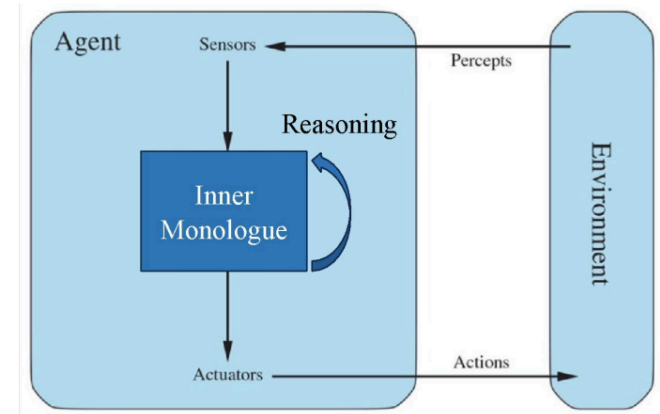

	Logical Agent	**Neural Agent**	**Language Agent**
Expressiveness	Low bounded by the logical language	Medium anything a (small) NN can encode	High almost anything, esp. verbalizable parts of the world
Reasoning	Logical inferences sound, explicit, rigid	Parametric inferences stochastic, implicit, rigid	Language-based inferences fuzzy, semi-explicit, flexible
Adaptivity	Low bounded by knowledge curation	Medium data-driven but sample inefficient	High strong prior from LLMs + language use

Image sources: https://www.scaler.com/topics/artificial-intelligence-tutorial/knowledge-based-agent/
Mnih et al., “Human-level control through deep reinforcement learning.” Nature (2015)

Part II: Foundations: Reasoning, Memory, and Planning

Cognitive Language Agents

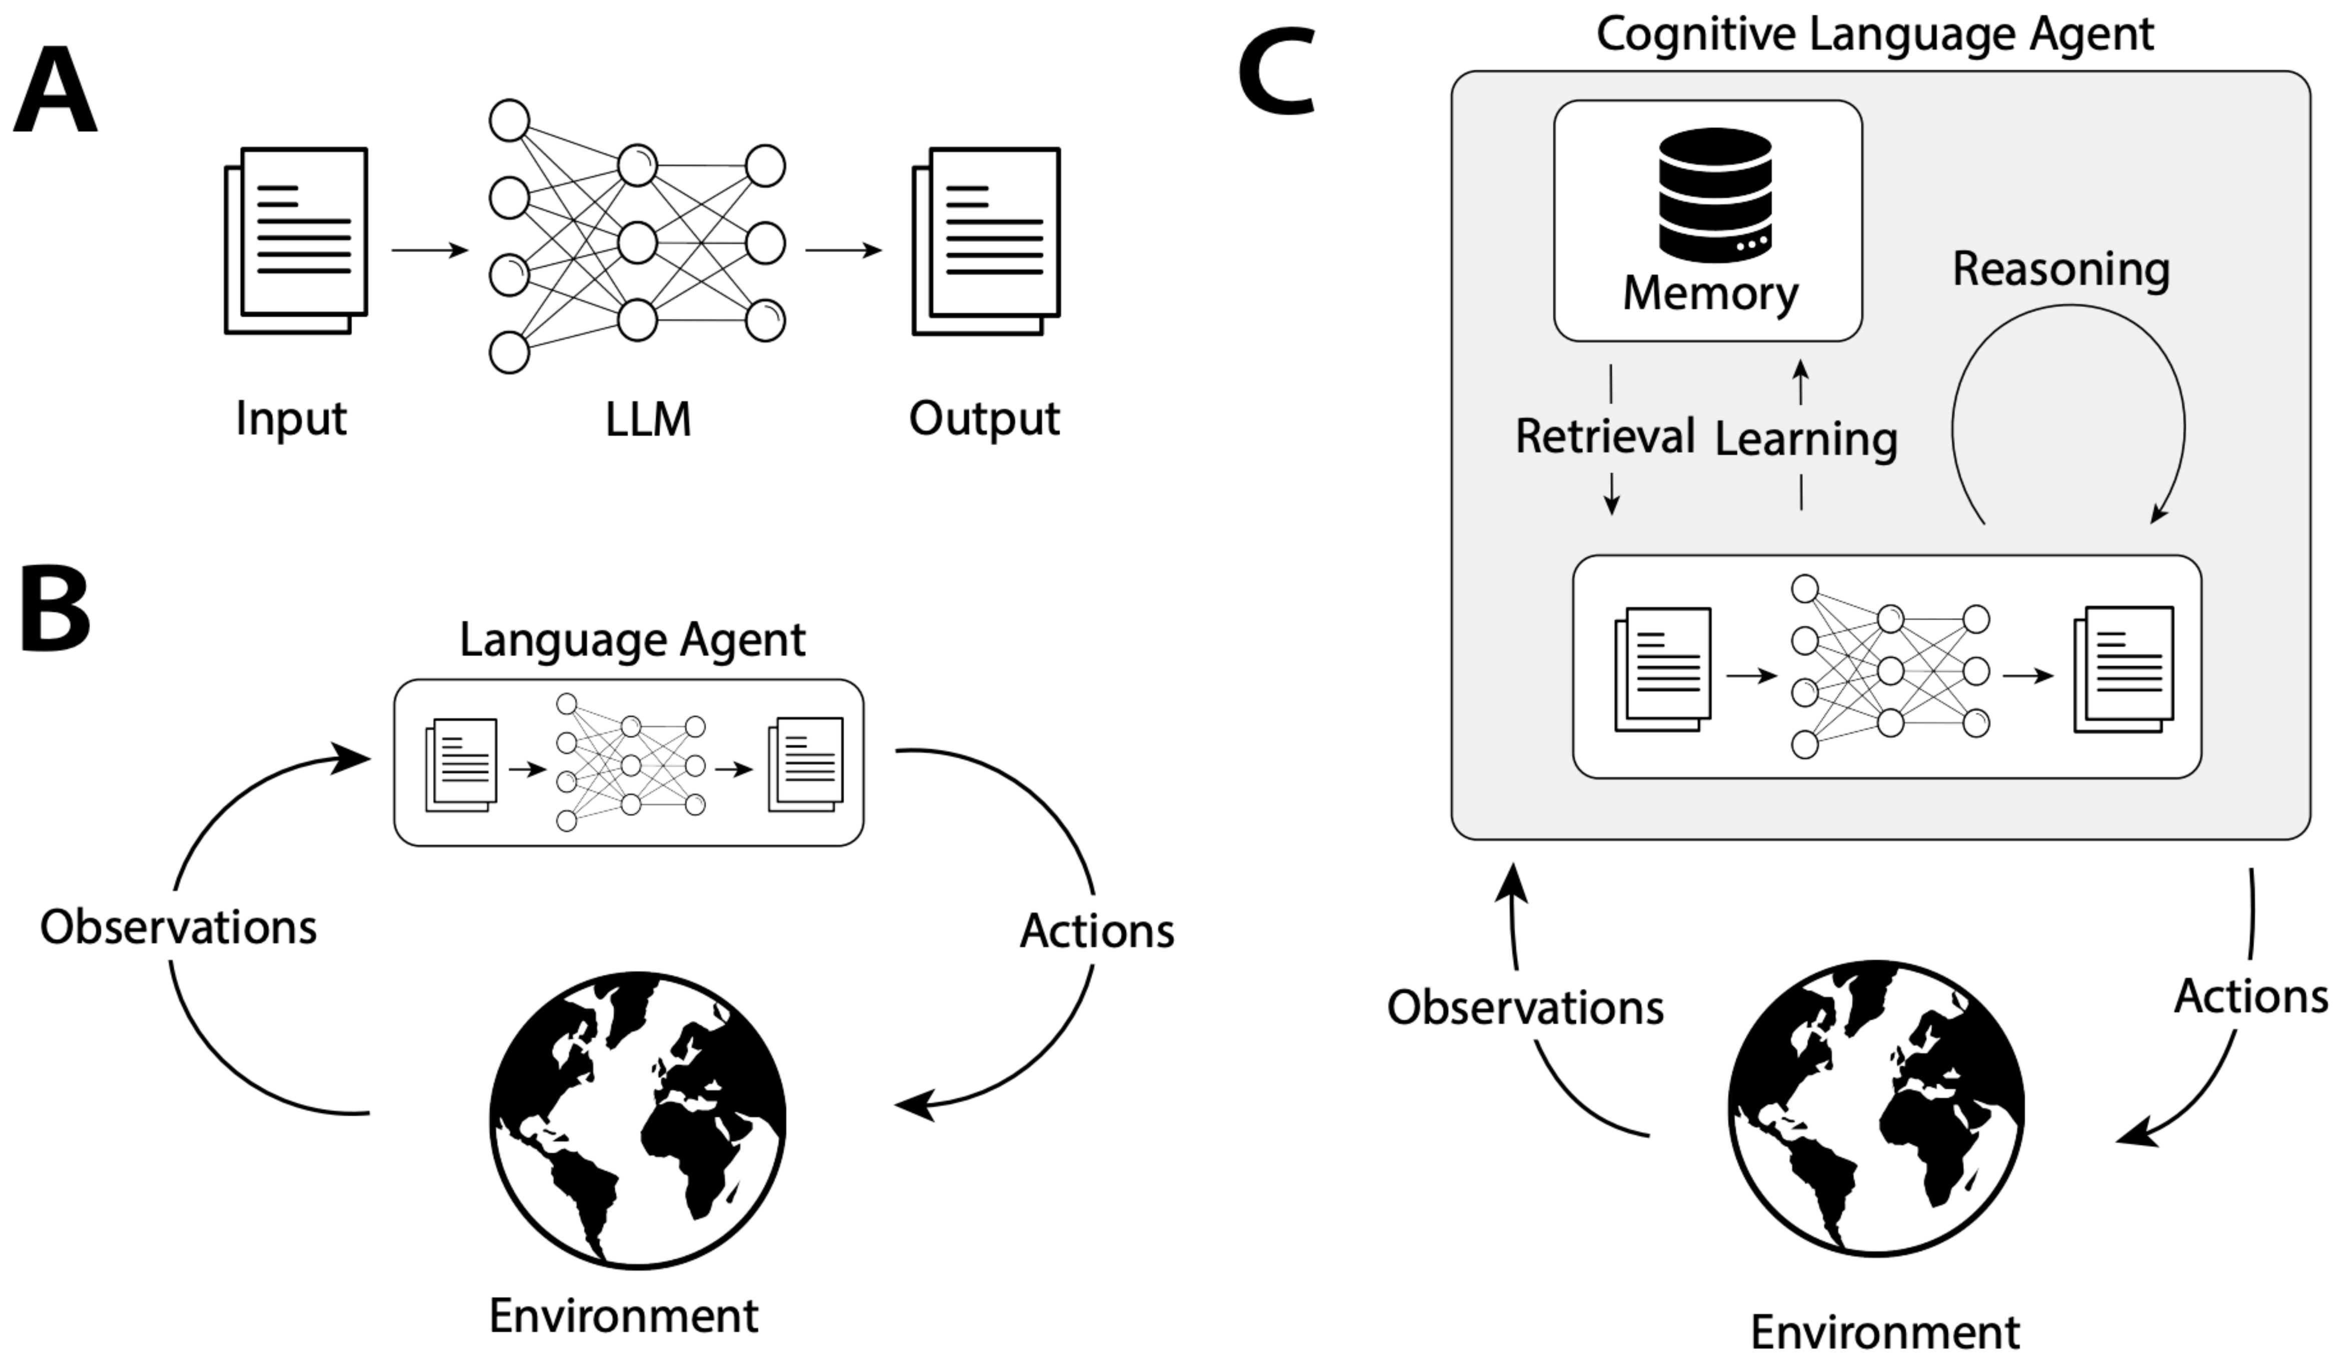

Sumers et al., 2024. Cognitive Architectures for Language Agents, TMLR.

Key Concepts for Language Agents

- **Action space** (beyond environment actions)
  - **Reasoning**: update short-term memory (context window)
  - **Retrieval/Learning**: read/write long-term memory (model weights, vector store, self-notes, event flows, etc)
- **Planning**: (inference-time) algorithm to choose an action from the action space
- **Environment**: receives an action from the agent and provides a reward

Sumers et al., 2024. Cognitive Architectures for Language Agents, TMLR.

Reasoning

- For humans: various mental processes

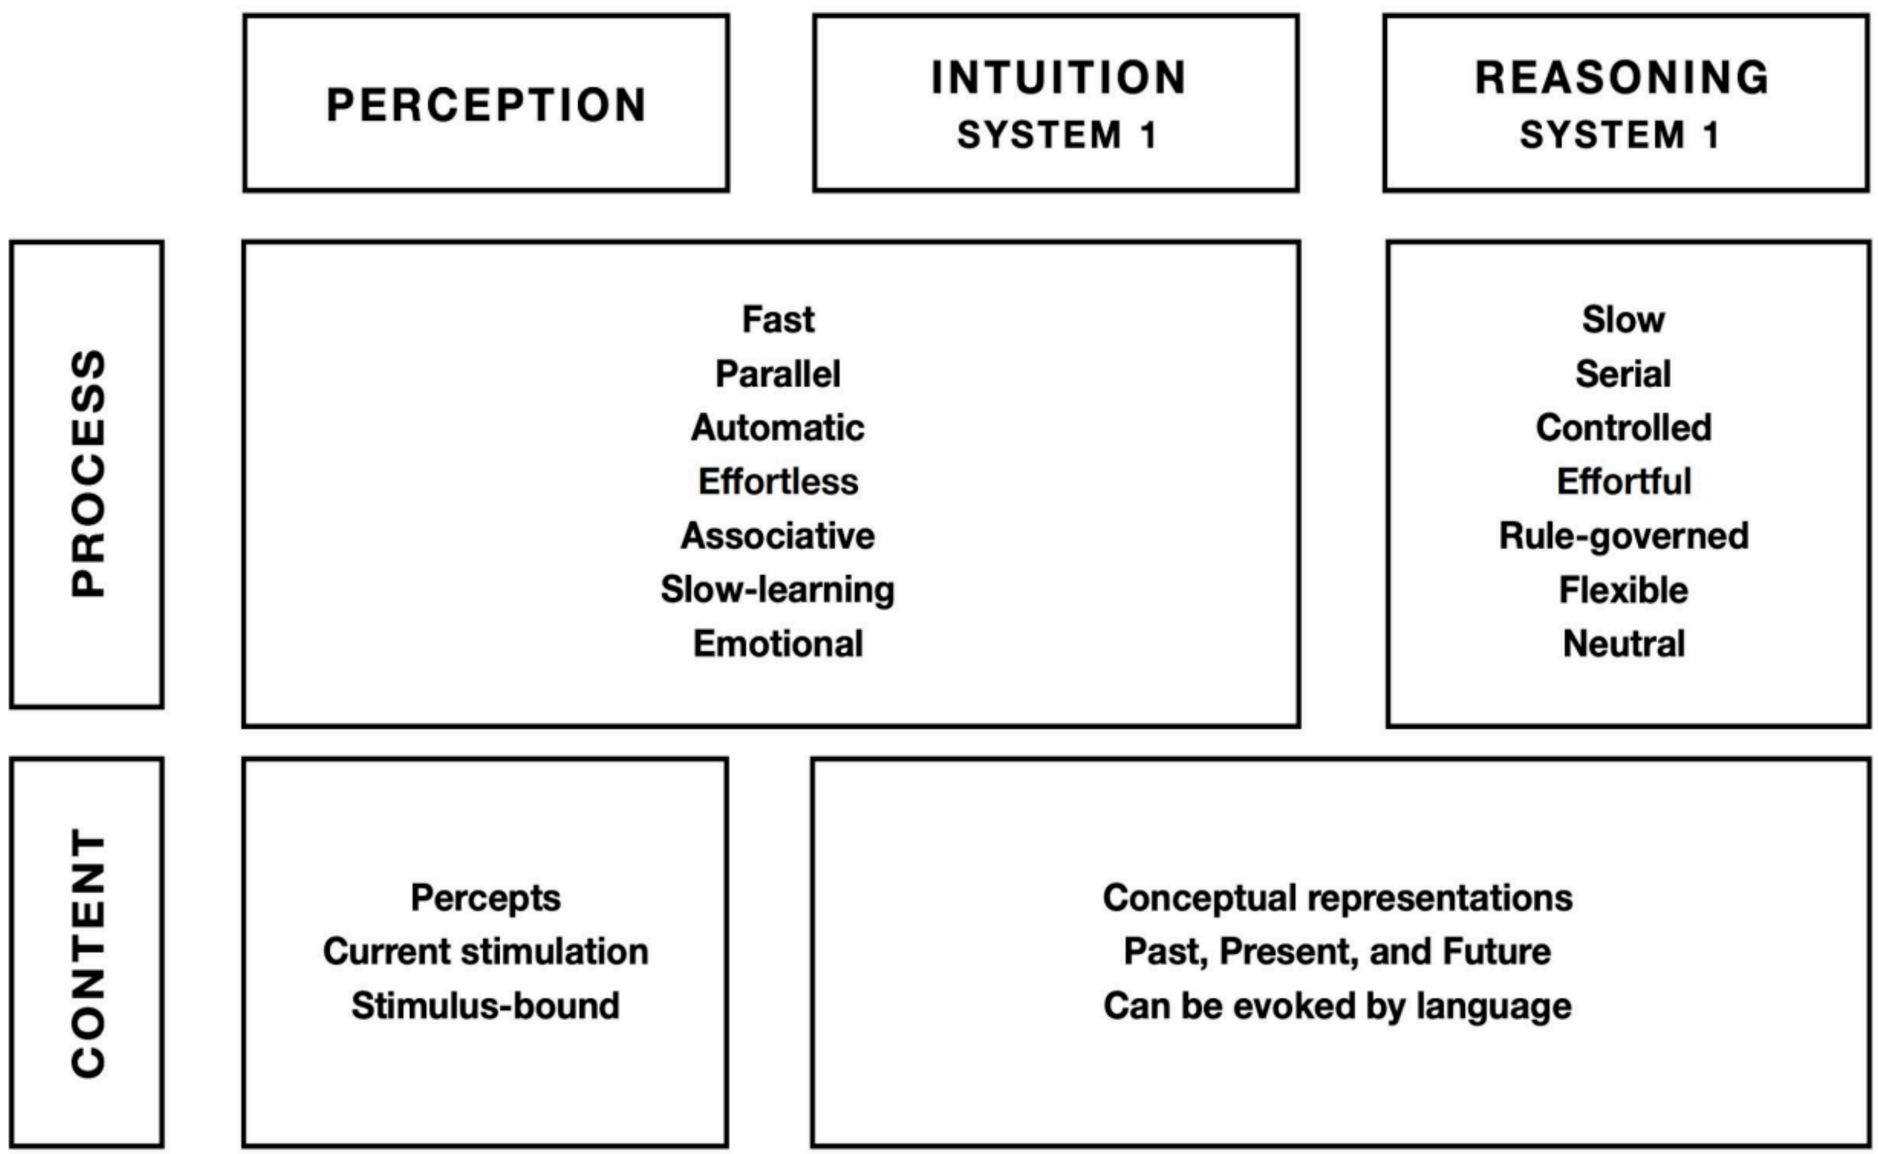

*(Kahneman, 2003), also note that this notion of dual process is put in question. See, e.g., Mercier and Sperber (2017)

Reasoning

- For humans: various mental processes
- For LMs: intermediate generation (Chain-of-Thought)

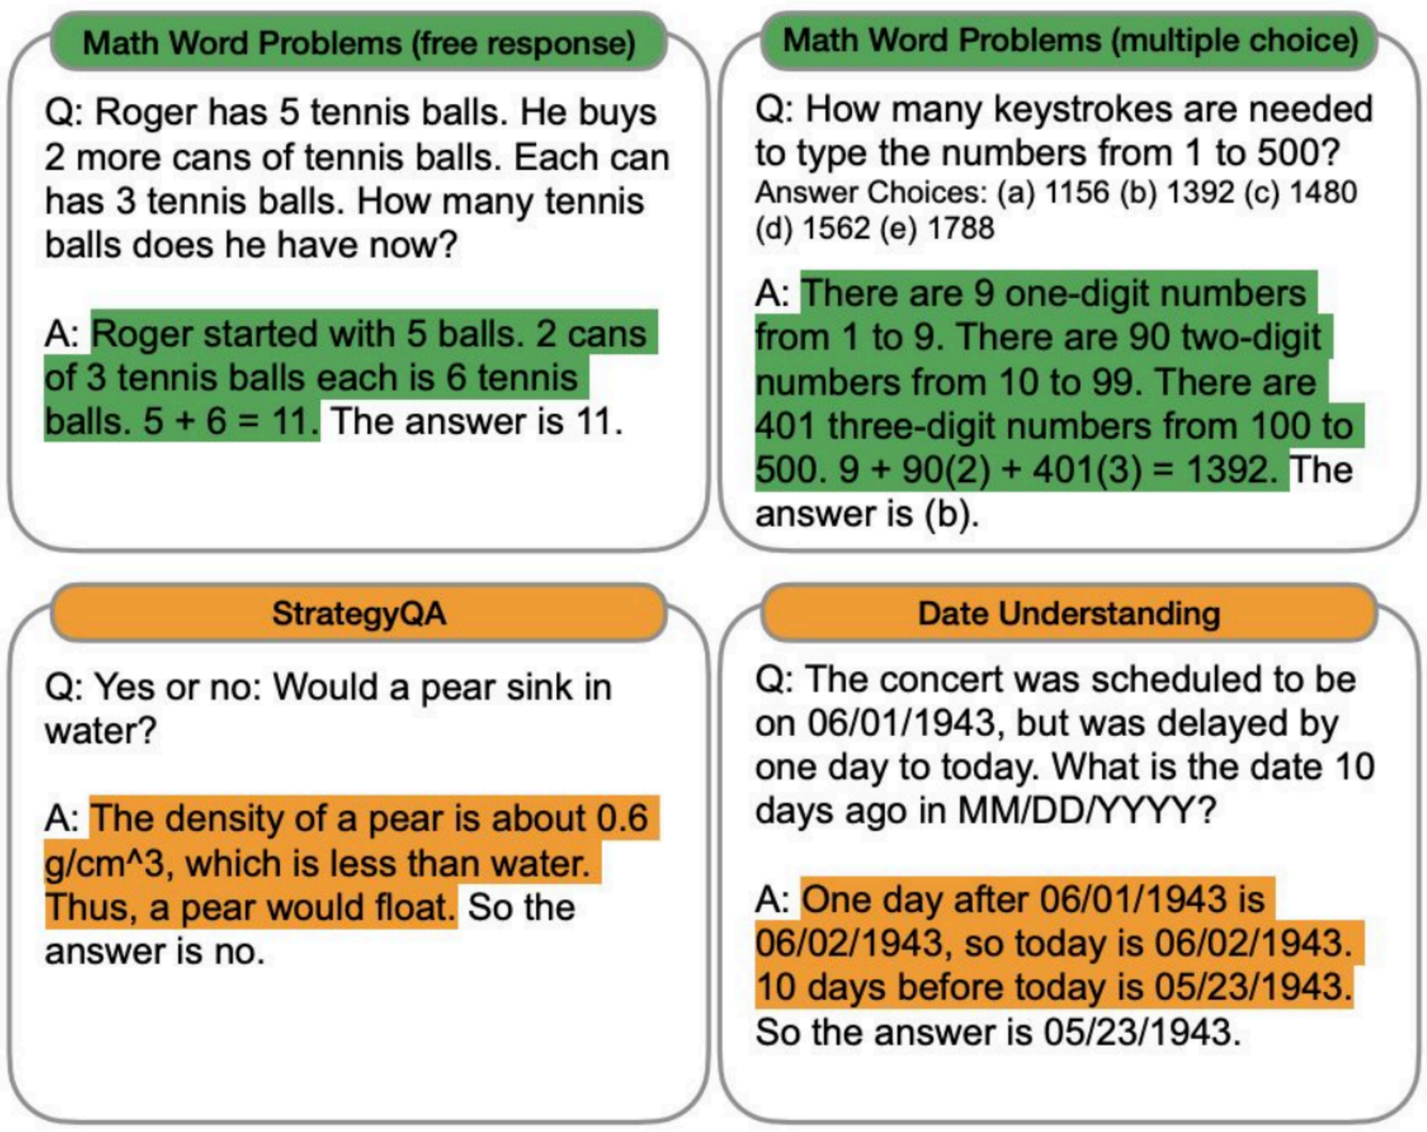

*(Wei et al., 2022), also see (Ling et al., 2017; Cobbe et al., 2021; Nye et al., 2021)

Reasoning

- For humans: various mental processes
- For LMs: intermediate generation (Chain-of-Thought)
- For agents: internal actions.

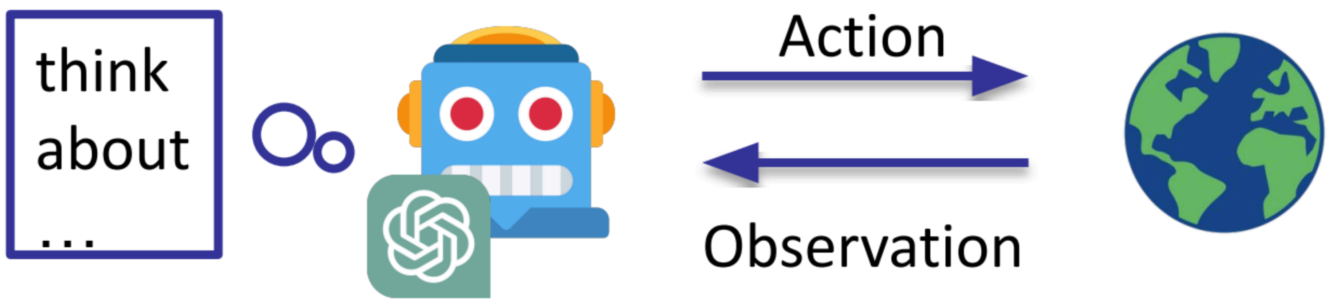

Why is reasoning helpful for agents?

Obs t — You are cooking a dish and seeing salt is out...

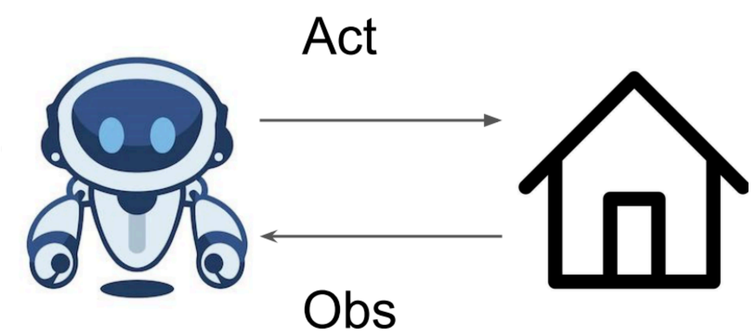

Act t — Turn right

Obs t+1 — You see a cabinet and a table...

Act t+1 — Open cabinet

......

Reason to act

1. Generalization
2. Alignment

Observation to Action mapping could be hard

Why is reasoning helpful for agents?

Obs t	You are cooking a dish and seeing salt is out...
Reasoning	"The dish should be savory, and since salt is out, I should find the soy sauce instead. It is in the cabinet to my right..."
Act t	Turn right
Obs t+1	You see a cabinet and a table...
Act t+1	Open cabinet
......	

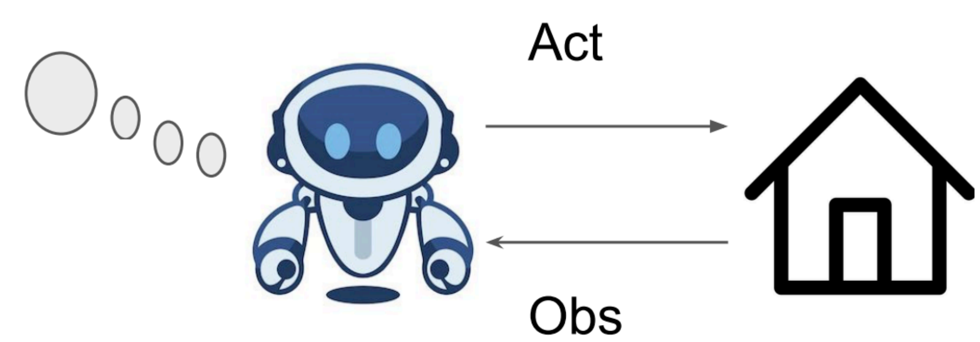

Reason to act

1. Generalization
2. Alignment

Observation to Action mapping could be hard

Conversely, acting also helps reasoning

SH who is PM of UK?

As of December 2022, the Prime Minister of the United Kingdom is Boris Johnson. However, since I am a large language model trained by OpenAI, my knowledge is limited to what was available at the time of my training in 2021. My ability to provide current information is limited, and I am not able to browse the internet for up-to-date information.

Act to reason

1. Obtain knowledge
2. Update knowledge

Hallucination due to limited / outdated knowledge

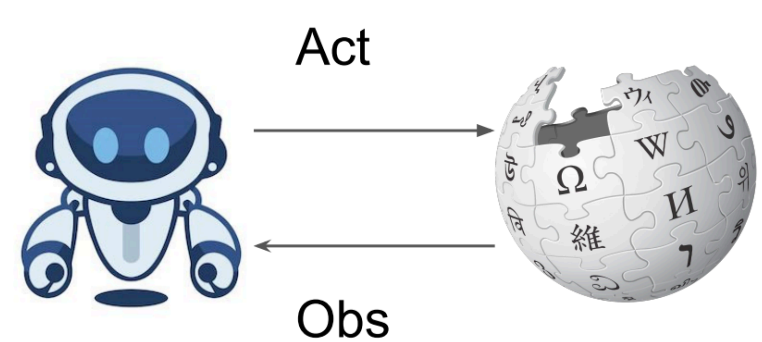

ReAct [Yao et al, 2022]

ReAct prompt

Question Which magazine was started first Arthur's Magazine or First for Women?
Thought 1 I need to search Arthur's Magazine and First for Women, and find which was started first.
Action 1 Search[Arthur's Magazine]
Observation 1 Arthur's Magazine (1844-1846) was an American literary periodical published in Philadelphia …
Thought 2 Arthur's Magazine was started in 1844. I need to search First for Women next.
Action 2 Search[First for Women]
Observation 2 First for Women is a woman's magazine published by Bauer Media Group … was started in 1989…
Thought 3 First for Women was started in 1989. 1844 (Arthur's Magazine) < 1989 (First for Women), so Arthur's Magazine was started first.
Action 3 Finish[Arthur's Magazine]

Reason-only prompt
(Chain-of-thought, CoT)

Question Which magazine was started first Arthur's Magazine or First for Women?
Thought Arthur's Magazine was started in 1844. First for Women was started in 1989. 1844 (Arthur's Magazine) < 1989 (First for Women), so Arthur's Magazine was started first.
Answer Arthur's Magazine

Act-only prompt
(similar to WebGPT)

Question Which magazine was started first Arthur's Magazine or First for Women?
Action 1 Search[Arthur's Magazine]
Observation 1 Arthur's Magazine (1844-1846) was an American literary periodical published in Philadelphia …
Action 2 Search[First for Women]
Observation 2 First for Women is a woman's magazine published by Bauer Media Group … was started in 1989…
Action 3 Finish[Arthur's Magazine]

Standard prompt

Question Which magazine was started first Arthur's Magazine or First for Women?
Answer Arthur's Magazine

32

Yao et al., 2022. ReAct: Synergizing Reasoning and Acting in Language Models

Reasoning without acting

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

Standard

Answer: iPod ✗

Reasoning without acting

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

Reasoning with acting

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

ReAct (Reason + Act)

Thought 1: I need to search Apple Remote and find the program it was originally designed to interact with.
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 ……(long text omitted)

Reasoning with acting

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

ReAct (Reason + Act)

Thought 1: I need to search Apple Remote and find the program it was originally designed to interact with.
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 ……(long text omitted)

Thought 2: Apple Remote was originally designed to control the Front Row media center program. I need to search Front Row next and find what other device can control it.
Act 2: **Search**[Front Row]
Obs 2: Could not find [Front Row]. Similar: ['Front Row Seat to Earth', 'Front Row Motorsports','Front Row (software)', …]

Thought 3: Front Row is not found. I need to searchFront Row (software) .
Act 3: **Search**[Front Row (software)]
Obs 3: Front Row is a discontinued media center software (long text omitted)

Thought 4: Front Row (software) is controlled by an Apple Remote or the keyboard function keysSo the answer is keyboard function keys.
Act 4: **Finish**[keyboard function keys]

Acting without reasoning

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]

✗

Acting in language

Embodied Env
(ALFRED game)

```
You are in the middle of a room. Looking quickly around you, you see a drawer 2, a shelf
5, a drawer 1, a shelf 4, a sidetable 1, a drawer 5, a shelf 6, a shelf 1, a shelf 9, a
cabinet 2, a sofa 1, a cabinet 1, a shelf 3, a cabinet 3, a drawer 3, a shelf 11, a shelf
2, a shelf 10, a dresser 1, a shelf 12, a garbagecan 1, a armchair 1, a cabinet 4, a shelf
7, a shelf 8, a safe 1, and a drawer 4.

Your task is to: put some vase in safe.

> go to shelf 6
You arrive at loc 4. On the shelf 6, you see a vase 2.

> take vase 2 from shelf 6
You pick up the vase 2 from the shelf 6.

> go to safe 1
You arrive at loc 3. The safe 1 is closed.

> open safe 1
You open the safe 1. The safe 1 is open. In it, you see a keychain 3.

> put vase 2 in/on safe 1
You won!
```

Text Env
(ALFWorld)

Shridhar et al. "ALFWorld: Aligning Text and Embodied Environments for Interactive Learning. "

The flexibility of “acting”

- **Tool usage**: use external tools as action

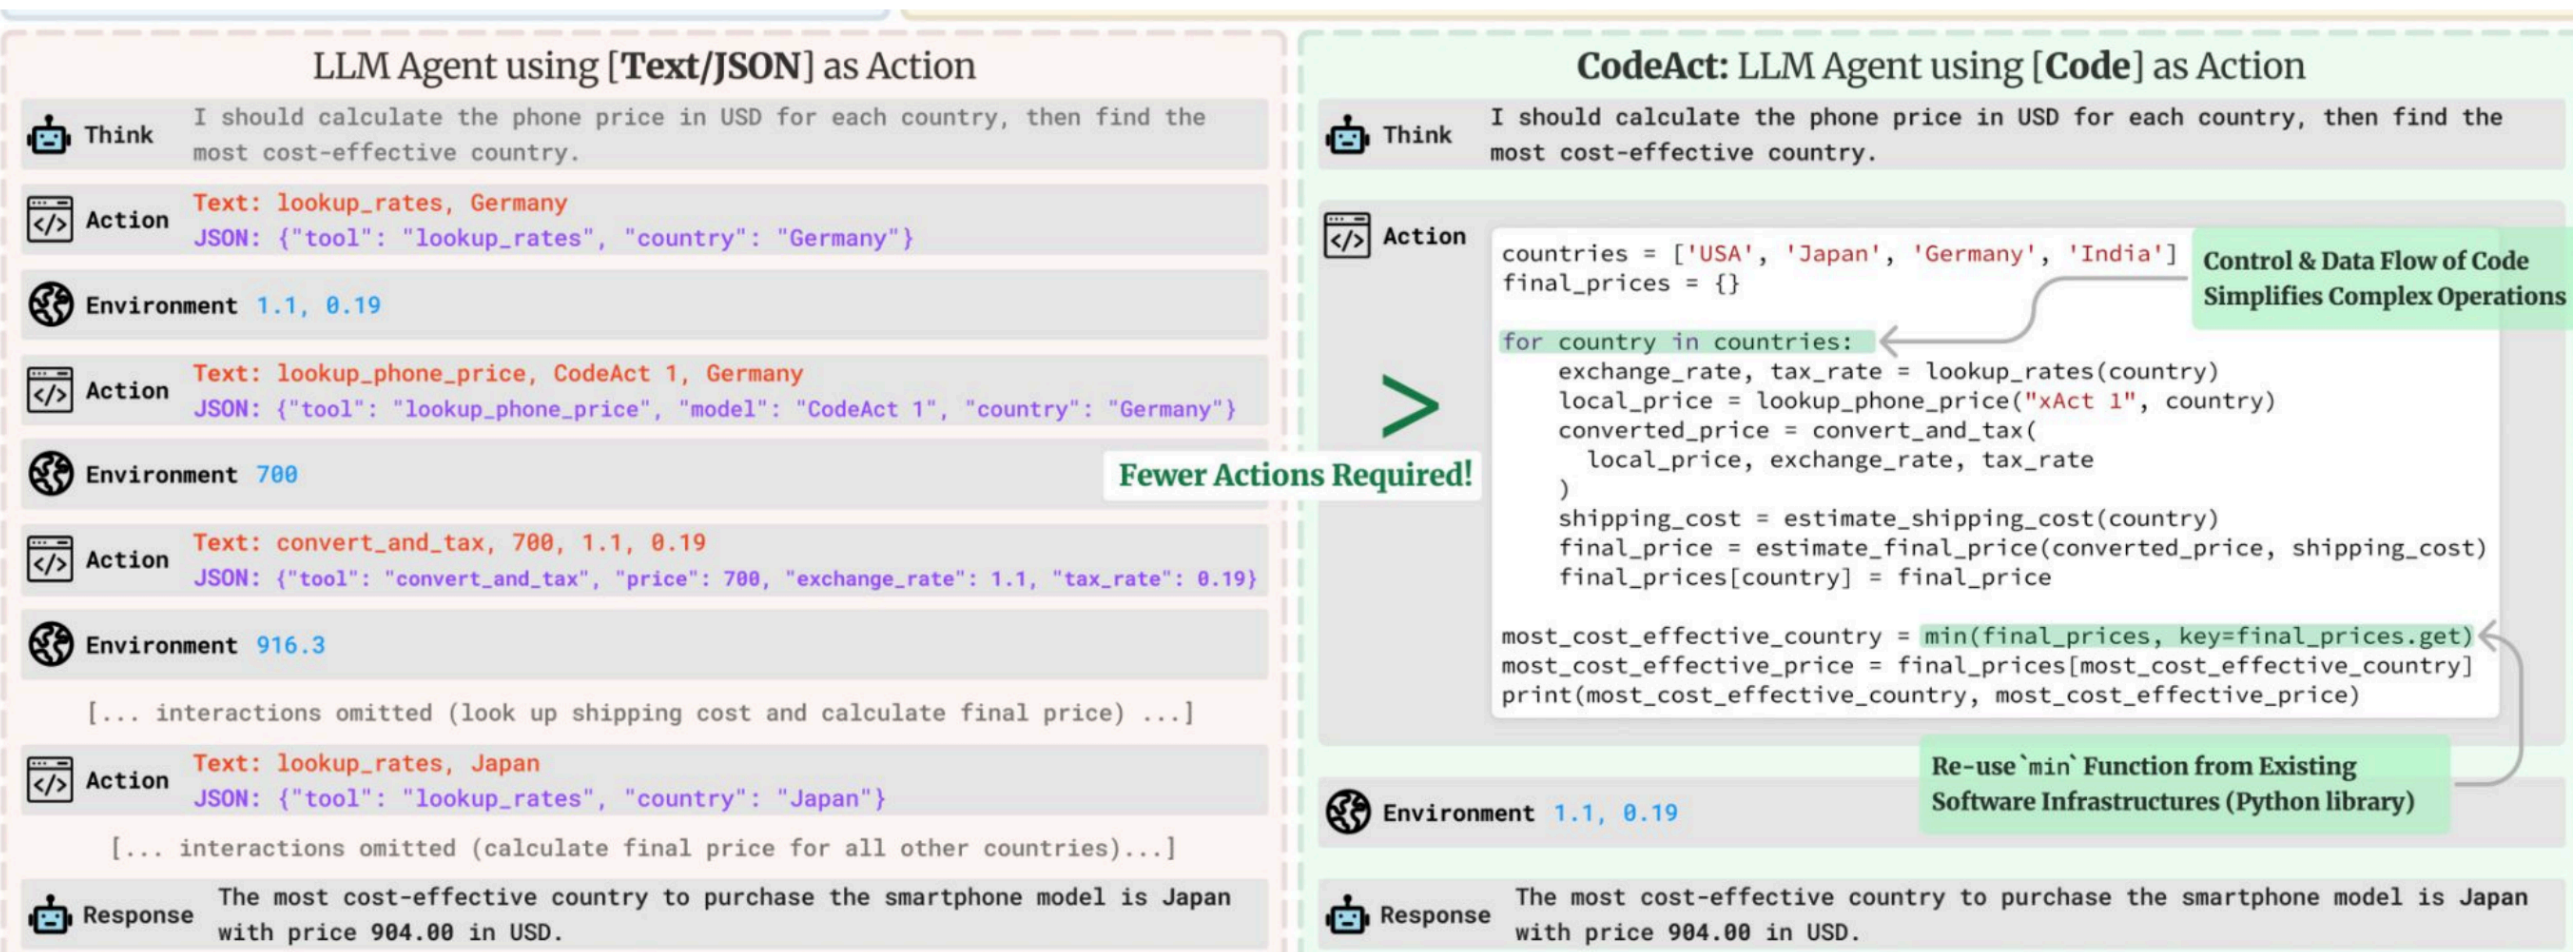

Wang, et al., ICLR 2024. Executable Code Actions Elicit Better LLM Agents.

The flexibility of “acting”

- **Tool usage**: use external tools as action

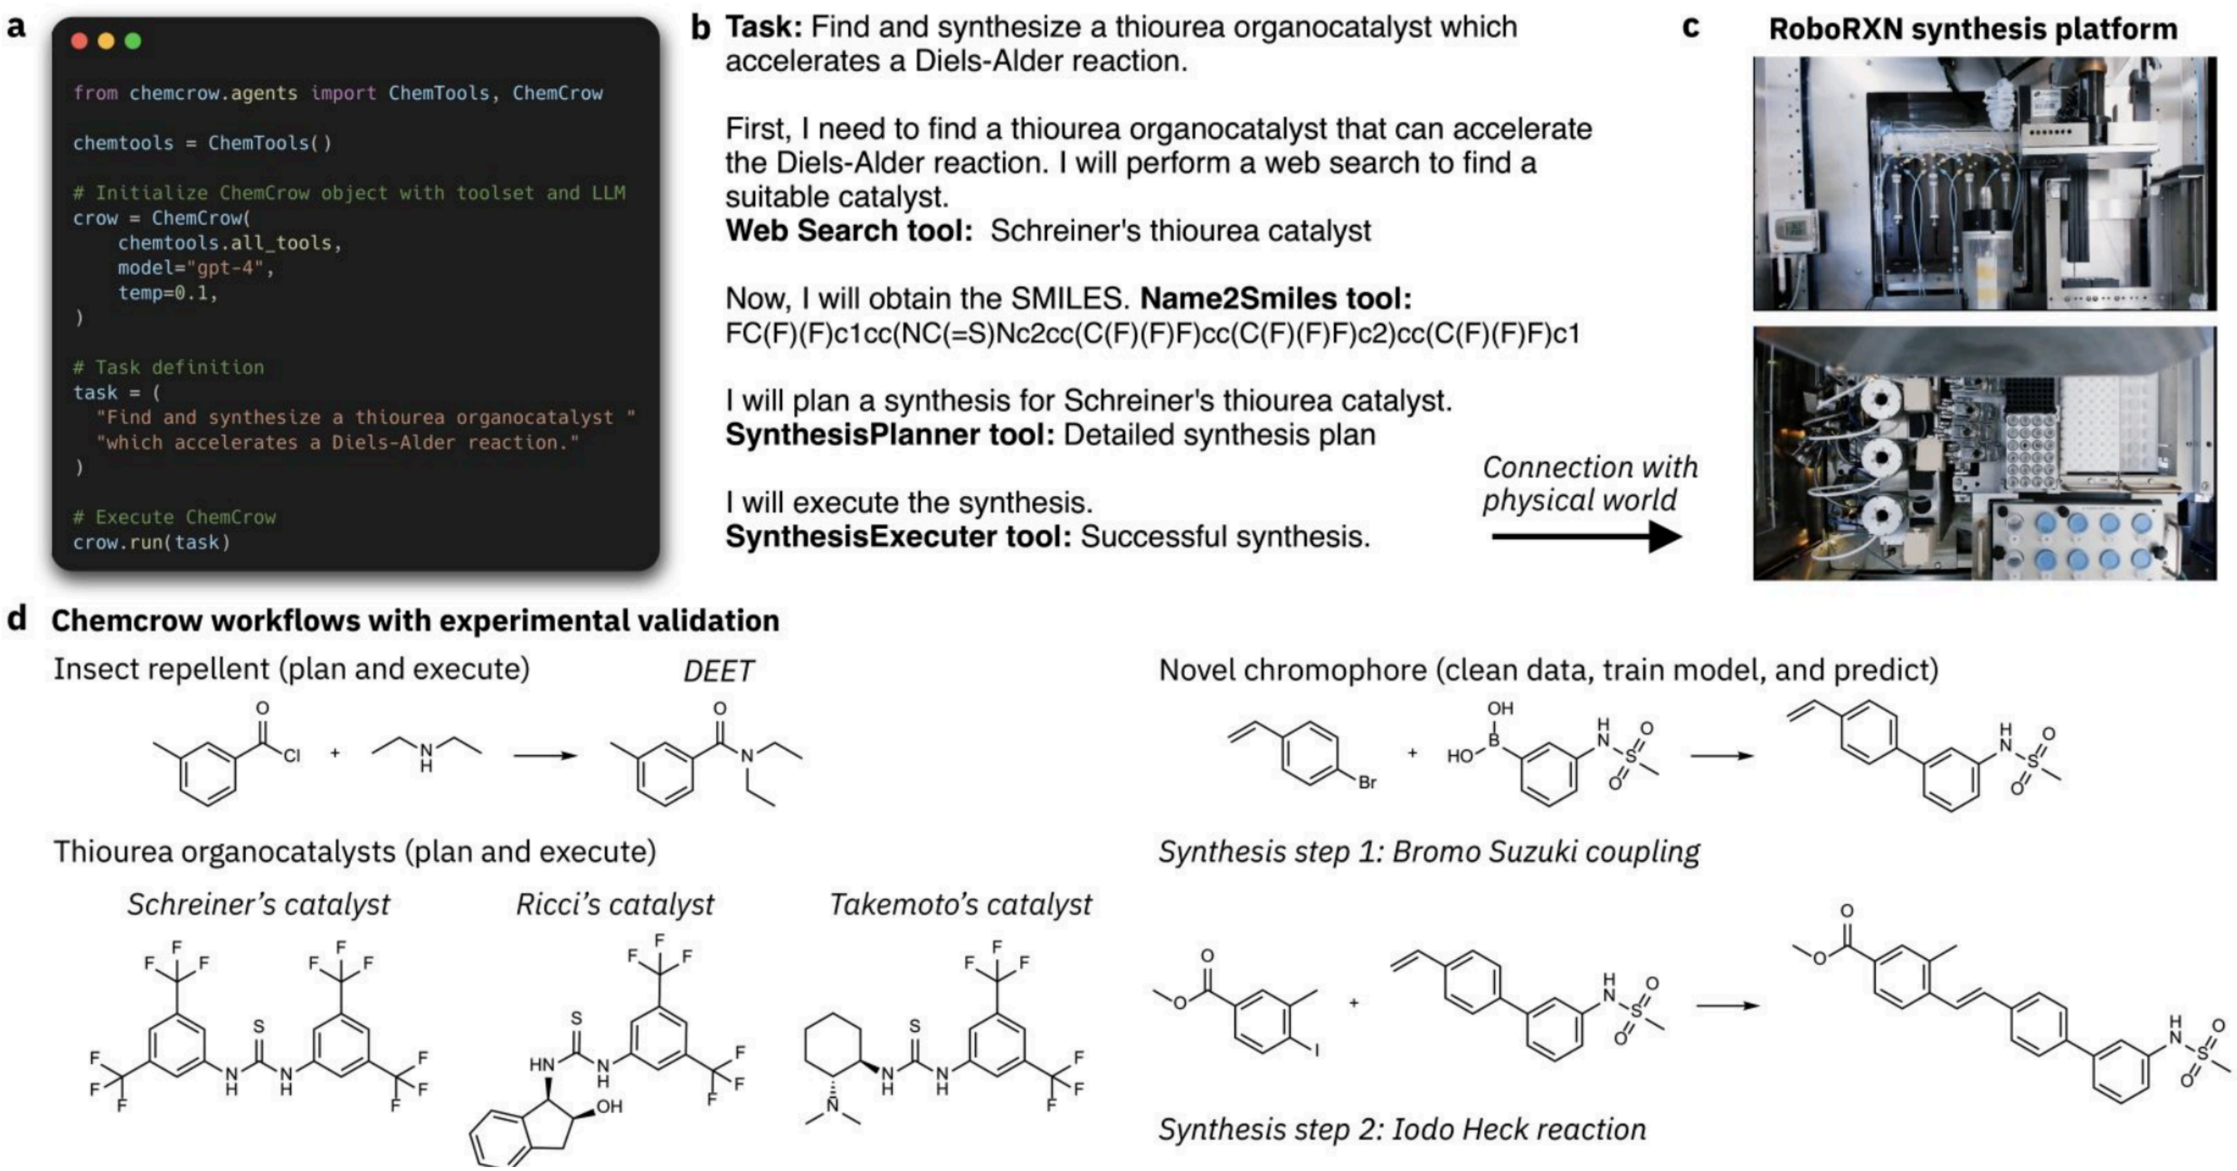

Bran, et al., 2023. ChemCrow: Augmenting large-language models with chemistry tools

Why is reasoning special for agents?

Traditional agents: action space A defined by the environment

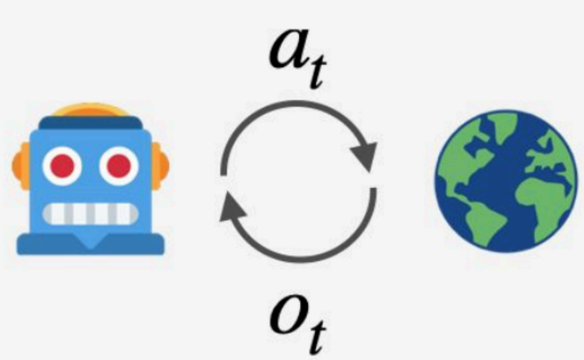

- **External feedback** o_t
- Agent context $c_t = (o_1, a_1, o_2, a_2, \cdots, o_t)$
- Agent action $a_t \sim \pi(a \mid c_t) \in A$

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

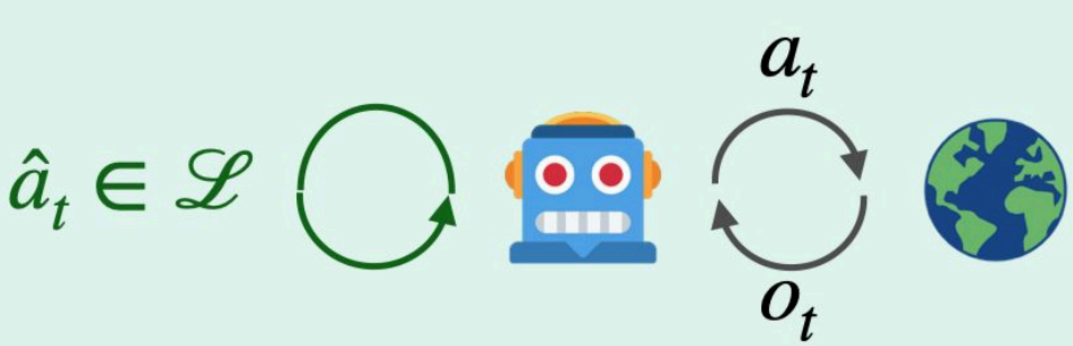

- $\hat{a}_t \in \mathcal{L}$ can be any language sequence
- Agent context $c_{t+1} = (c_t, \hat{a}_t, a_t, o_{t+1})$
- $\hat{a}_t \in \mathcal{L}$ only updates **internal context**

Why is reasoning just now for agents?

- Bigger action space -> 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 $\hat{A} = A \cup \mathscr{L}$ augmented by reasoning

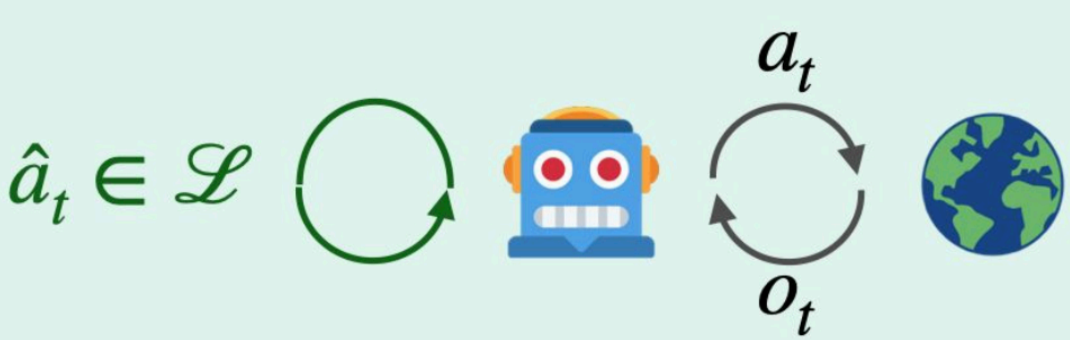

- $\hat{a}_t \in \mathscr{L}$ can be any language sequence
- Agent context $c_{t+1} = (c_t, \hat{a}_t, a_t, o_{t+1})$
- $\hat{a}_t \in \mathscr{L}$ only updates **internal context**

Reasoning: Takeaways

- Reasoning as internal actions for language agents
- Reasoning guides acting & acting updates reasoning

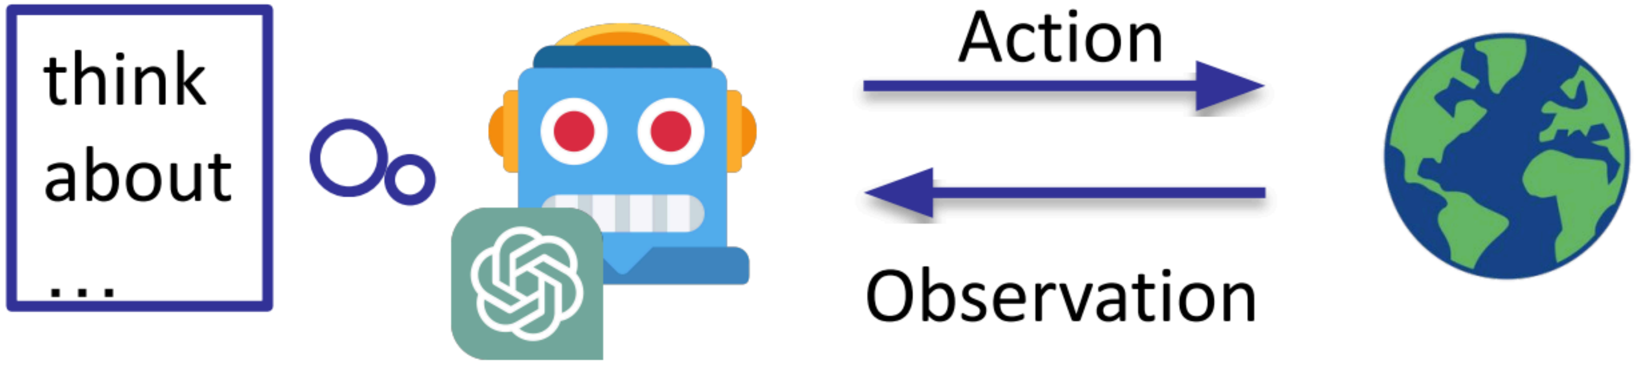

Memory

Memory is everything. Without it, we are nothing.

— Eric Kandel

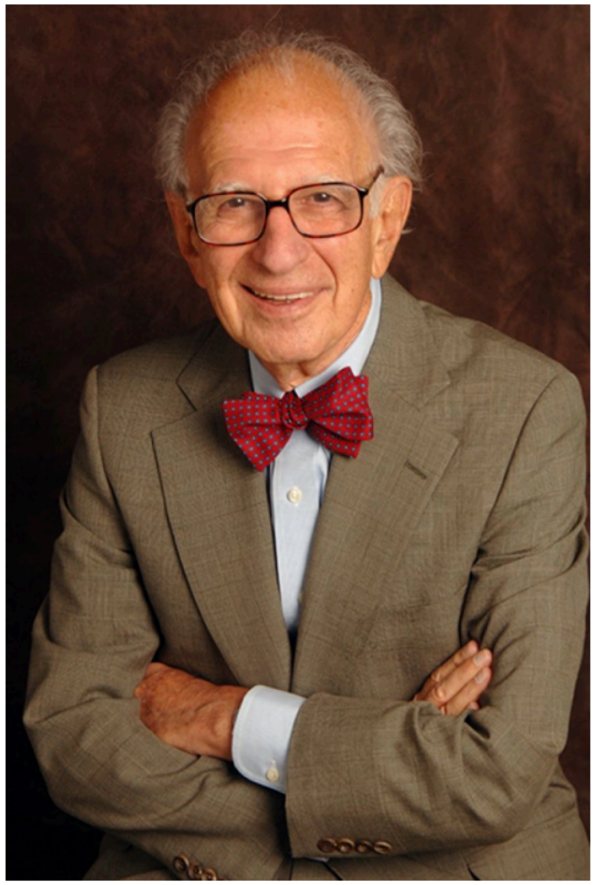

Memory

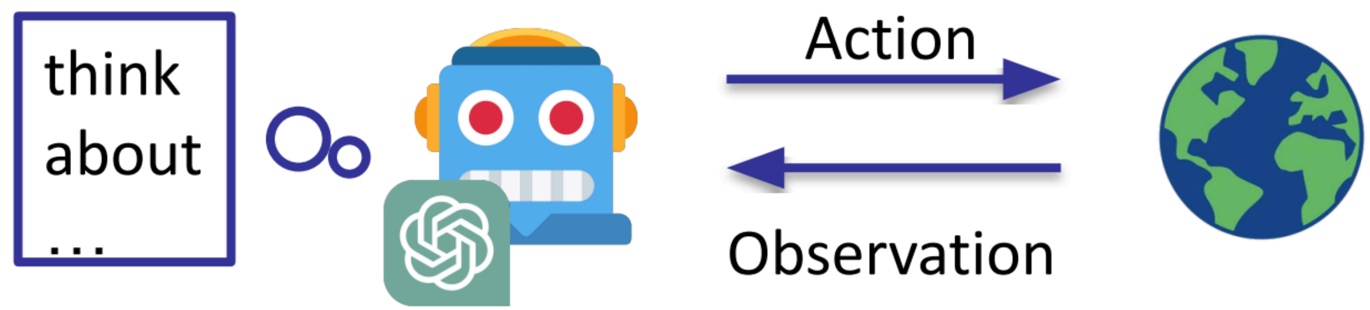

Memory

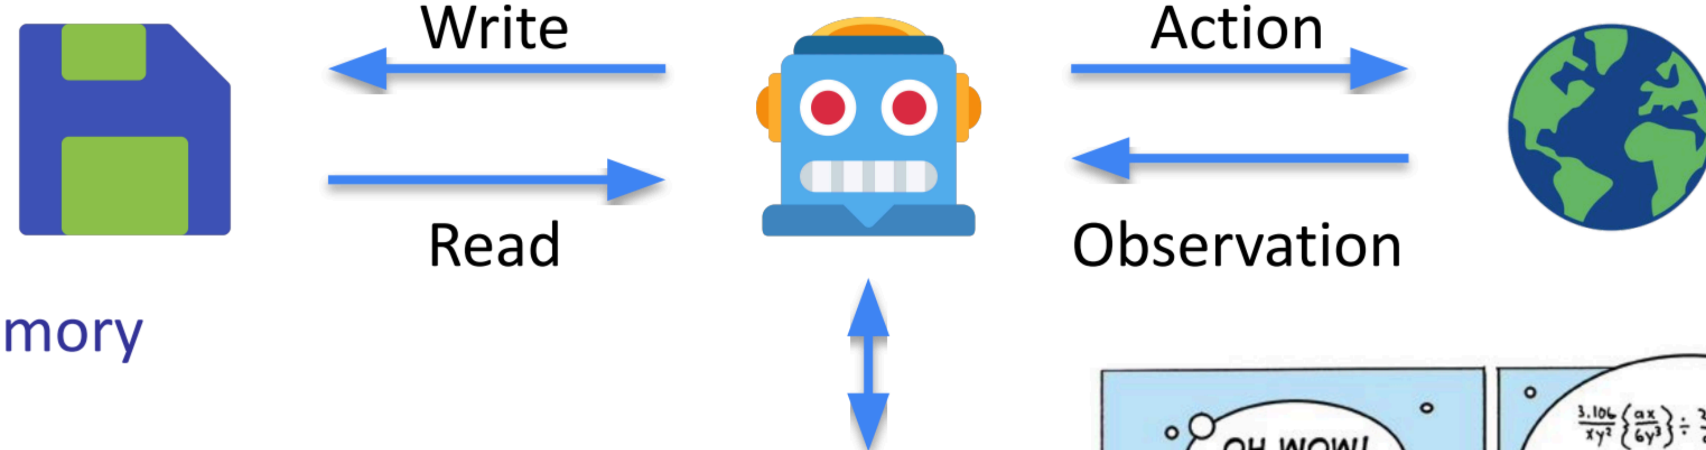

A long-term memory

- Read and write
- Stores experience, knowledge, skills, …
- Persist over new experience

Instruction: …
Thought: …
Action: …
Obs: …
Thought: …
……

A short-term memory

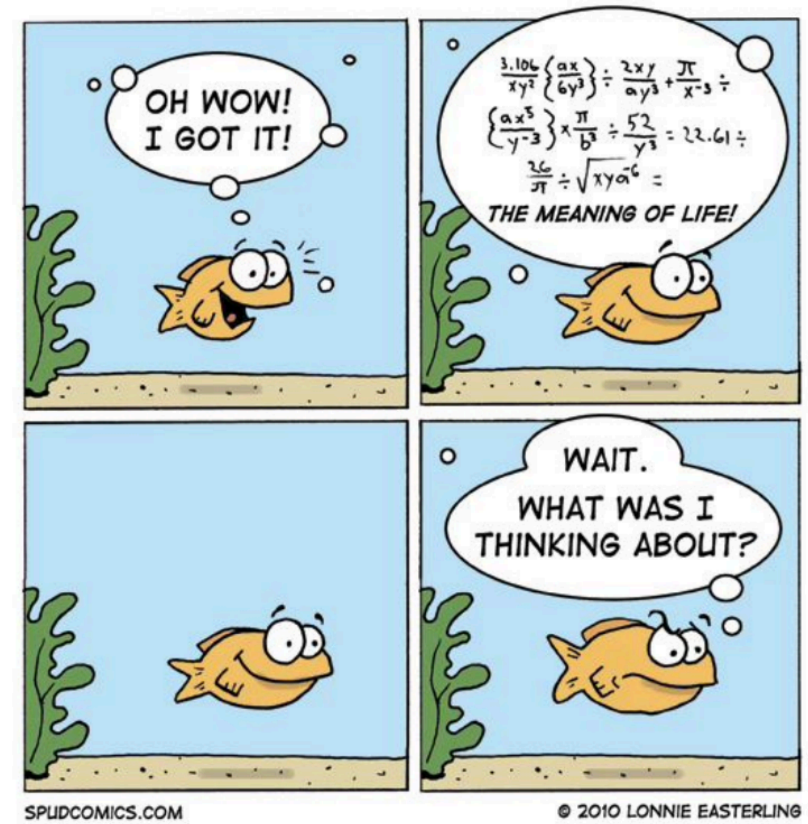

THE TRAGEDY OF A THREE SECOND MEMORY

What about retrieval and RAG?

- We can think of the retrieval corpus as “read-only” long-term memory
  - Written by others (e.g., Wikipedia editors), not the agent itself
- Limitations
  - Can only live “others’ experience”, which might not be optimal for the agent
  - The way corpus is written might not be optimal for agent usage
- Agent memory: also be able to autonomously write to it!

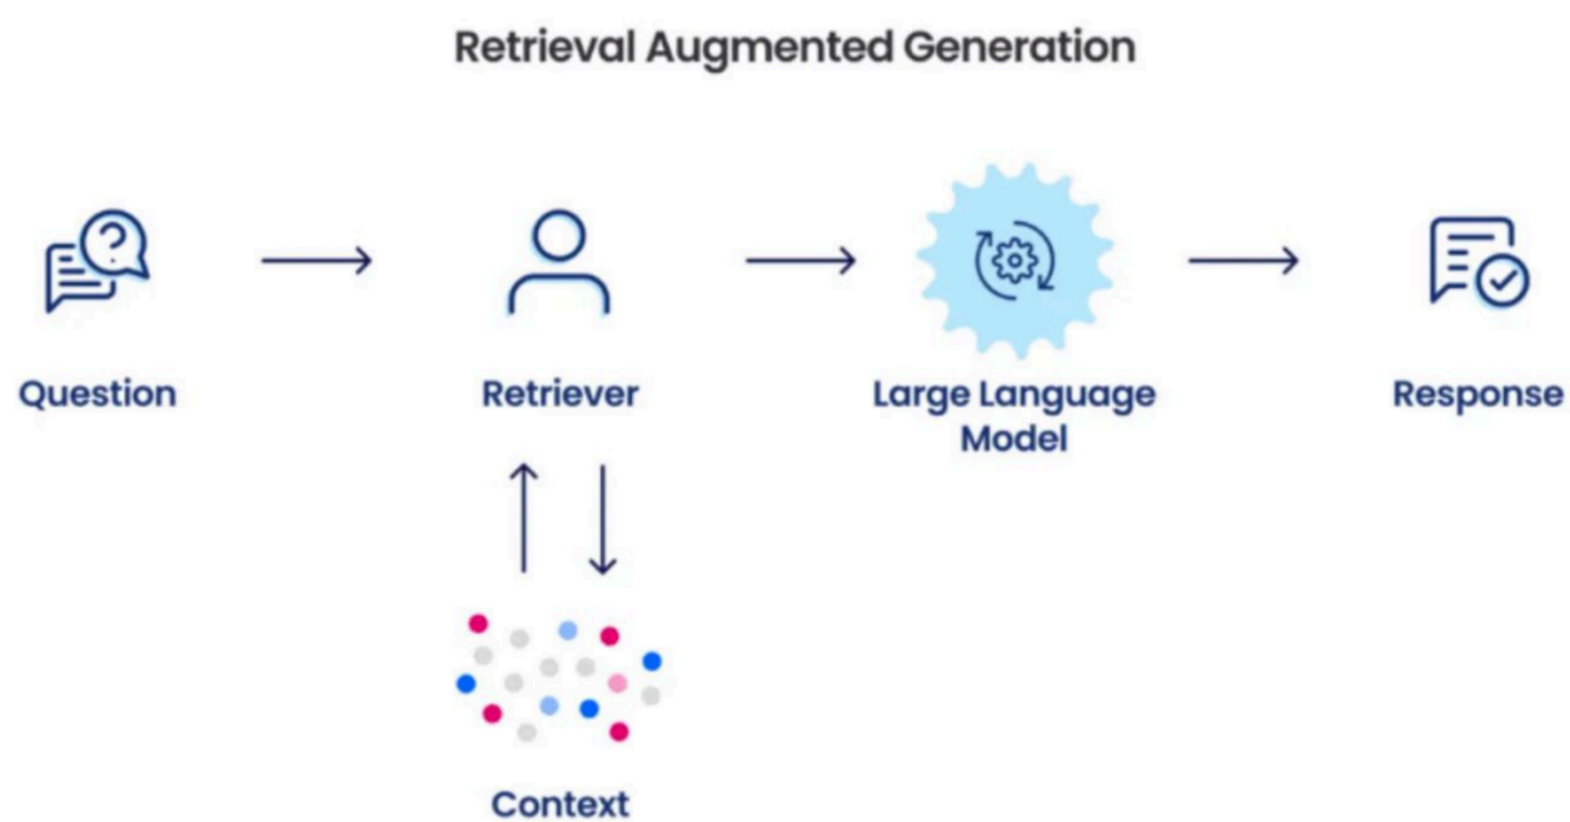

Long-term memory (LTM): Content

- What content is stored in LTM?
- Note: here we categorize three LTMs based on memory content

Type by content	Definition	Examples
Episodic memory	Stores experience	Generative agents [Park et al., 2023]
Semantic memory	Stores knowledge	
Procedural memory	Stores skills	Voyager [Wang et al., 2023]

Episodic memory

- Generative agents for social simulations

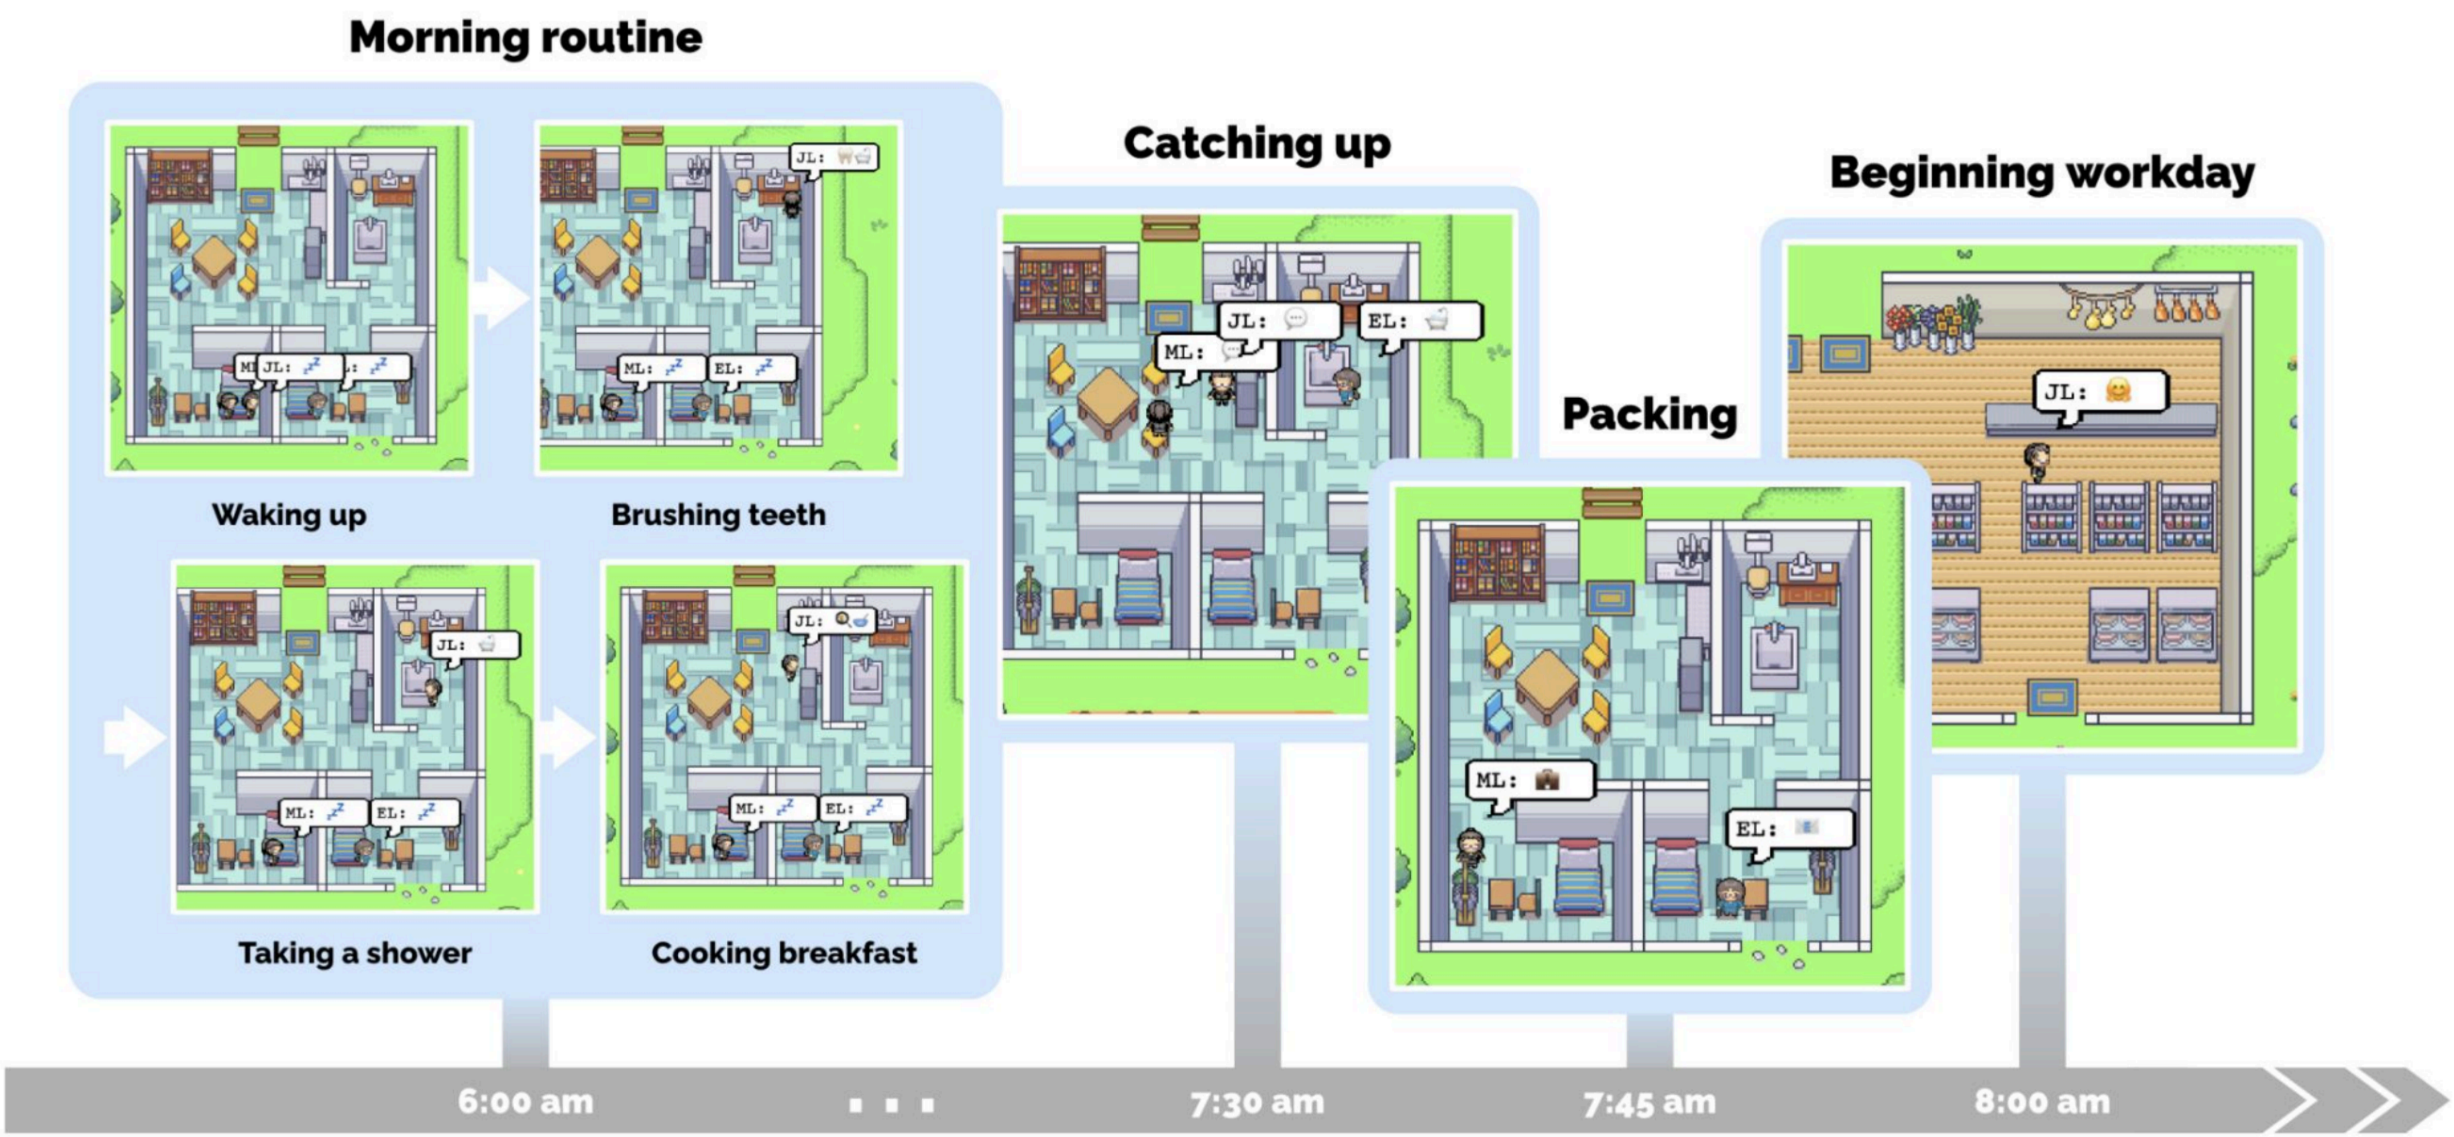

Episodic memory

- Write: append-only event streams
- Read: retrieval based on heuristic scores

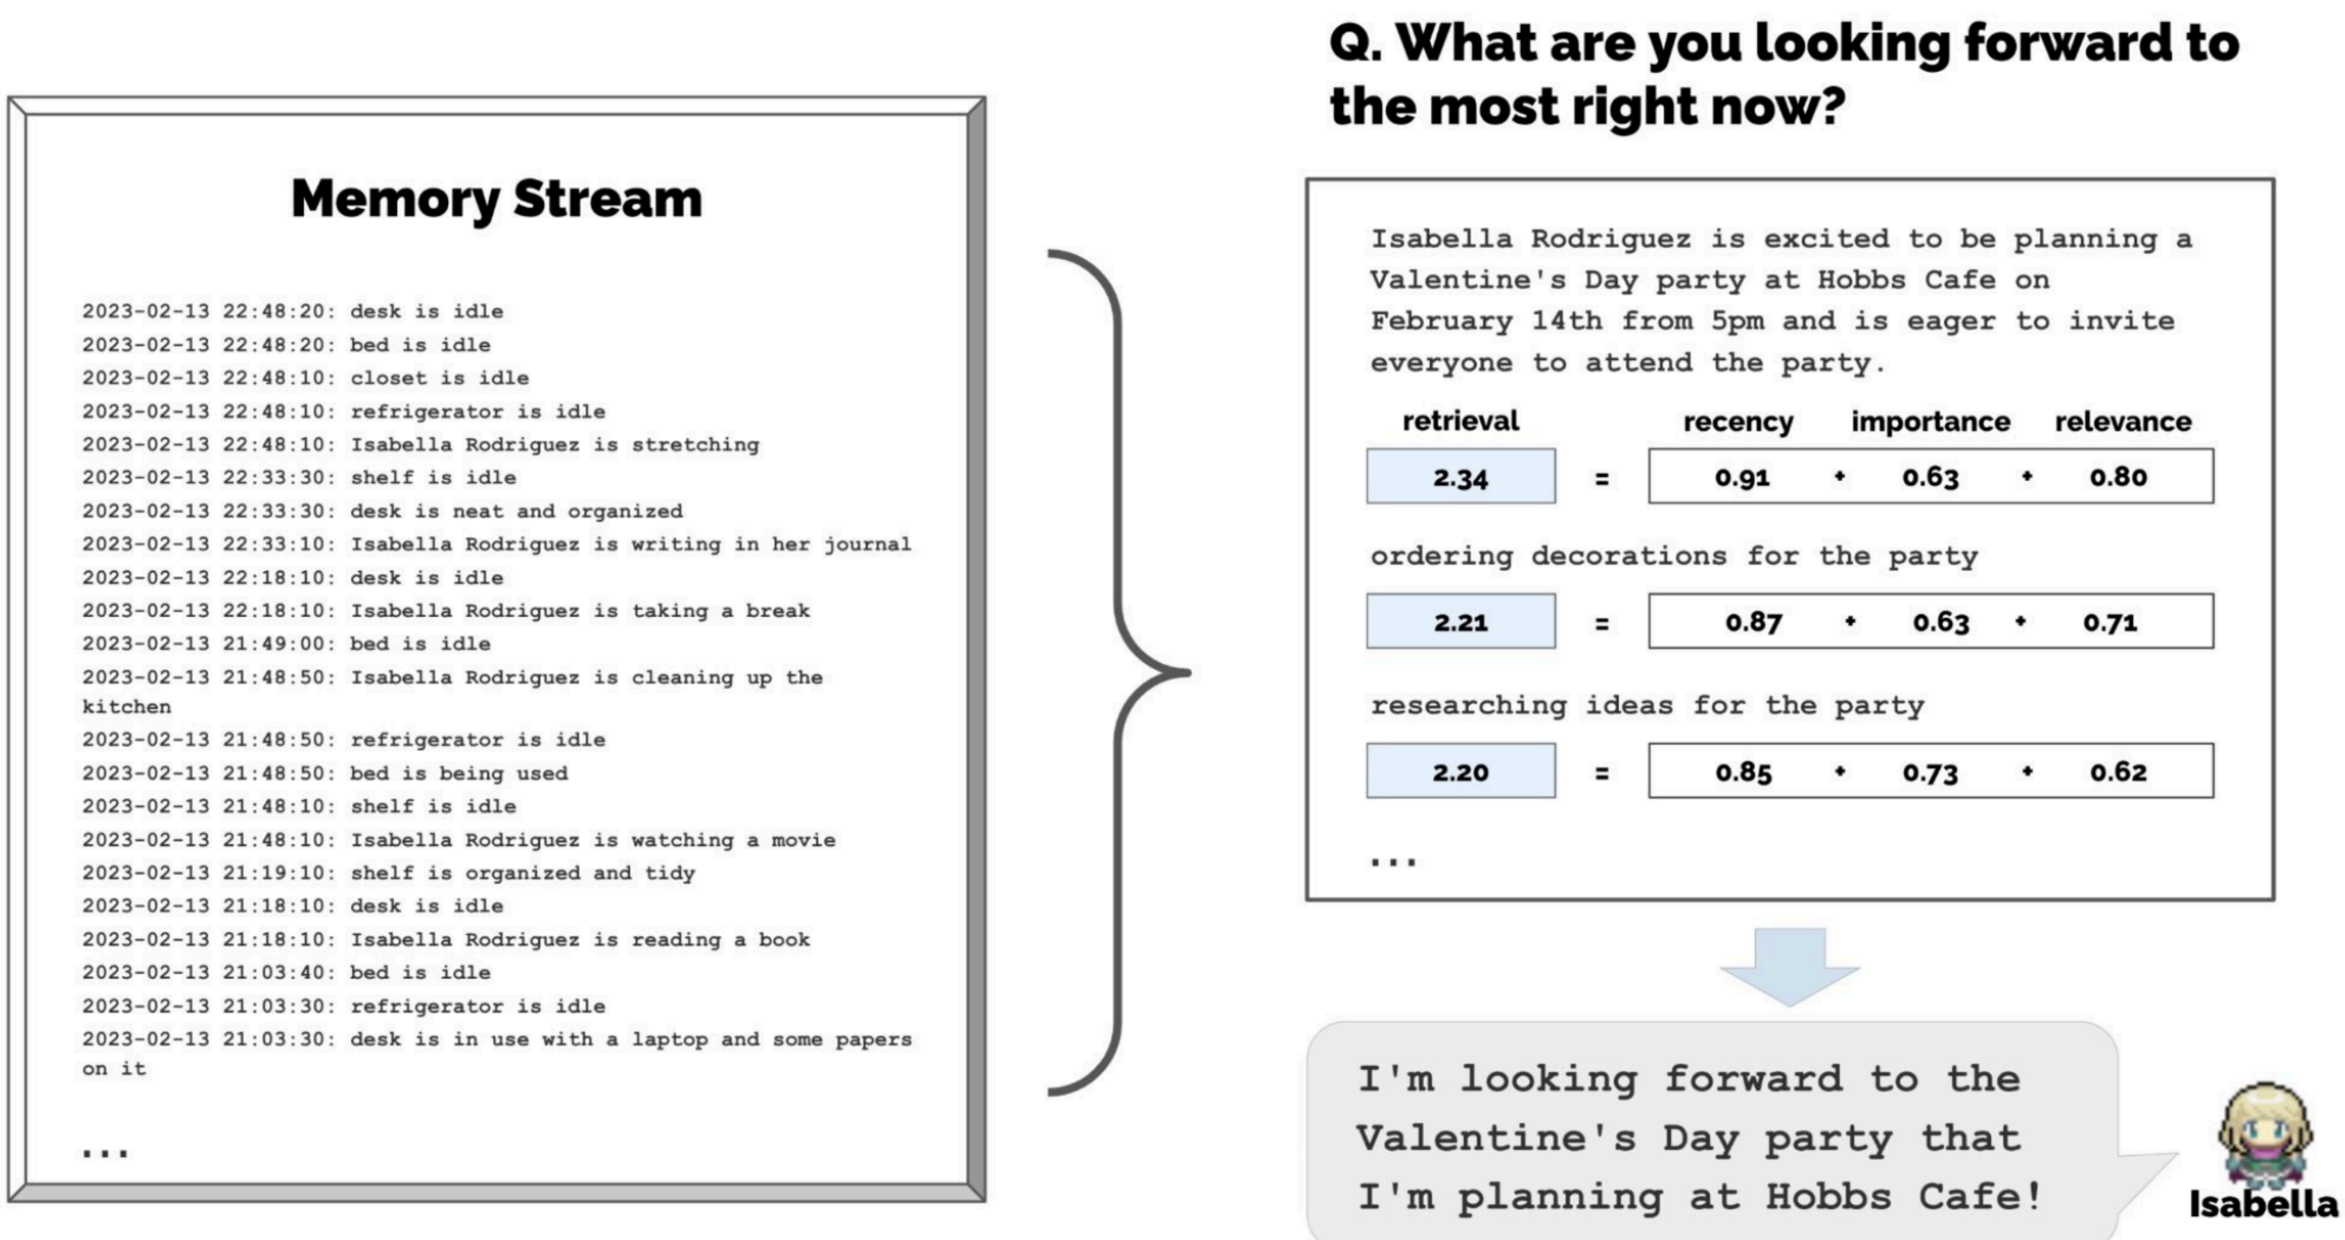

Semantic memory

- Write: LLM reasoning over events
- Read: retrieval

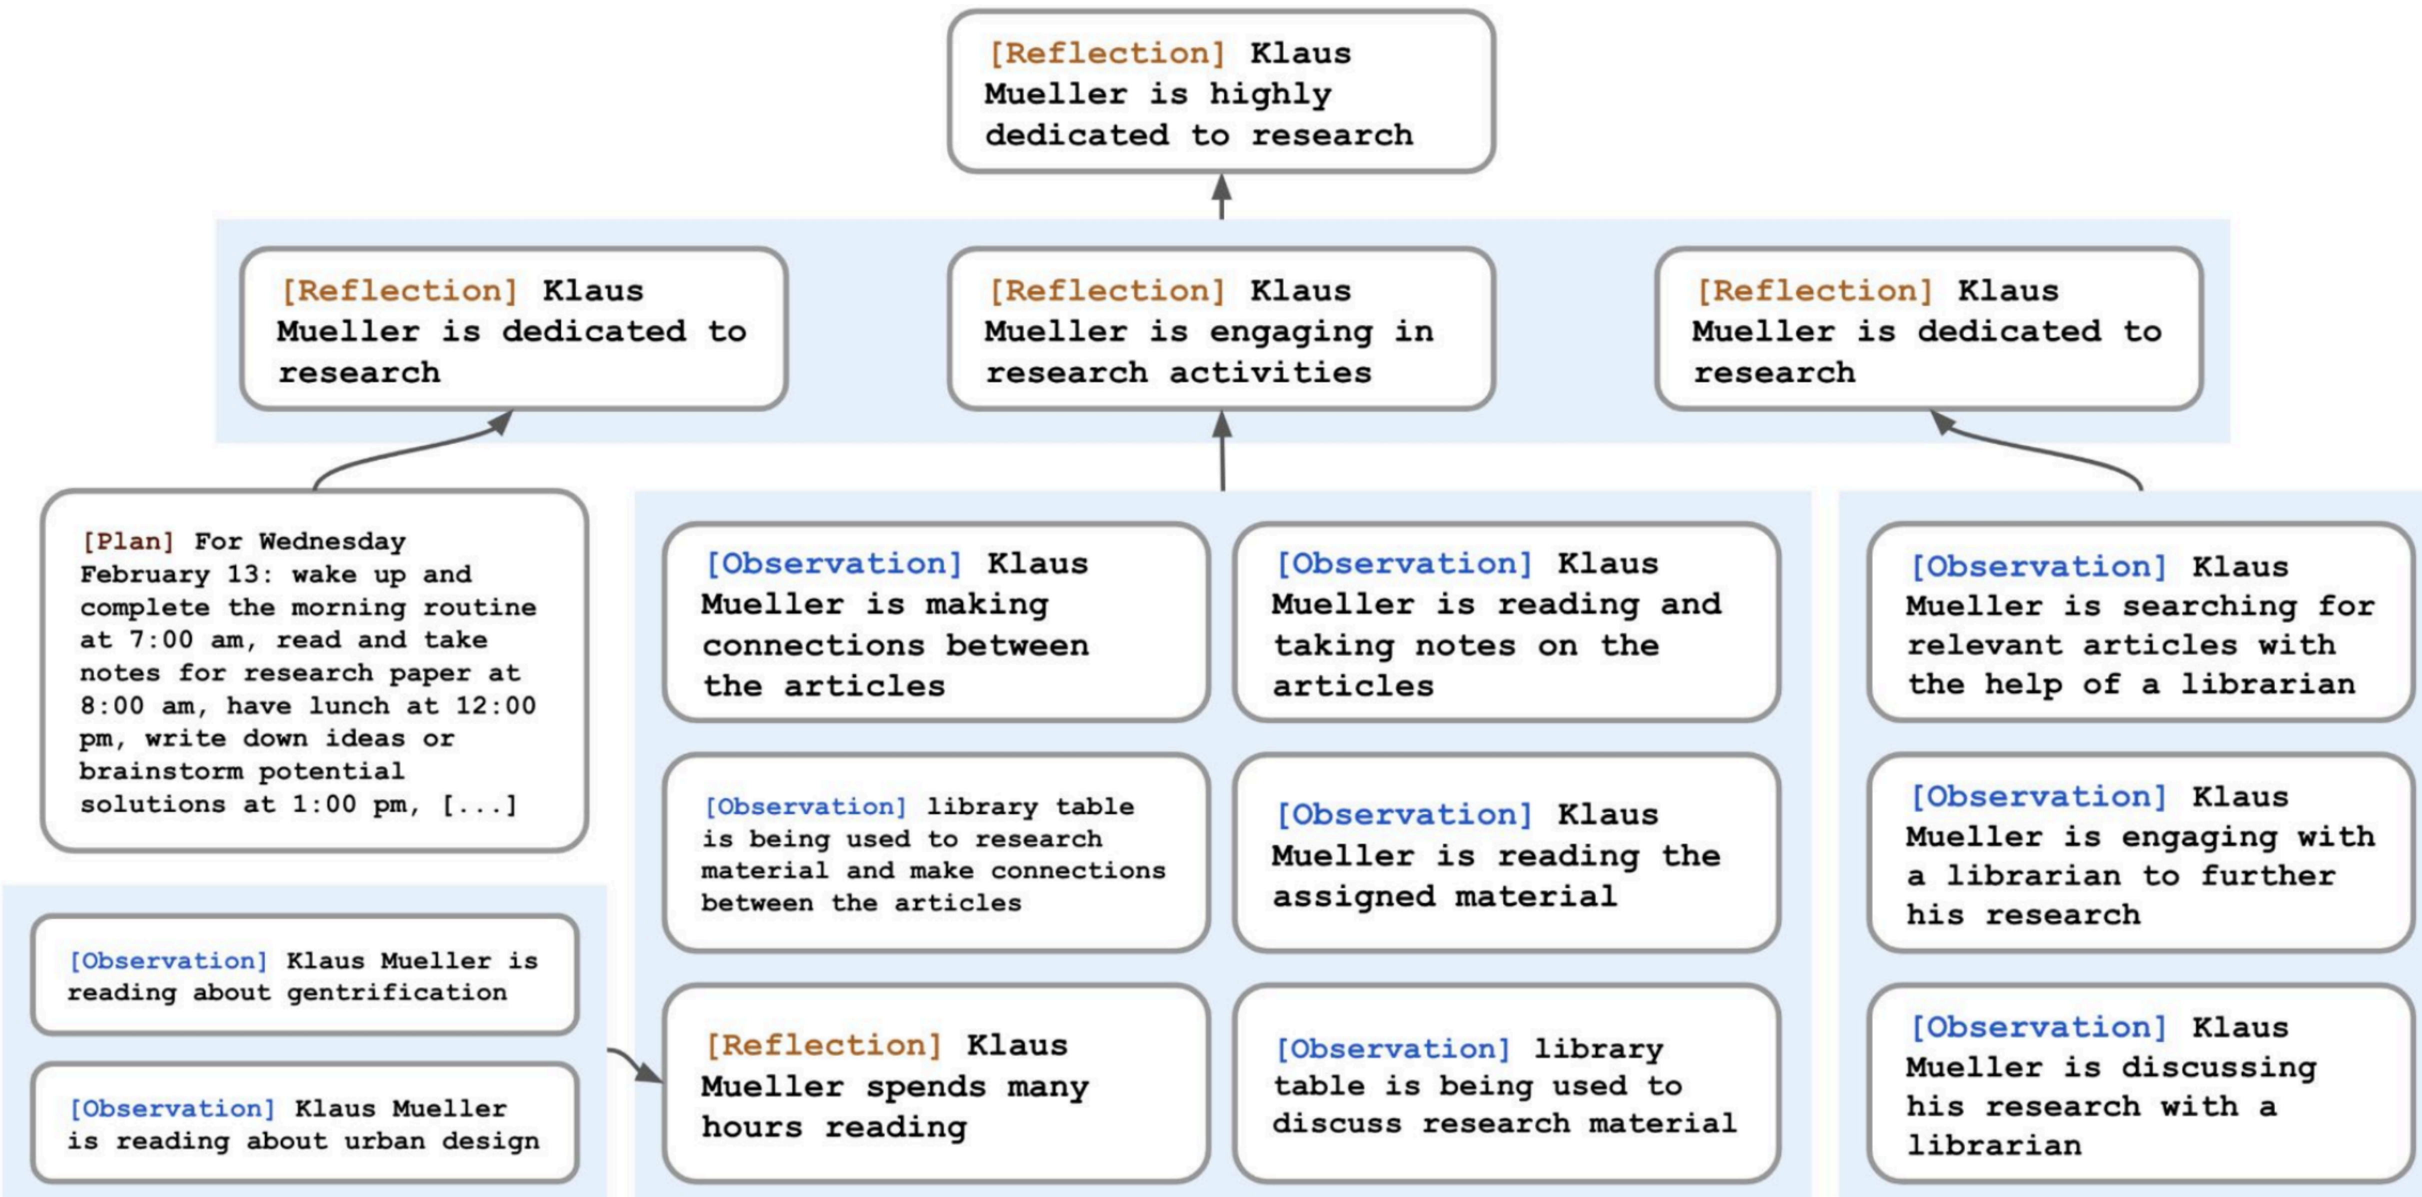

Procedural memory

- Write: Coding-based skills
- Read: embedding retrieval

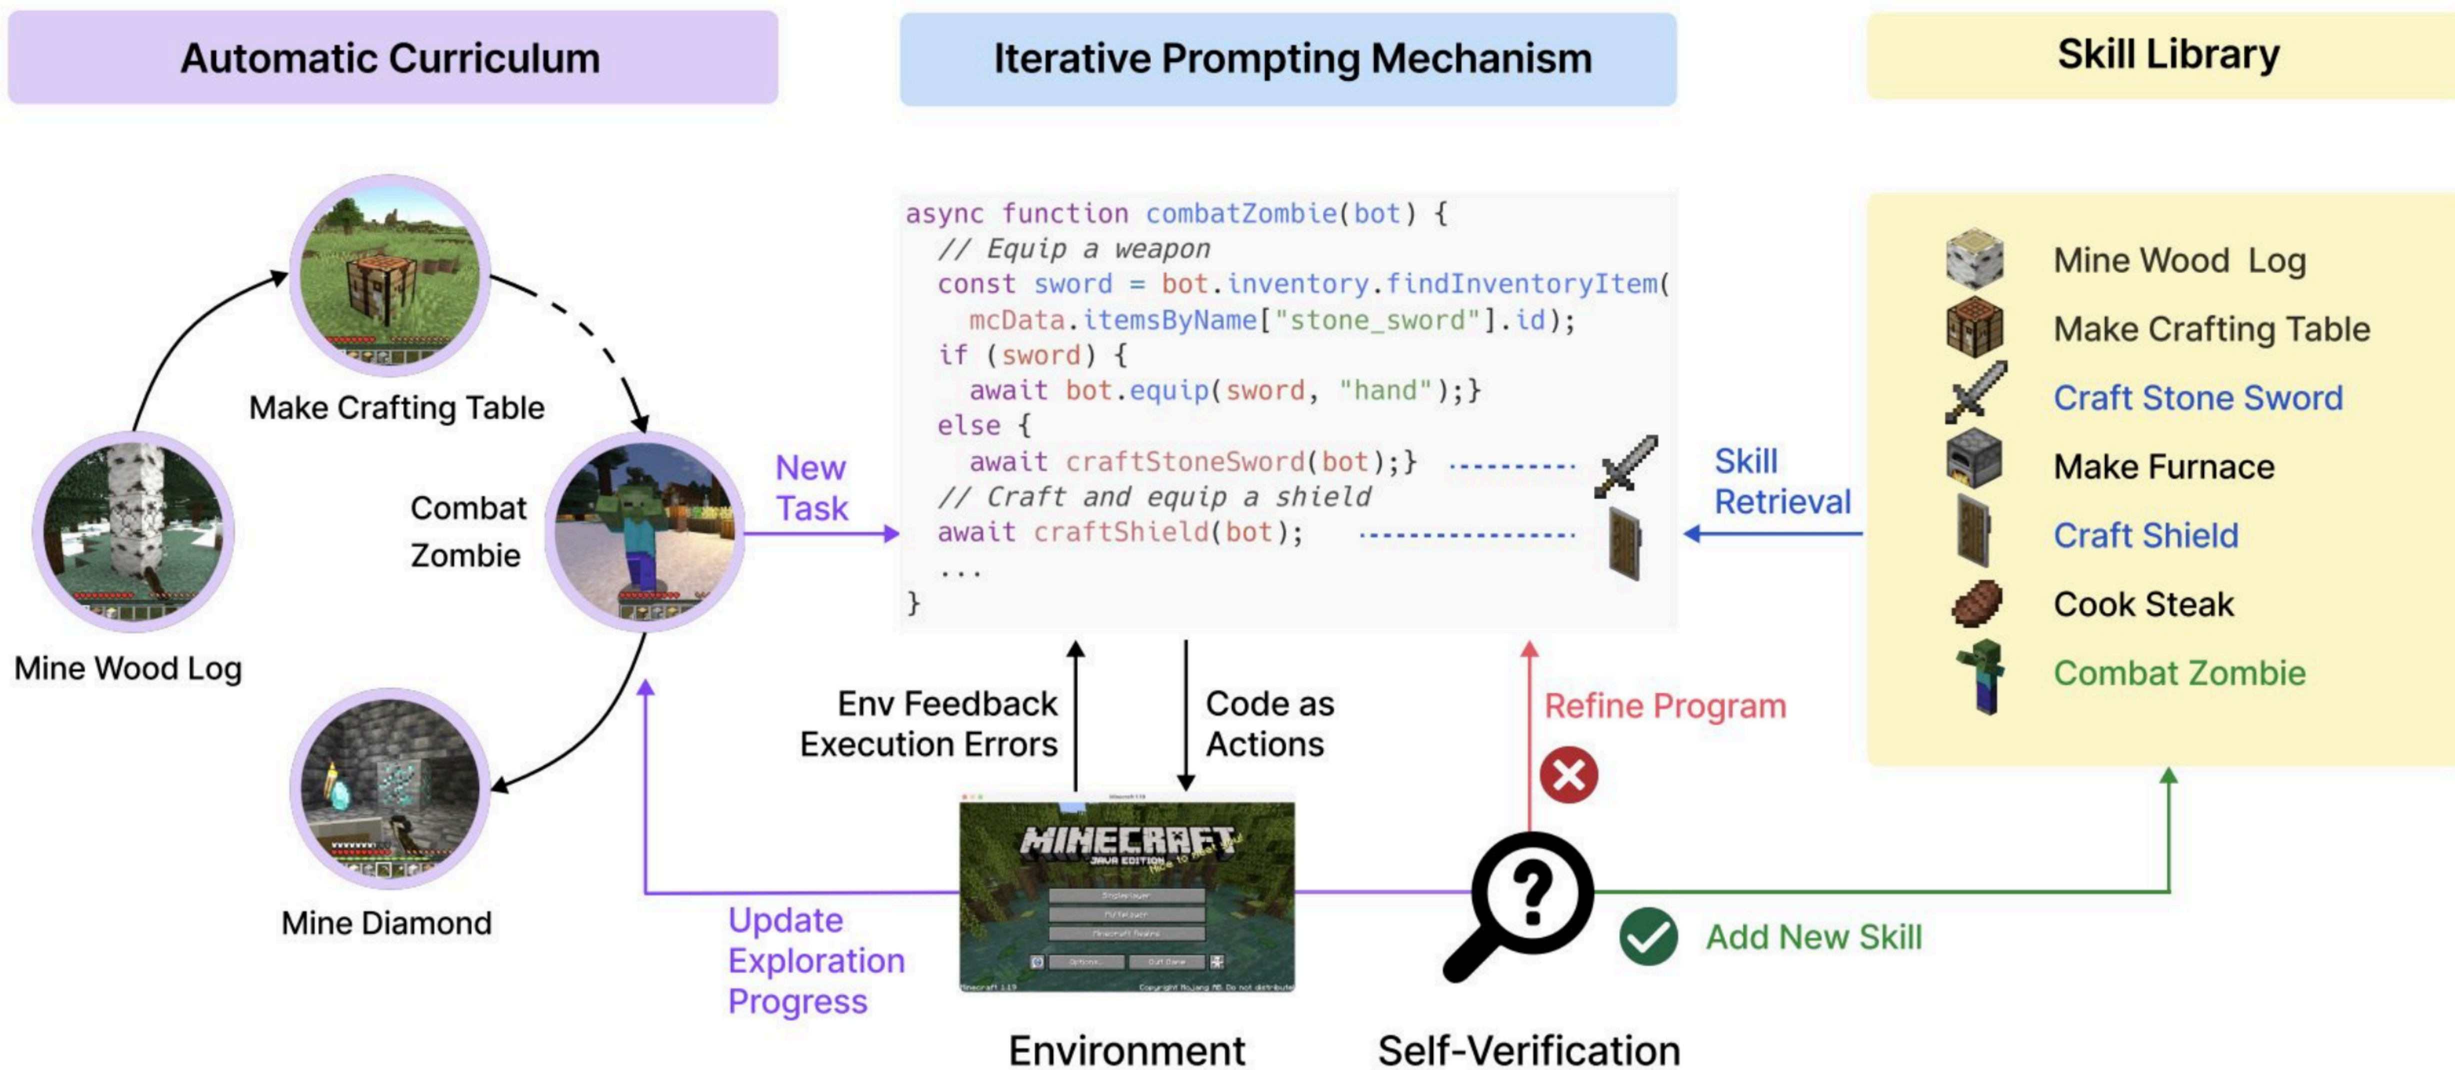

Wang, et al., 2023. Voyager: An Open-Ended Embodied Agent with Large Language Models

Exercise

- Where does a cognitive language agent store long-term memory?

- Where does a agent store short-term memory?

- What's the difference between external environment vs internal memory then?

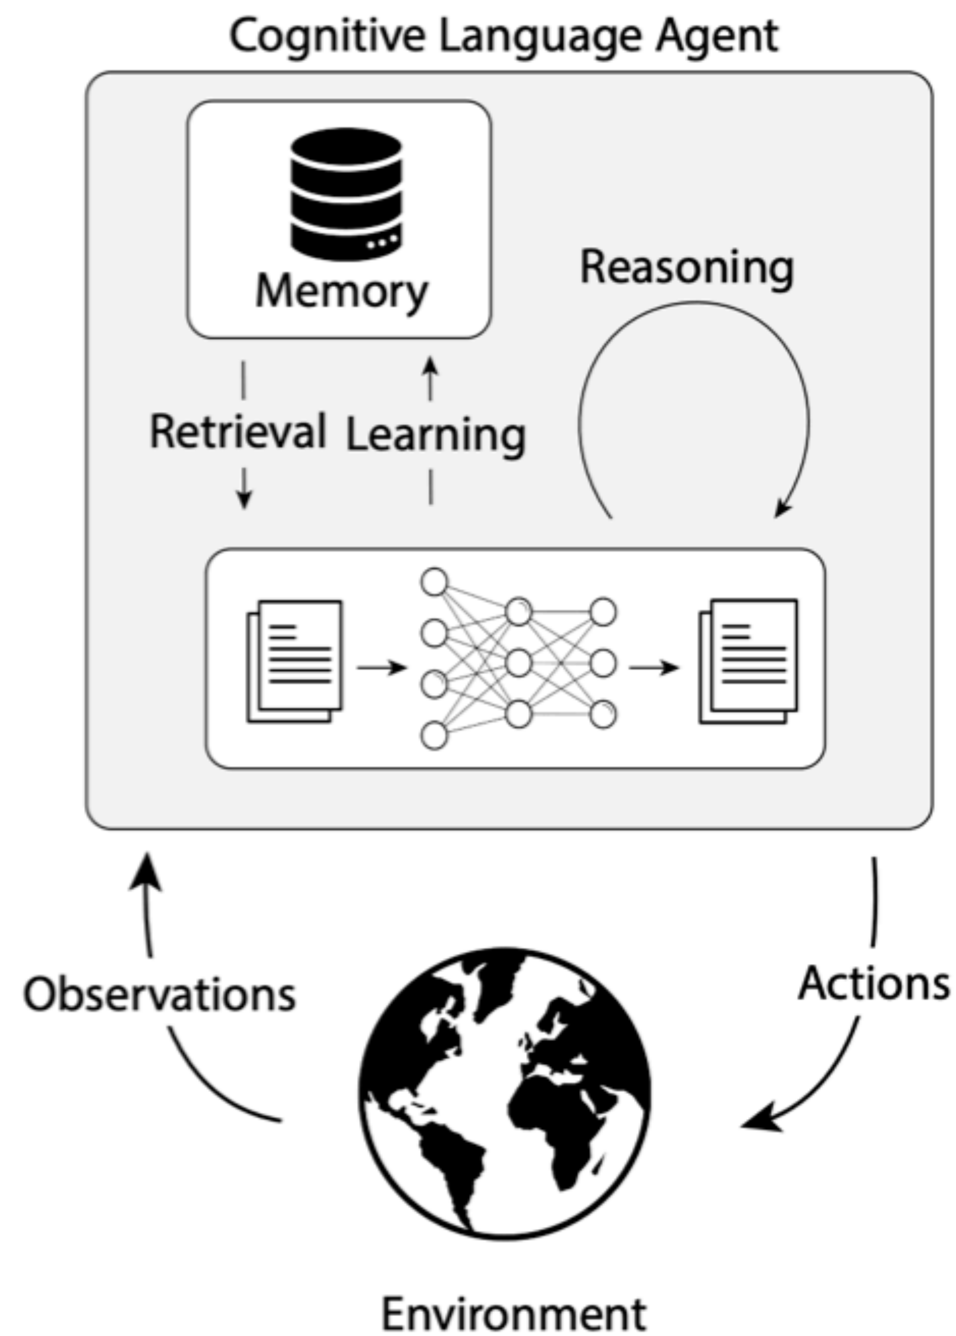

Exercise

- Where does a cognitive language agent store long-term memory?
  - External database
  - LLM's parameters

- Where does a agent store short-term memory?
  - Prompt

- What's the difference between external environment vs internal memory then?
  - Memory contains agent-specific contents (experiences, knowledge, skills)
  - Environment is independent of all agents

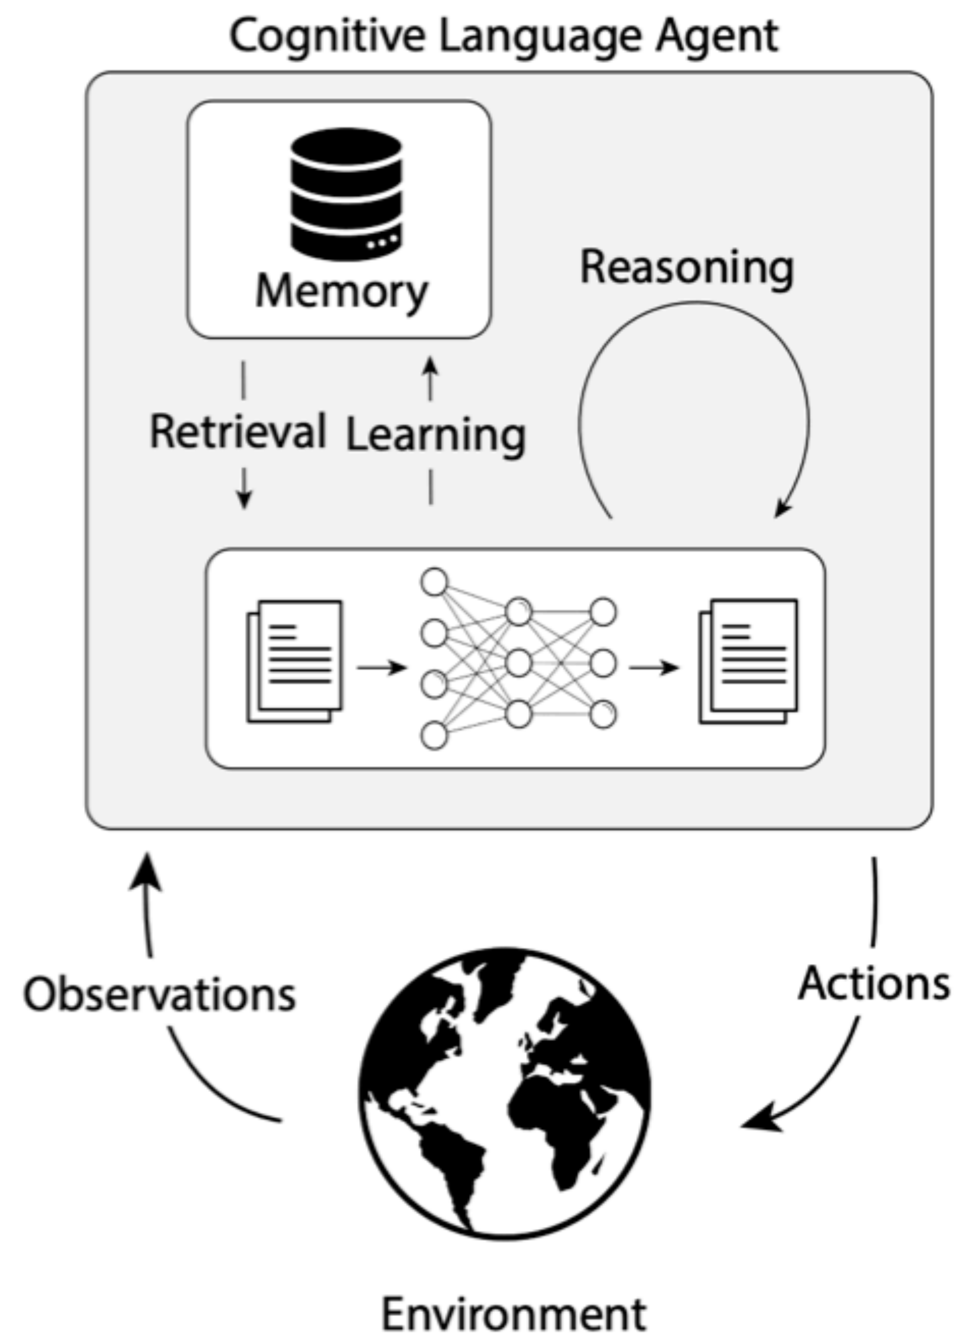

Planning: (simplified) definition

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

- General trends in planning settings for language agents
  - Increasing expressiveness in **goal specification**, e.g., in natural language as opposed to formal language
  - Substantially expanded or open-ended **action space**
  - Increasing difficulty in automated **goal test**

Language agent planning: web agents

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:

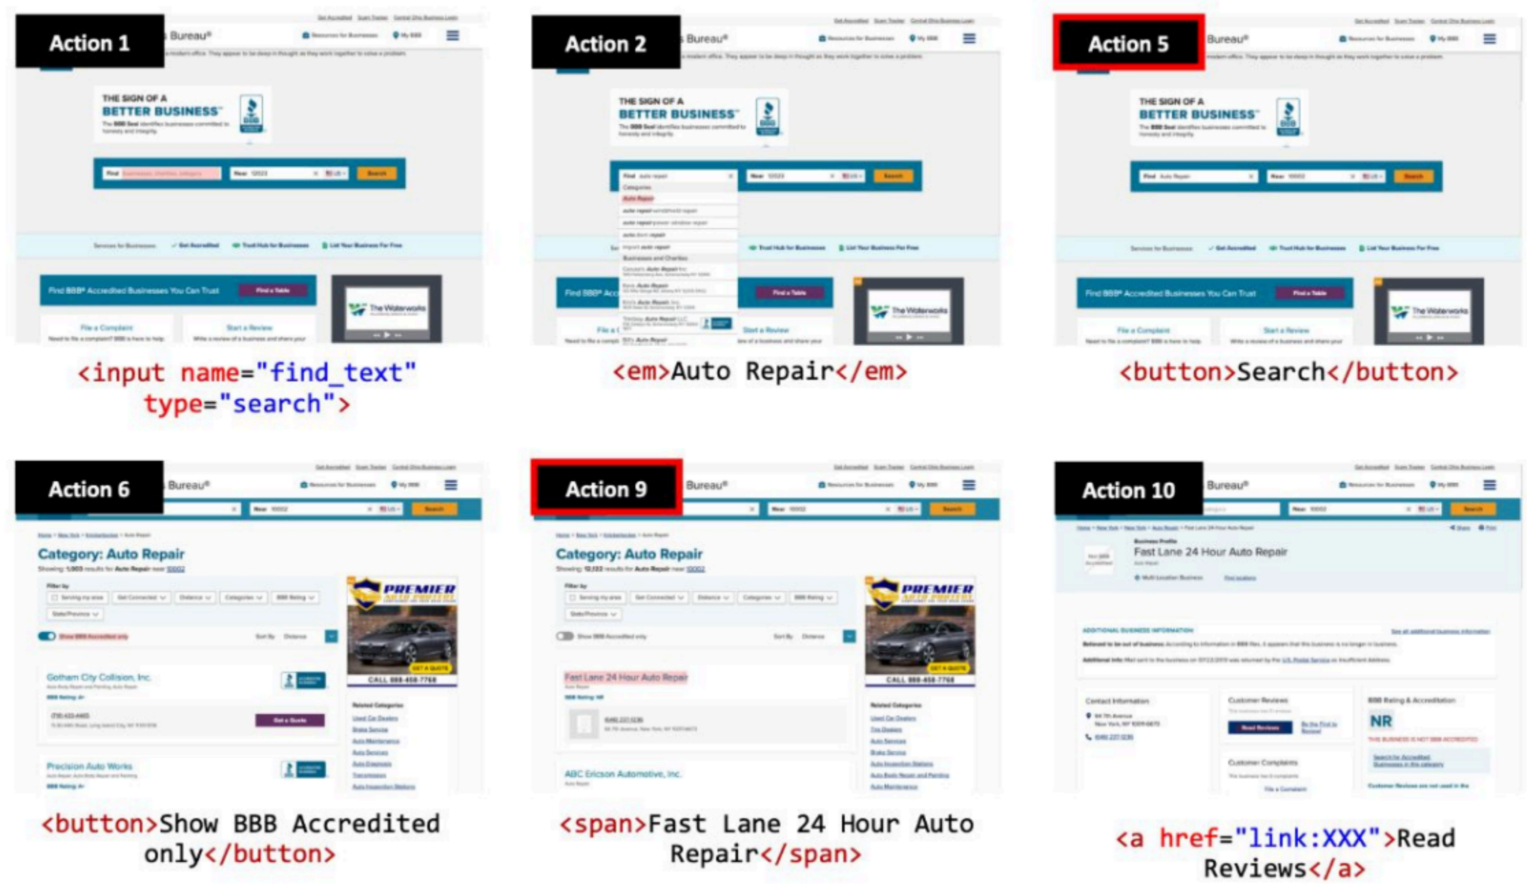

Deng et al., "Mind2Web: Towards a Generalist Agent for the Web." NeurIPS (2024)

Planning paradigms for language agents

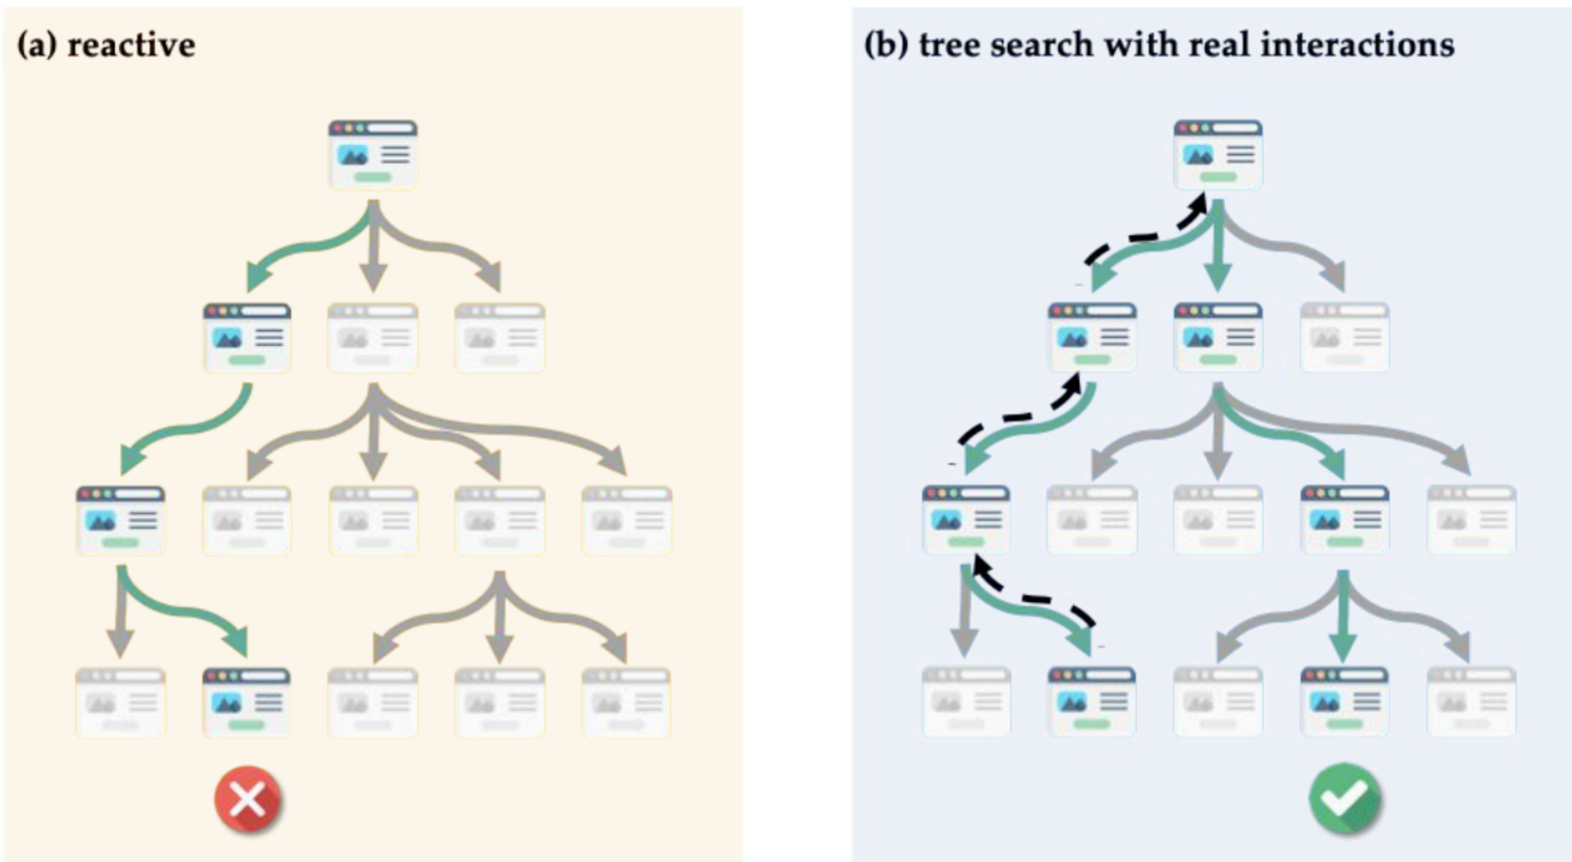

fast, easy to implement

greedy,

Tree search with real interactions

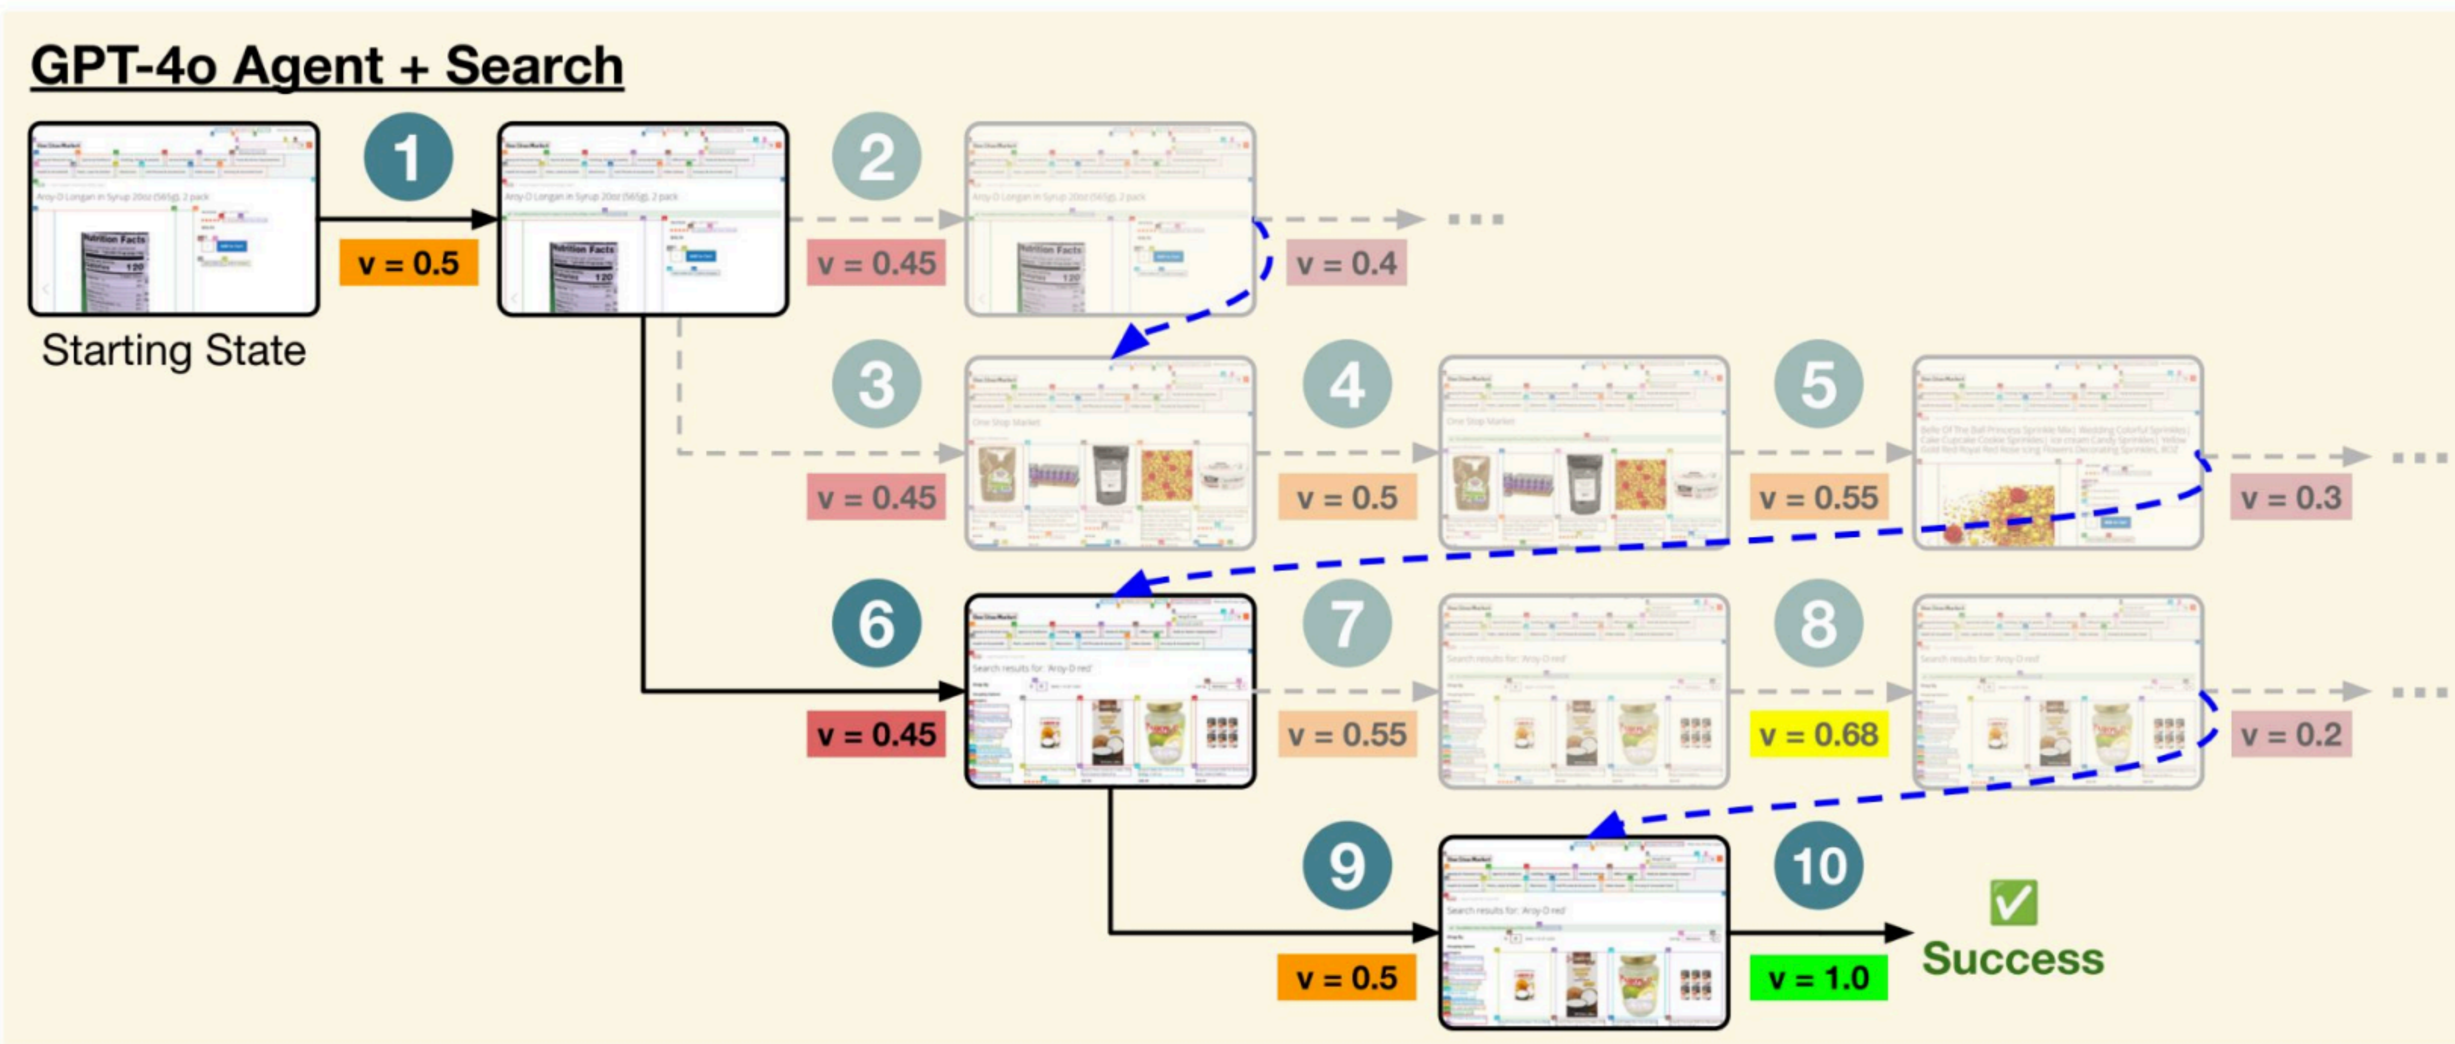

Jing Yu Koh, Stephen McAleer, Daniel Fried, Ruslan Salakhutdinov. "Tree Search for Language Model Agents." arXiv preprint arXiv:2407.01476 (2024).
Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, Yu-Xiong Wang. “Language Agent Tree Search Unifies Reasoning Acting and Planning in Language Models.” ICML (2024).

Challenges with tree search in the real world

- Many actions are state-changing and irreversible → backtracking ❌
- Safety/privacy risks
- Inference-time exploration could be slow and costly

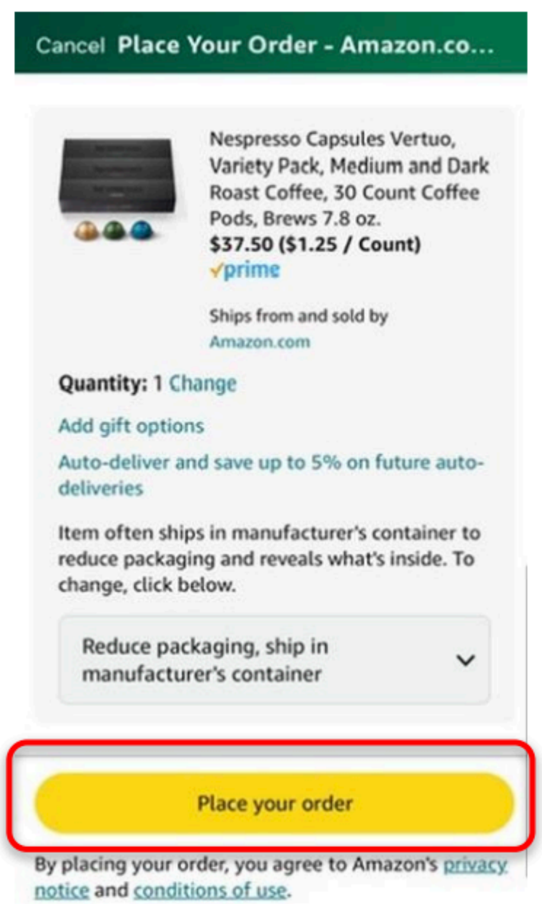

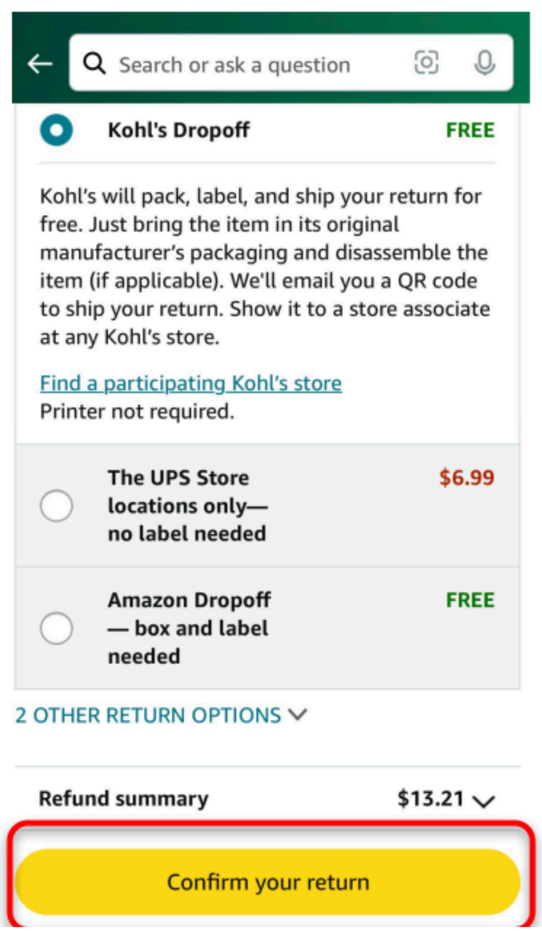

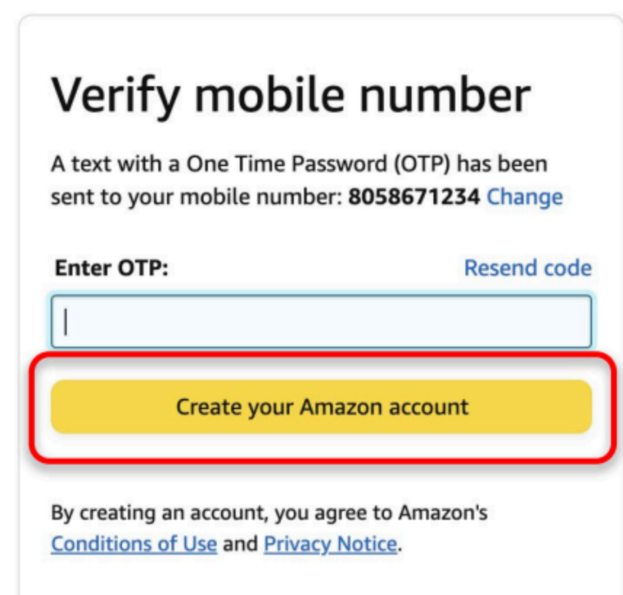

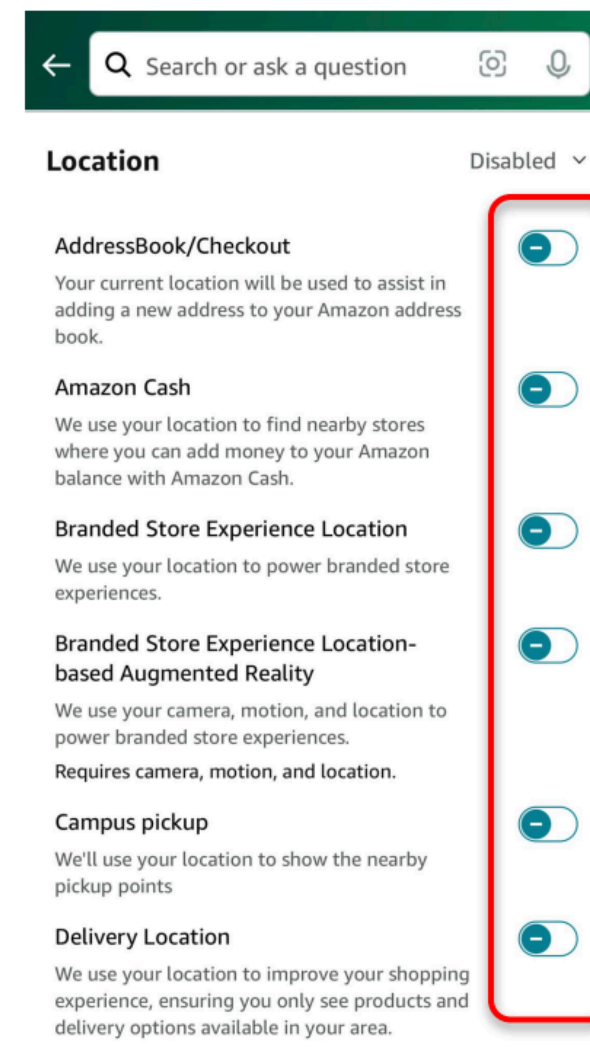

Planning paradigms for language agents

- Train a world model that knows what is going to happen

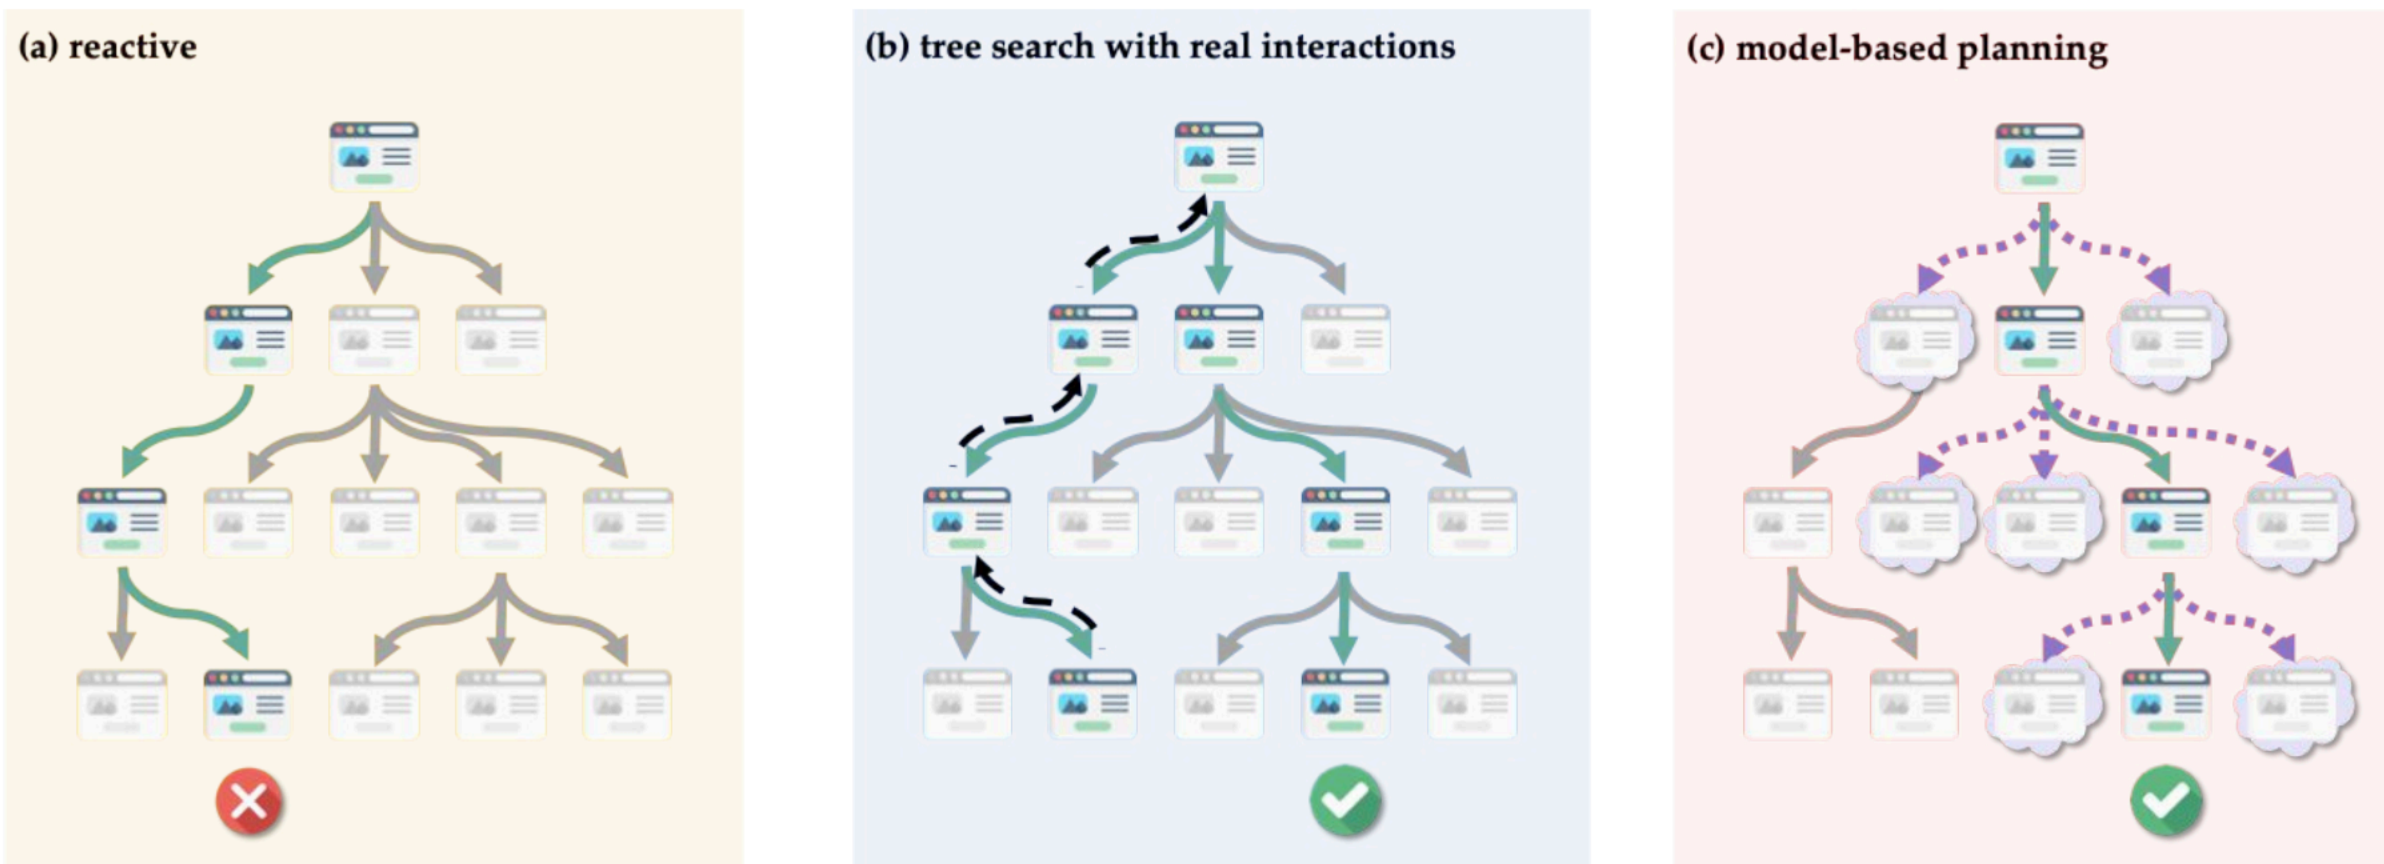

(a) reactive	(b) tree search with real interactions	(c) model-based planning
fast, easy to implement	systematic exploration	faster, safer, systematic exploration
greedy,	irreversible actions, unsafe, slow	how to get a world model?

Still open question?

Part III: LLM Coding Agents for Software Development

More and more major businesses and industries are being run on software and delivered as online services—from movies to agriculture to national defense. […] Over the next 10 years, I expect many more industries to be disrupted by software […].

— Marc Andreessen - Why Software is Eating the World (2011)

If we gave everyone the ability to quickly write software to achieve their goals, what could they do?

What is Involved in Developing Software?

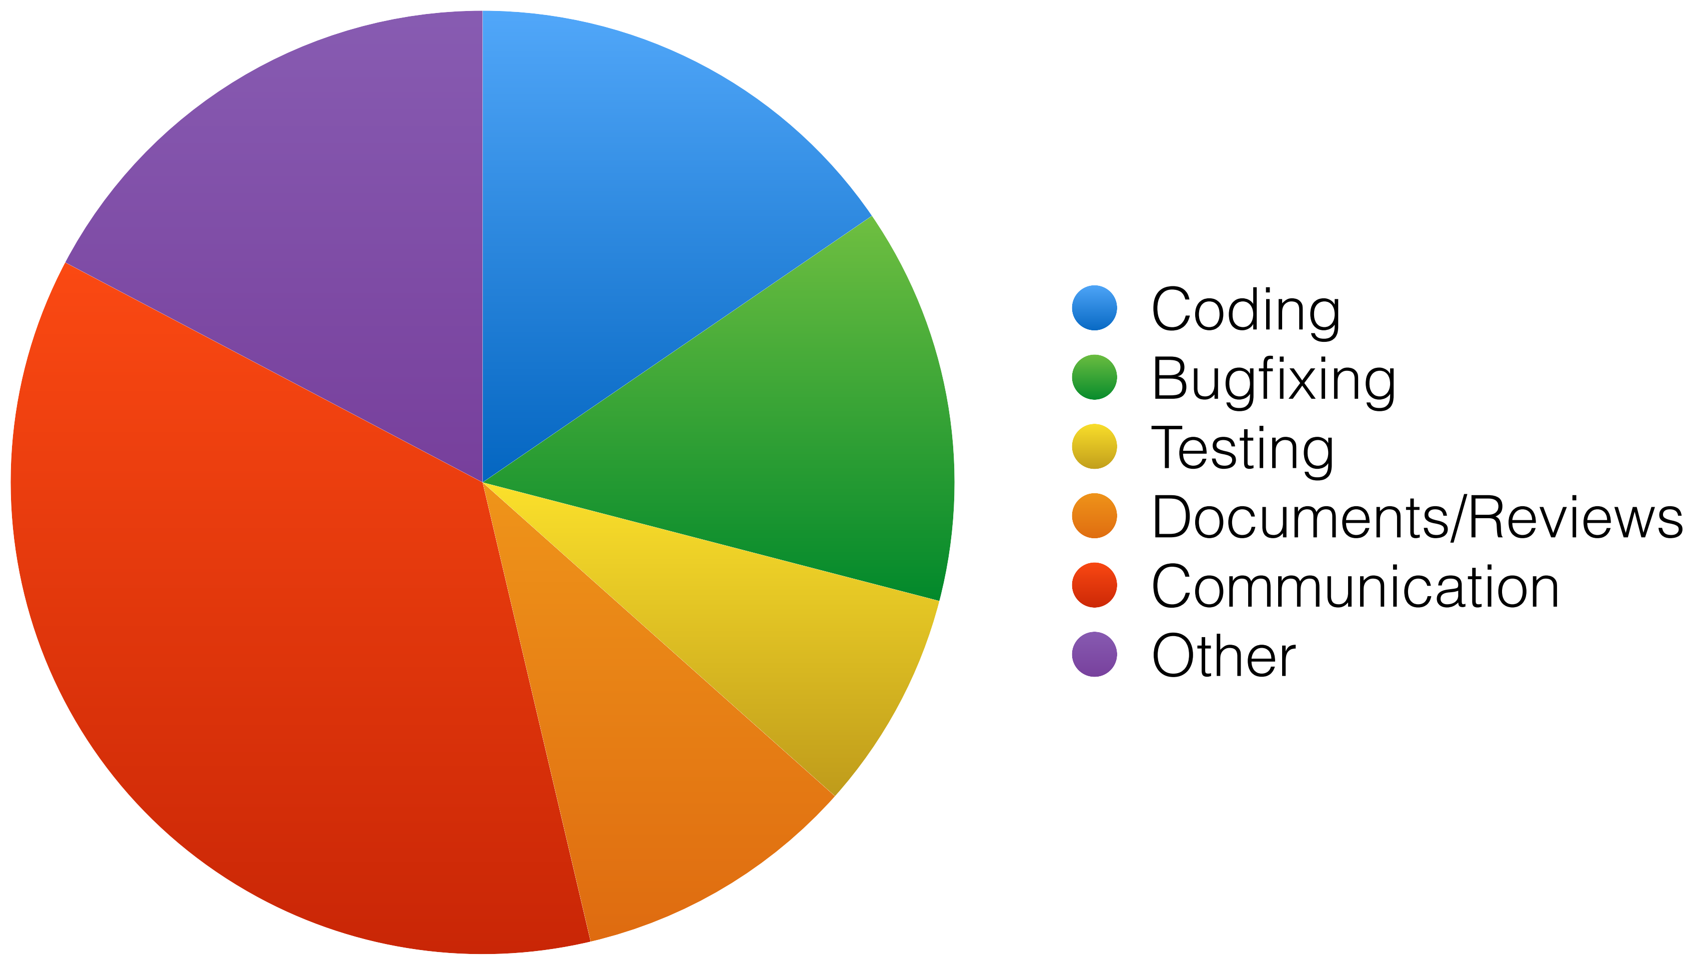

Today was a Good Day: The Daily Life of Software Developers
Meyer et al. 2019

How Can We Support Developers?

(Neubig 2024)

Level	Self Driving	Software Development
0: No Automation	Manual driving	Manual Coding
1: Driver Assistance/ Code Completion	Adaptive cruise control/braking	Copilot/Cursor code completion
2: Partial Automation	Tesla's autopilot	Copilot chat refactoring
3: Conditional Automation	Mercedes-Benz drive pilot	DiffBlue test generation, Transcoder code porting
4: High Automation	Cruise self-driving vehicles	Devin/OpenDevin end-to-end development
5: Full Automation	...	...

Development Copilots

- Work synchronously with the developer to ease writing code
- e.g. **Github Copilot/Cursor**

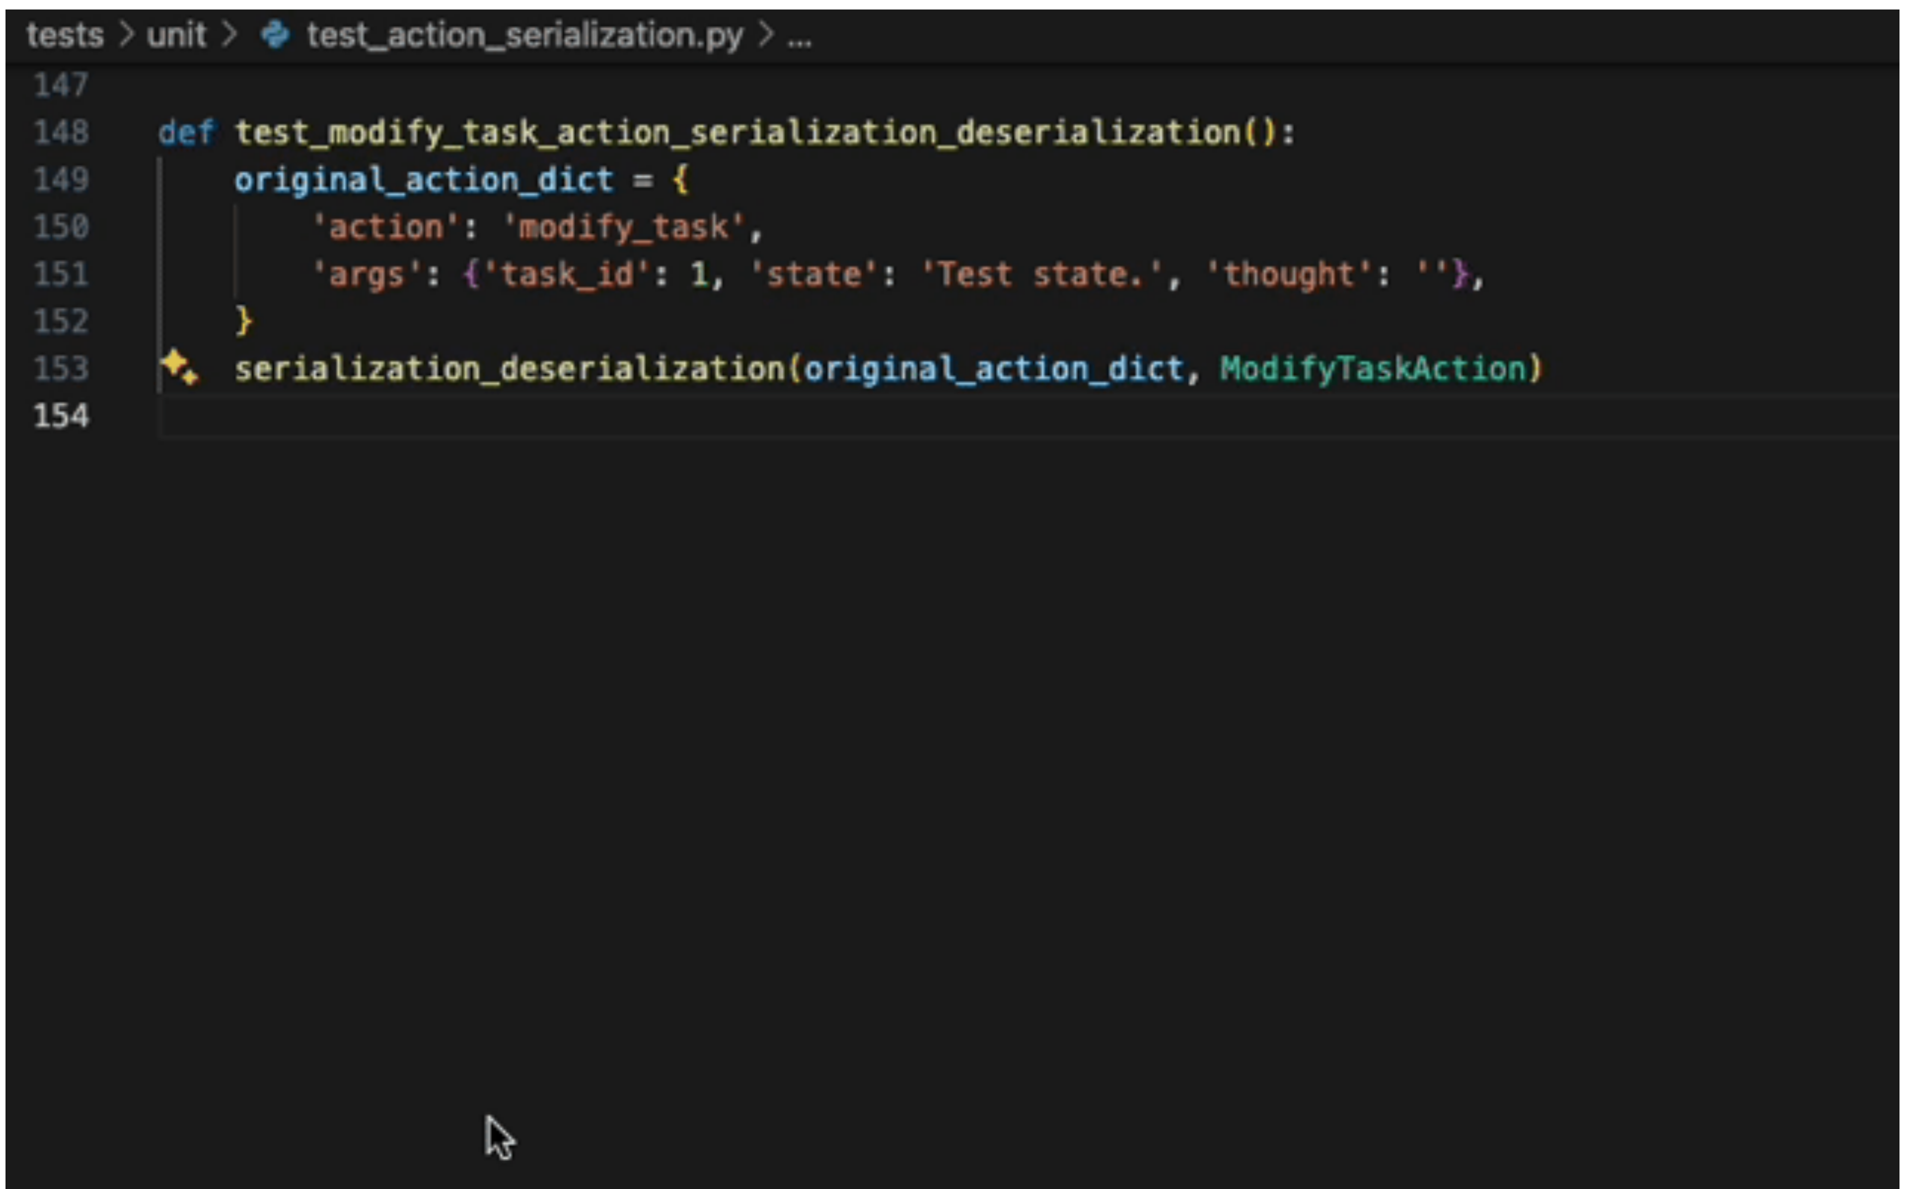

Development Agents

- For coding (e.g. SWE-Agent, Aider)
- For broader development (e.g. Devin, OpenHands)

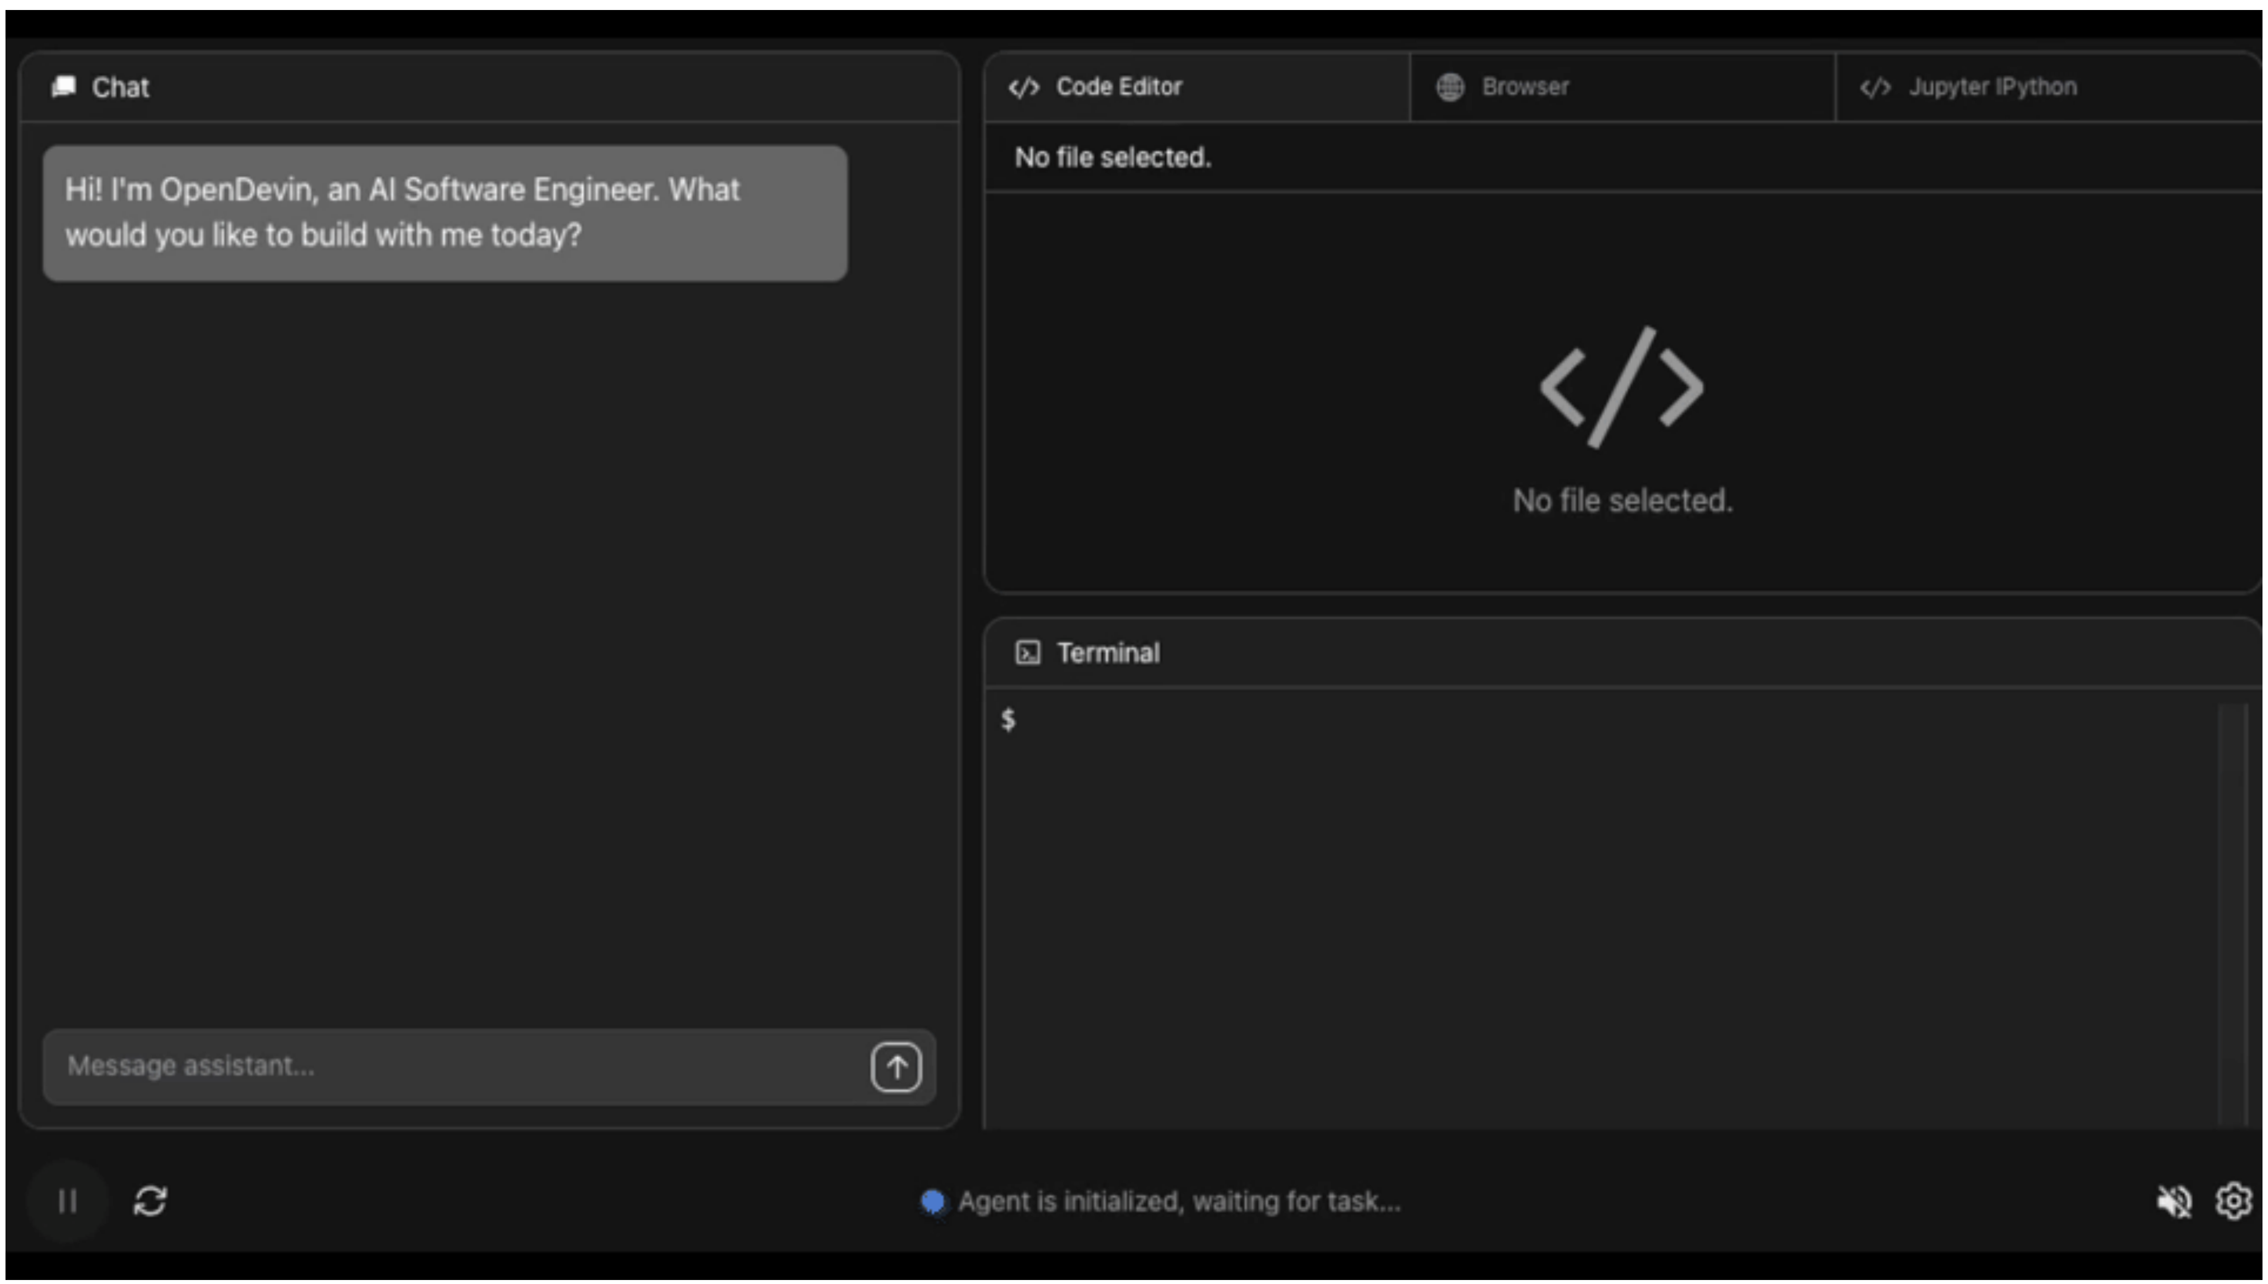

Non-coding Tasks

- Gathering information from Github
- Managing task resolution software
- Setting up web infrastructure

How Promising?

- Code generation leads to large improvements in productivity (Github 2022)

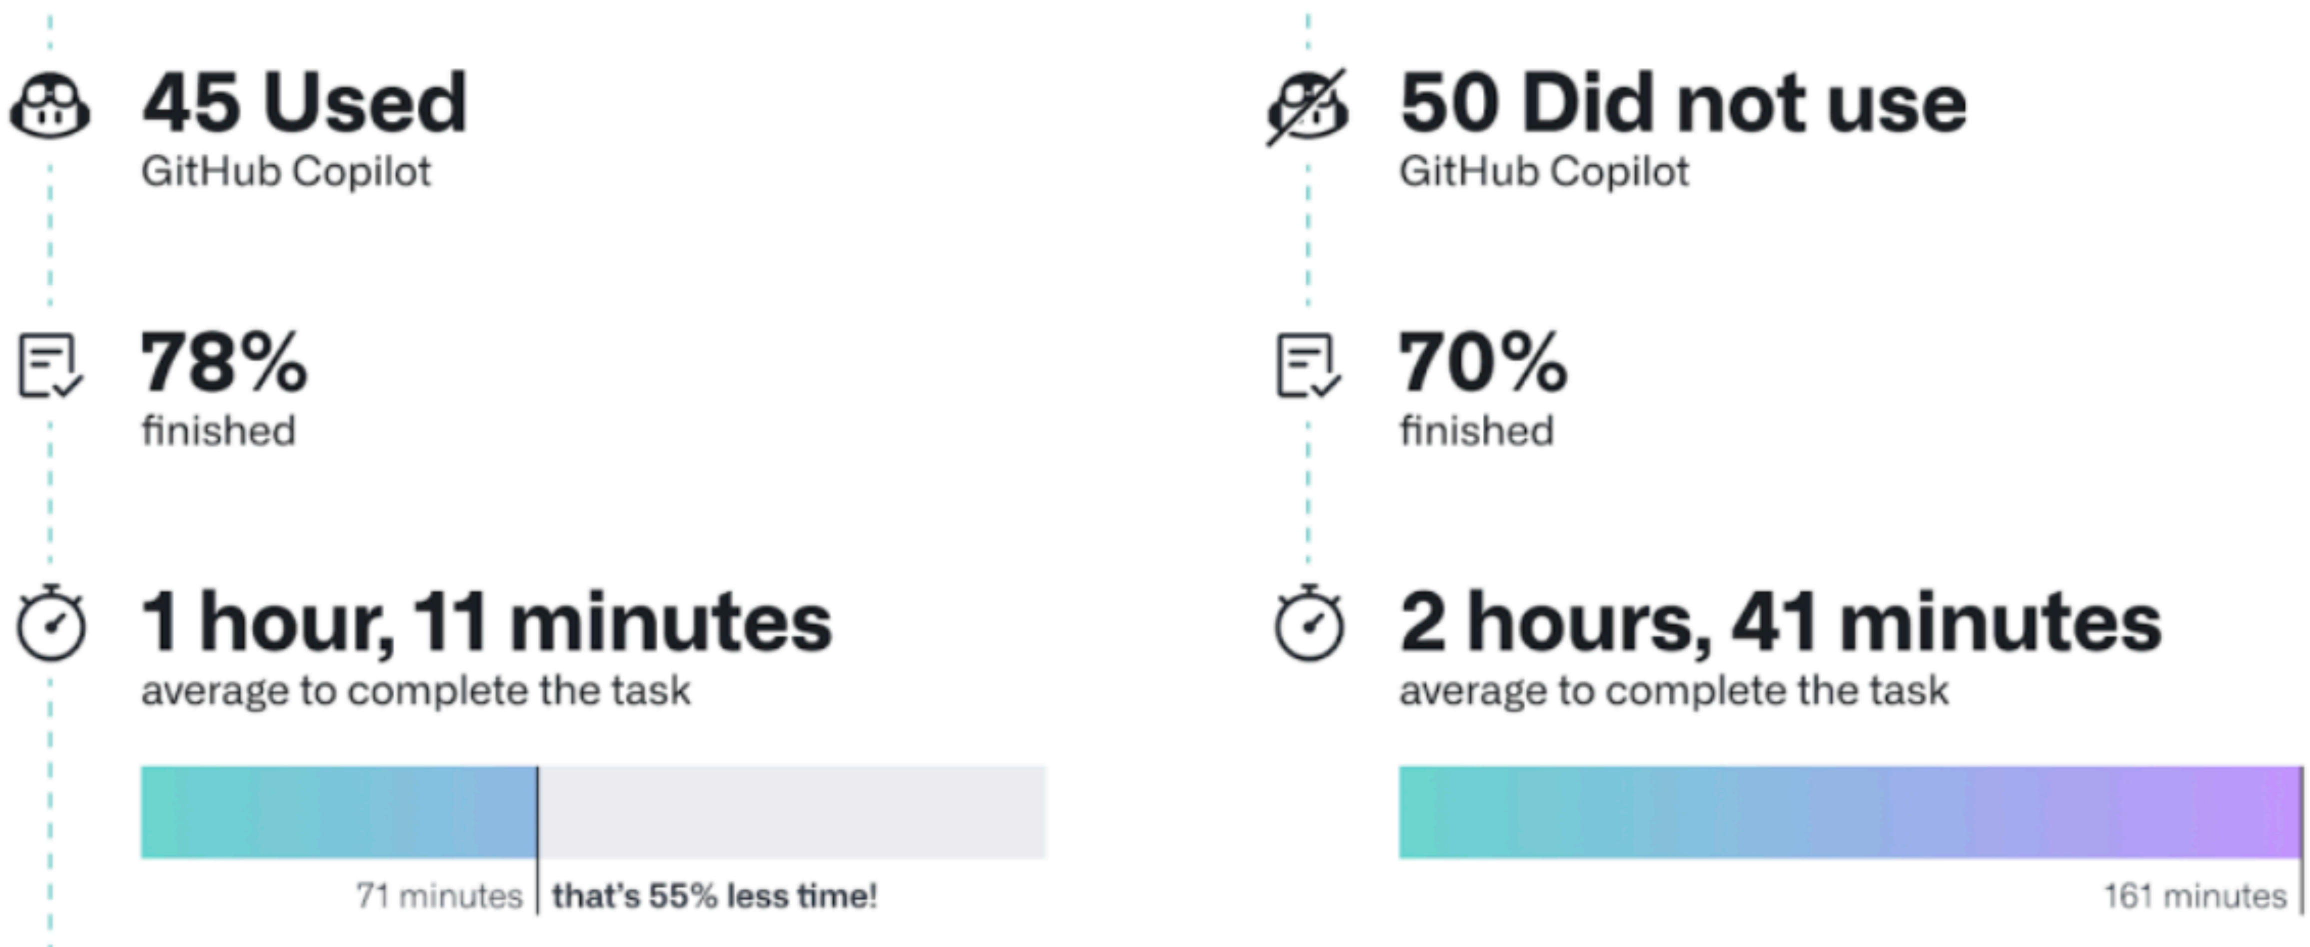

Challenges in Development Agents

- Defining the Environment
- Designing an Observations/Actions
- Code Generation (atomic actions)
- File Localization (exploration)
- Planning and Error Recovery
- Safety

Types of Environments

- **Actual Environments:**
  - *Source Repositories:* Github, Gitlab
  - *Task Management Software:* Jira, Linear
  - *Office Software:* Google Docs, Microsoft Office
  - *Communication Tools:* Gmail, Slack
- **Testing Environments:**
  - Mostly focused on coding!
  - Developers do more, e.g. browse the web (next session)

Data Science Notebooks: ARCADE

(Yin et al. 2022)

- Data science notebooks (e.g. Jupyter) allow for incremental implementation
- Allows evaluation of code in context

```
[1] import pandas as pd
c1  df = pd.read_csv('dataset/Gamepass_Games_v1.csv')

[2] u1 Extract min and max hours as two columns
    def get_avg(x):
      try: return float(x[0]) , float(x[1])
      except: return 0, 0

    df['min'], df['max'] = zip(*df['TIME'].str.replace(
      ' hours','').str.split("-").apply(get_avg))
c2

[3] df['ADDED'] = pd.to_datetime(
c3      df['ADDED'],format="%d %b %y",errors='coerce')

[4] u2 In which year was the most played game added?
    df['GAMERS']=df['GAMERS'].str.replace(
                               ',','').astype(int)
c4  added_year=df[df['GAMERS'].idxmax()]['ADDED'].year

[5] u3 For each month in that year, how many games that
       has a rating of more than four?

    df[(df['ADDED'].dt.year== added_date.year) &
    (df['RATING']>4)].groupby(
               df["ADDED"].dt.month)['GAME'].count()
c5

[6] u4 What is the average maximum completion time for
       all fallout games added in 2021?
    fallout=df[df['GAME'].str.contains('Fallout')]
    fallout.groupby(fallout['ADDED'].dt.year).get_group(
c6    2021)['max'].mean()

[7] u5 What is the amount of games added in each year
       for each month? (show a table with index as years,
       columns as months and fill null values with 0)

    pd.pivot_table(df, index=df['ADDED'].dt.year, ...,
      aggfunc=np.count_nonzero,
      fill_value='0').rename_axis(
        index='Year', columns='Month')
c7
```

Figure 1: An example of a computational notebook adapted from our dataset, with examples of reading and preprocessing data (cell c_1), data wrangling (cell c_2, c_3), and data analysis (cells $c_3 - c_7$). Annotated NL intents are shown in green.

Dataset: SWEBench

(Jiminez et al. 2023)

- Issues from GitHub + codebases -> pull request

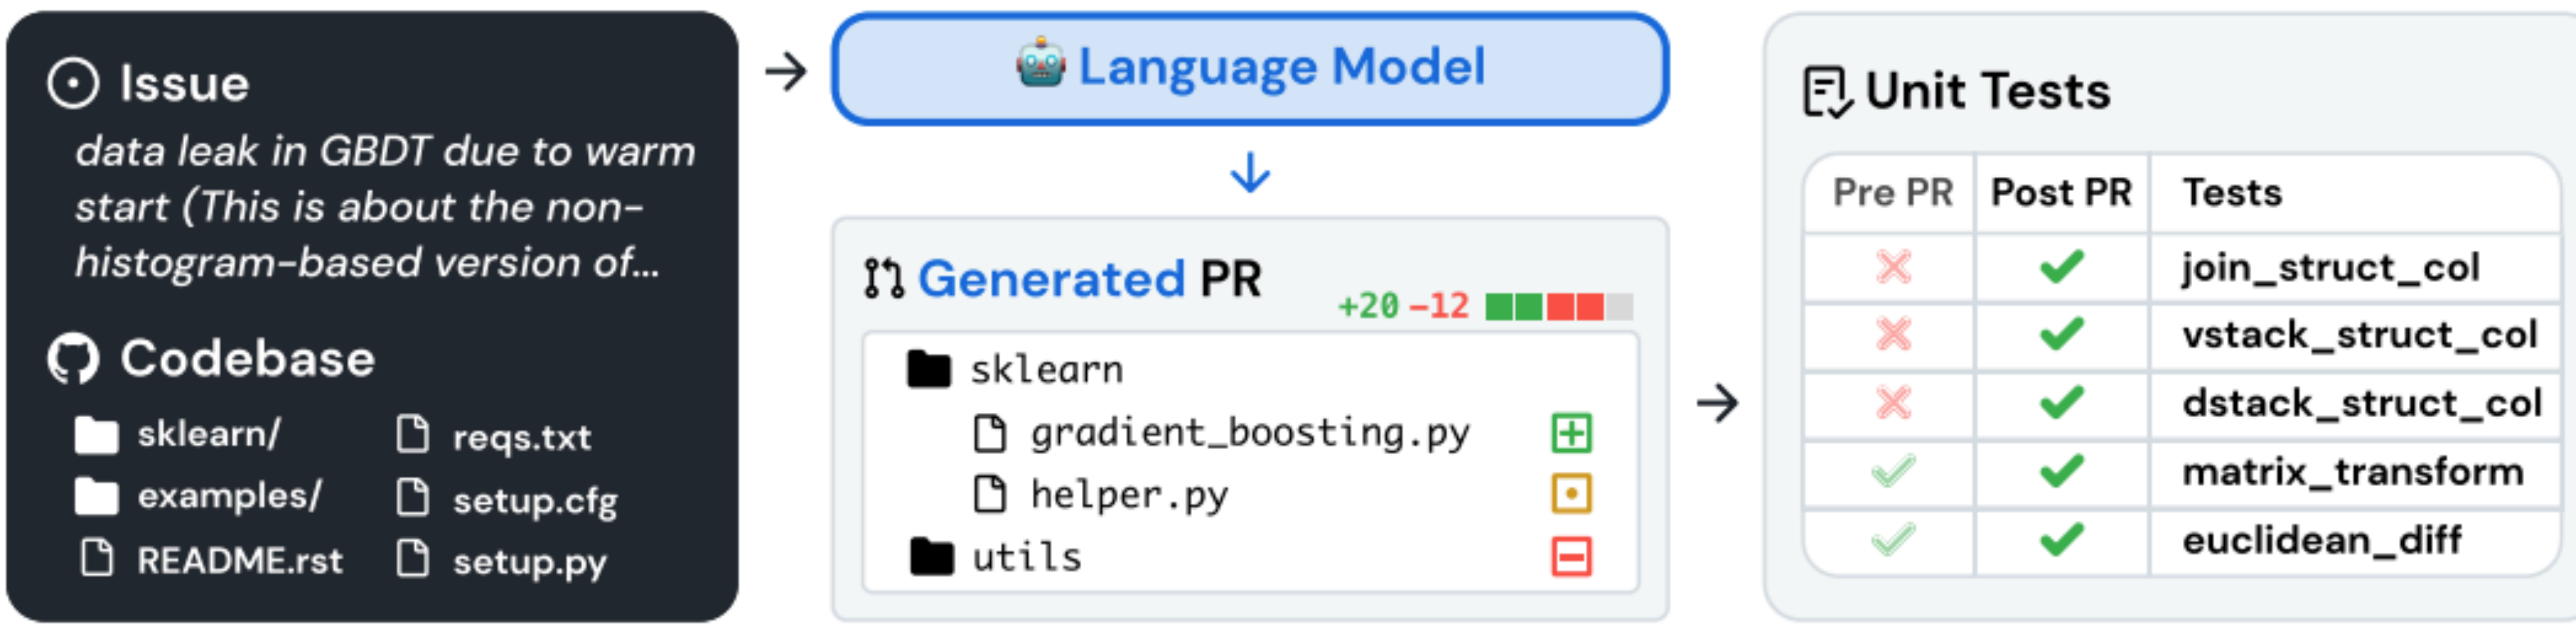

- Requires long-context understanding, precise implementation

Action Spaces

- Coding Agents
  - Understand repository structure
  - Read in existing code
  - Modify or produce code
  - Run code and debug

Example: CodeAct (Wang et al. 2024)

- Interact w/ the environment through code

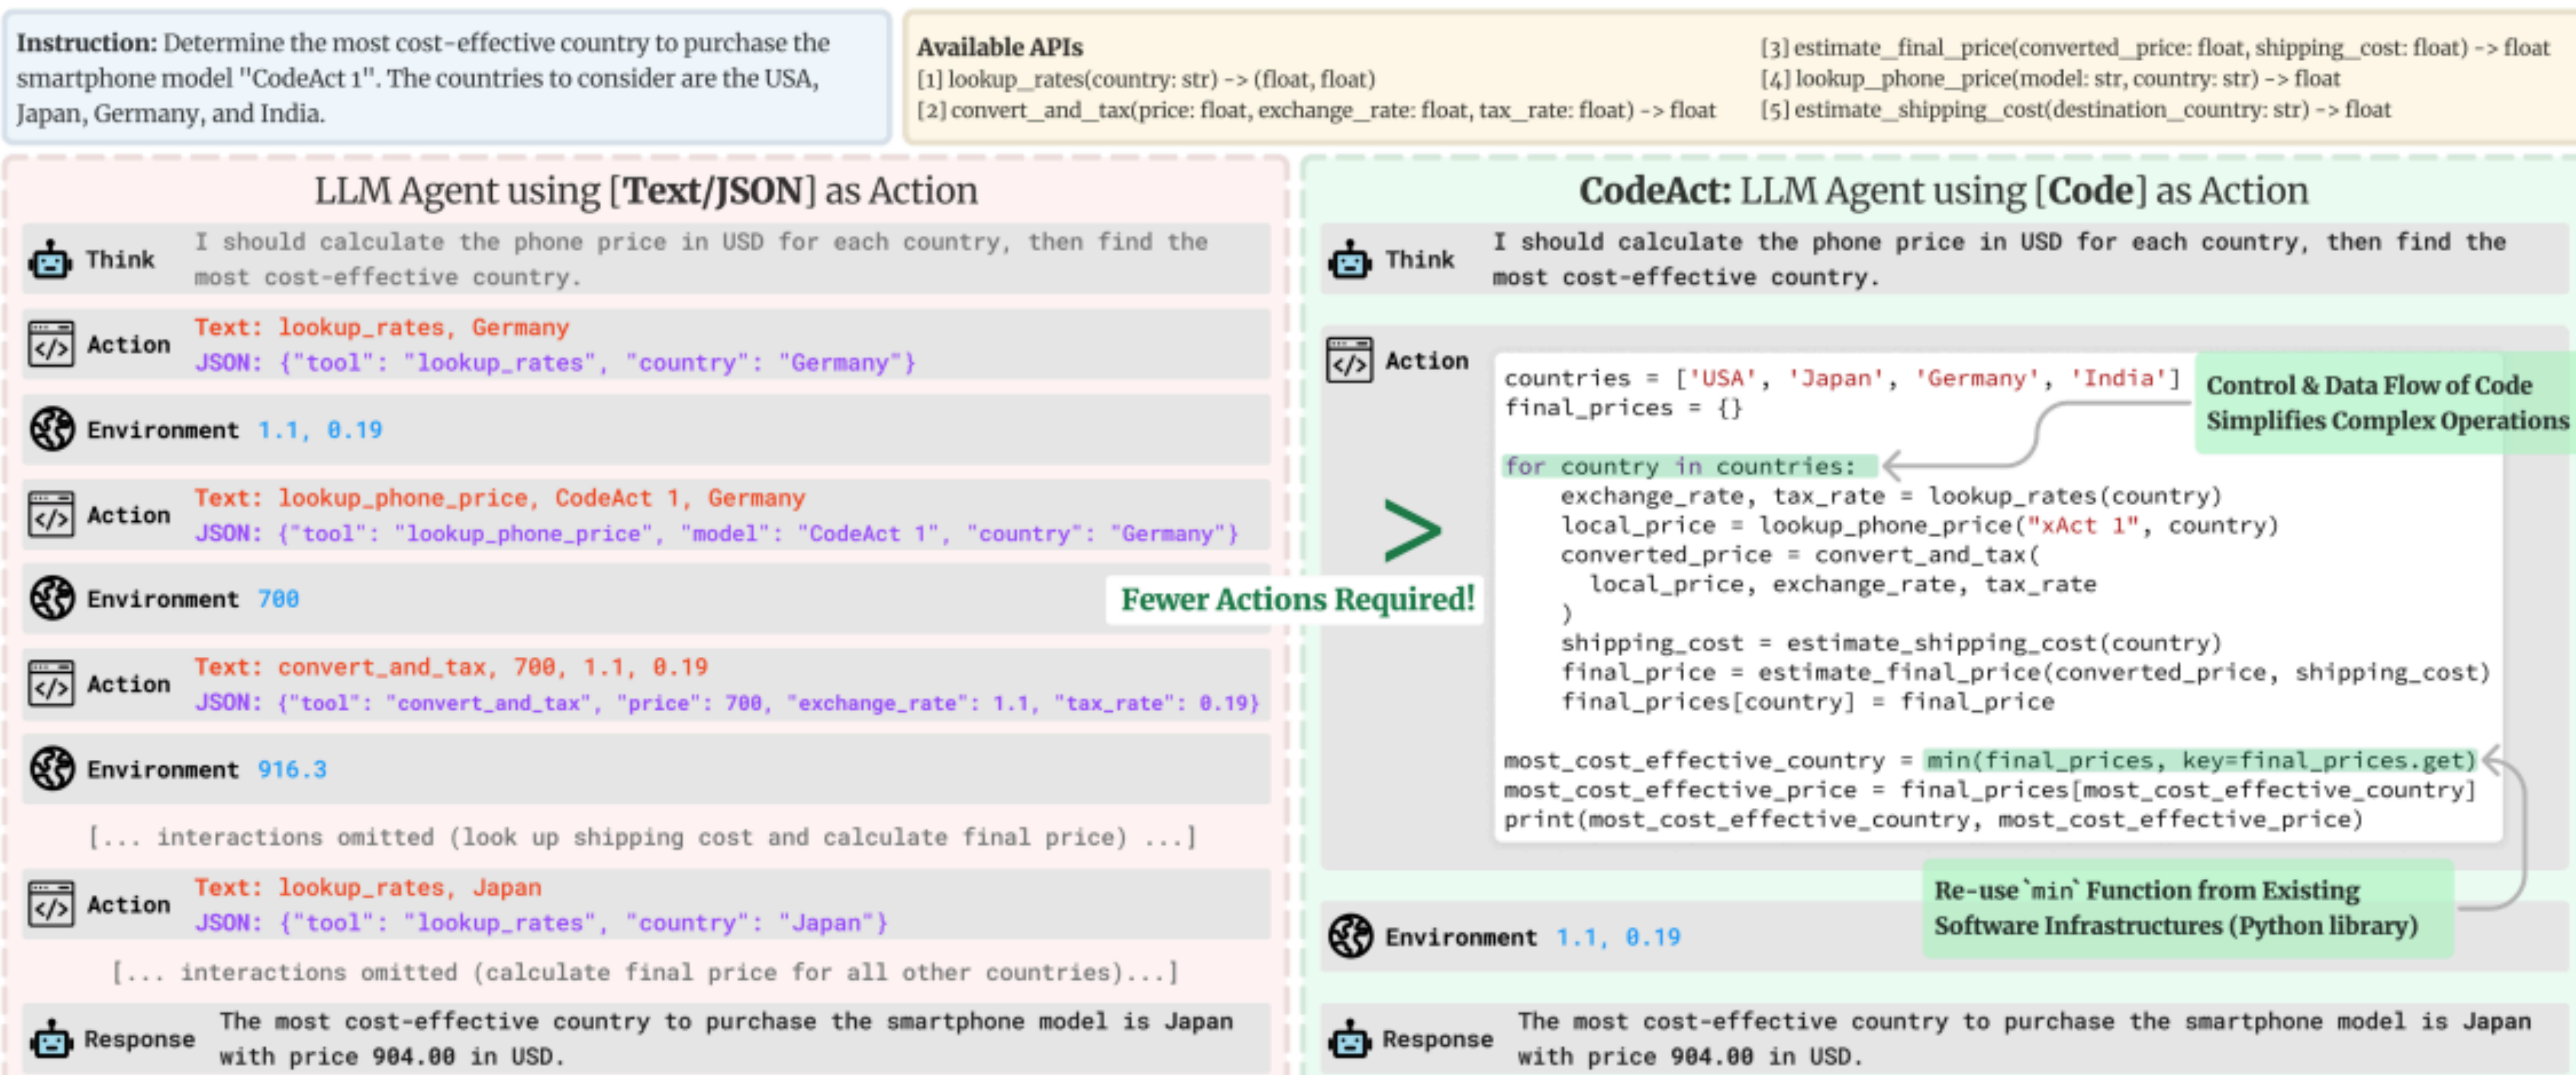

- Can execute bash commands, Jupyter commands
- Faster resolution, higher success than direct tool use

Example: SWE-Agent

(Yang+Jimenez et al. 2024)

- Define specialized tools that make it possible to efficiently explore repositories and edit code

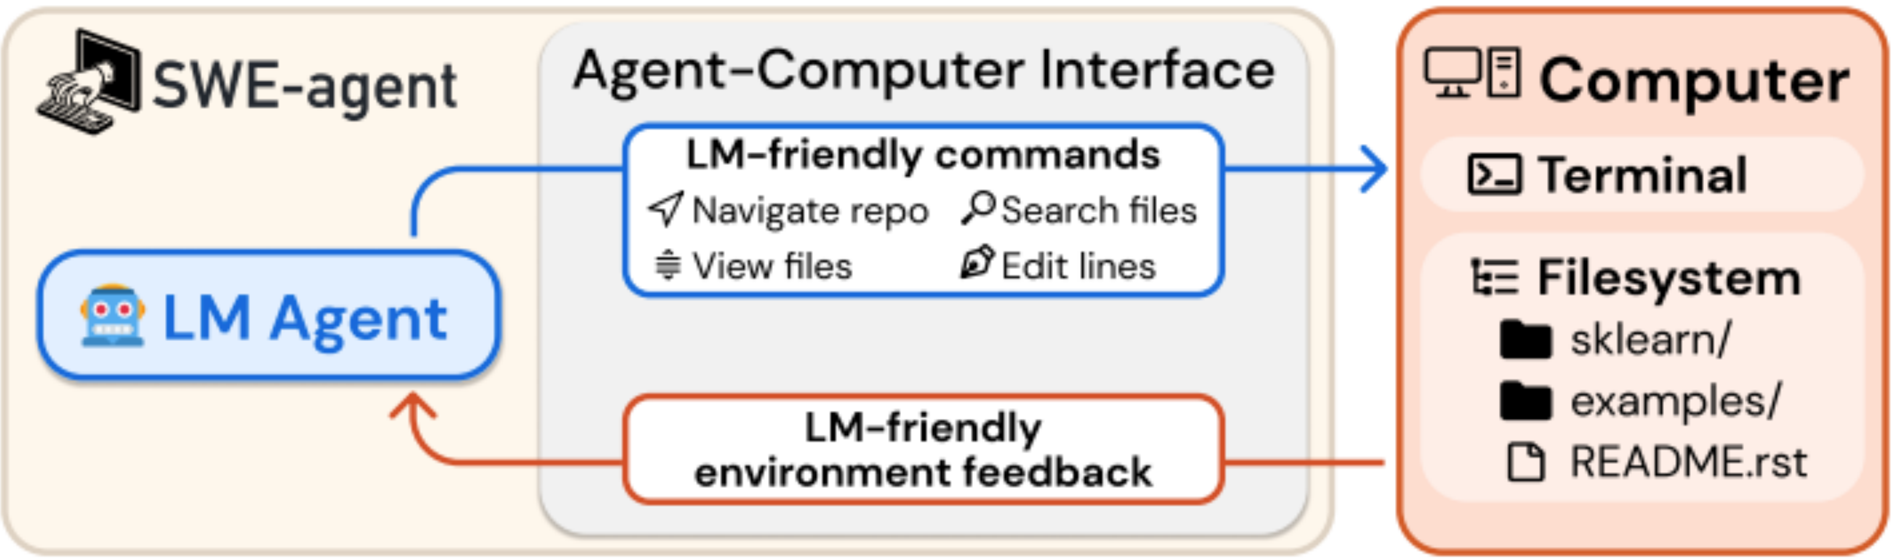

Observation (File Viewer)

```
[File: solvers/diophantine.py (3093 lines total)]
(400 lines above)
401:     elif total_degree == 3:
402:         diop_type = "cubic_thue"
403:
404:     elif (total_degree > 3):
405:         if all(coeff[k] == 1):
406:             diop_type = 'general_sum_even'
407:
408:     if diop_type is not None:
409:         return var, coeff, diop_type
(2684 lines below)
```

Response (Thought + Edit Action)

We need to add a check to ensure that all powers are the same. Let's edit the code to include this check.

```
edit 404:407
    elif (total_degree > 3 and
          len(set(k.exp[k.is_Pow])) == 1):
        if all(coeff[k] == 1):
            diop_type = 'general_sum_even'
end_of_edit
```

Example: OpenHands

(Wang et al. 2024)

- Defines “event stream” for coding, execution, and browsing actions/observations
- Implements SWE-agents style actions as “agent skills” that can be called

File Localization

LLM-based Localization

- Finding the correct files given user intent

> **What problem or use case are you trying to solve?**
> When in confirmation mode it's not possible to give instructions in between steps. You have to reject an action and it seems like it doesn't know that the action was rejected.
>
> **Describe the UX of the solution you'd like**
> The simplest would be to have a third option, confirm action and wait. This way the action is confirmed but before it tries to take the next step you are able to give some feedback. Also if it somehow knows the action was rejected that would be helpful as well so when you do reject an action it knows that action wasn't taken.

https://github.com/All-Hands-AI/OpenHands/issues/4259

- Which JavaScript file should I modify?
- Analogous to environment understanding / exploration problems in other agents

Solution 1:
Offload to the User

- Experienced users familiar with prompting and the project can specify which files to use

> In **`.github/workflows/openhands-resolver.yml`** and **`.github/workflows/openhands-resolver-experimental.yml`**, we should check to make sure that all required environment variables are set before running any additional workflows. If all of the variables are not set, we can fail immediately with an error.

https://github.com/All-Hands-AI/openhands-resolver/issues/146

Solution 2:
Prompt the Agent w/ Search Tools

- e.g. SWE-agent provides a tool for searching repositories

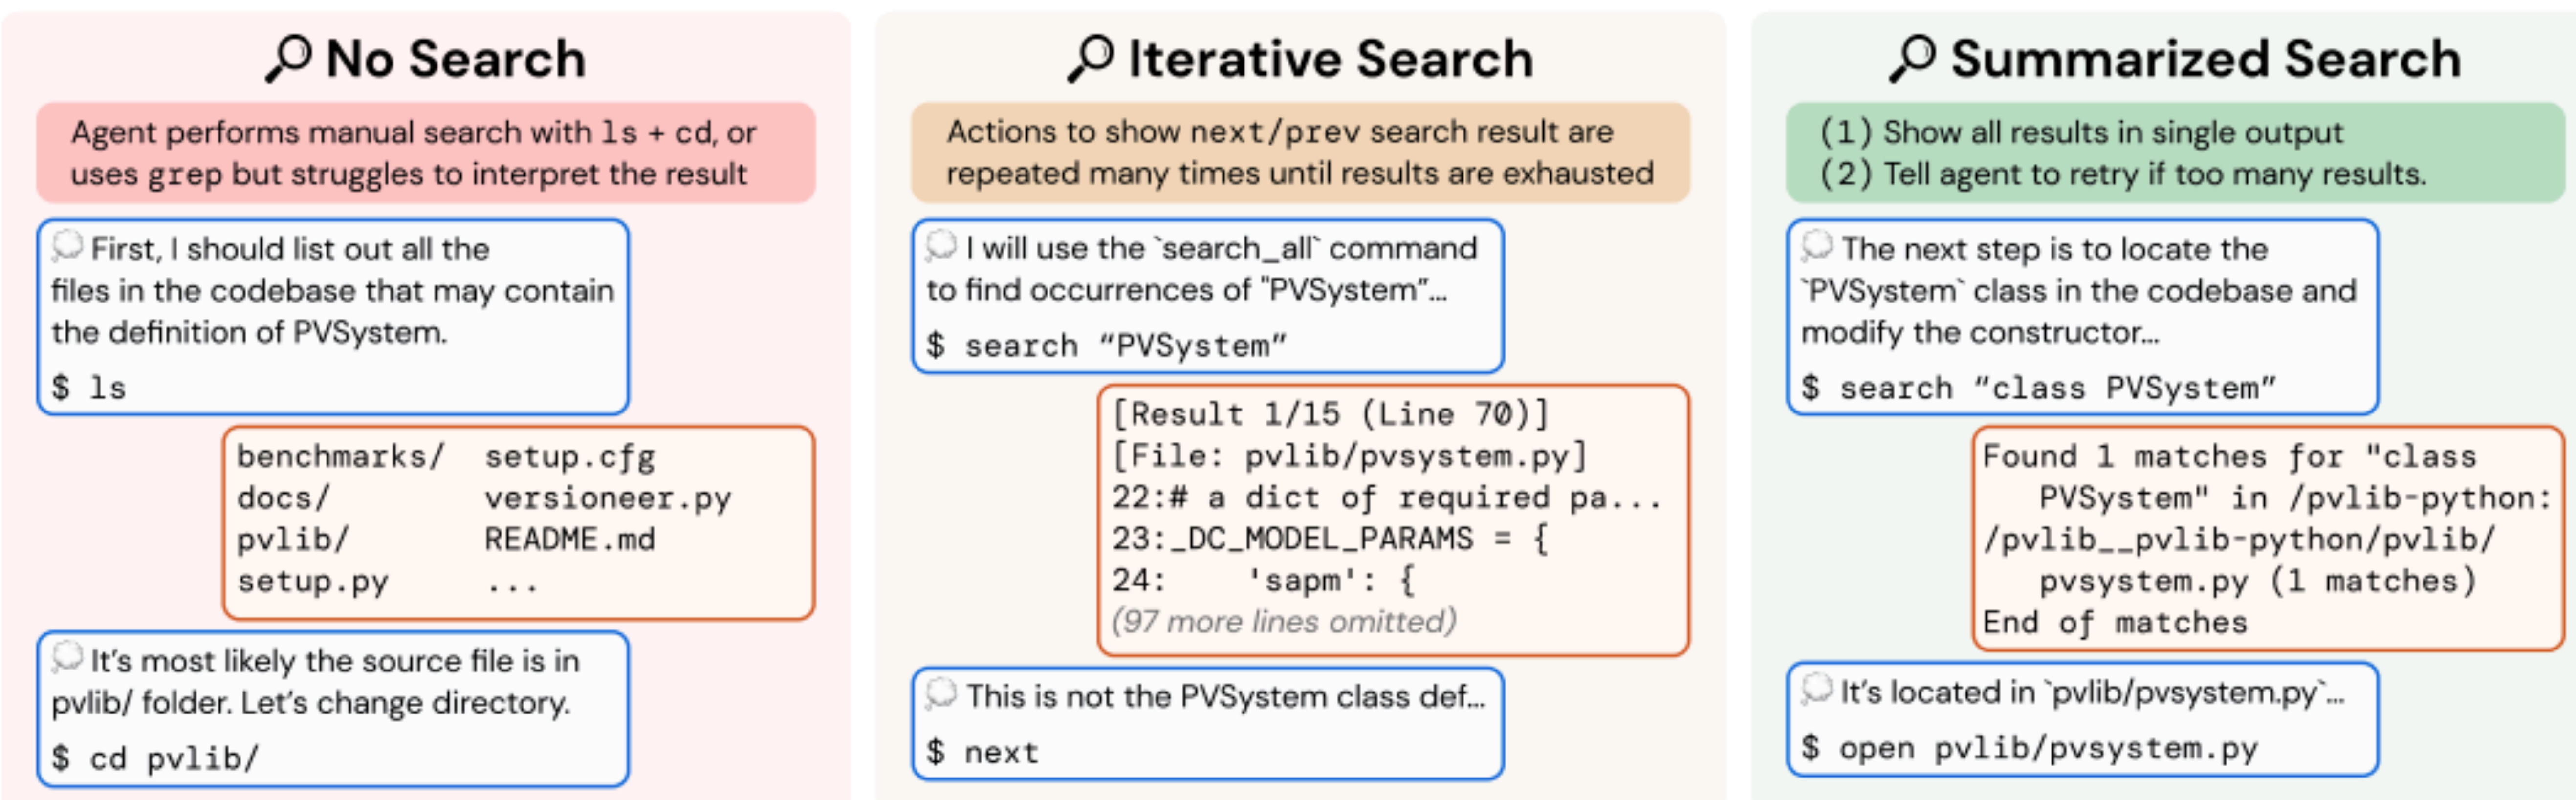

Solution 3: A-priori Map the Repo

- Create a map of the repo and prompt agent with it
- Aider repomap creates a tree-structured map of the repo
- Agentless (Xia et al. 2024) does a hierarchical search for every issue

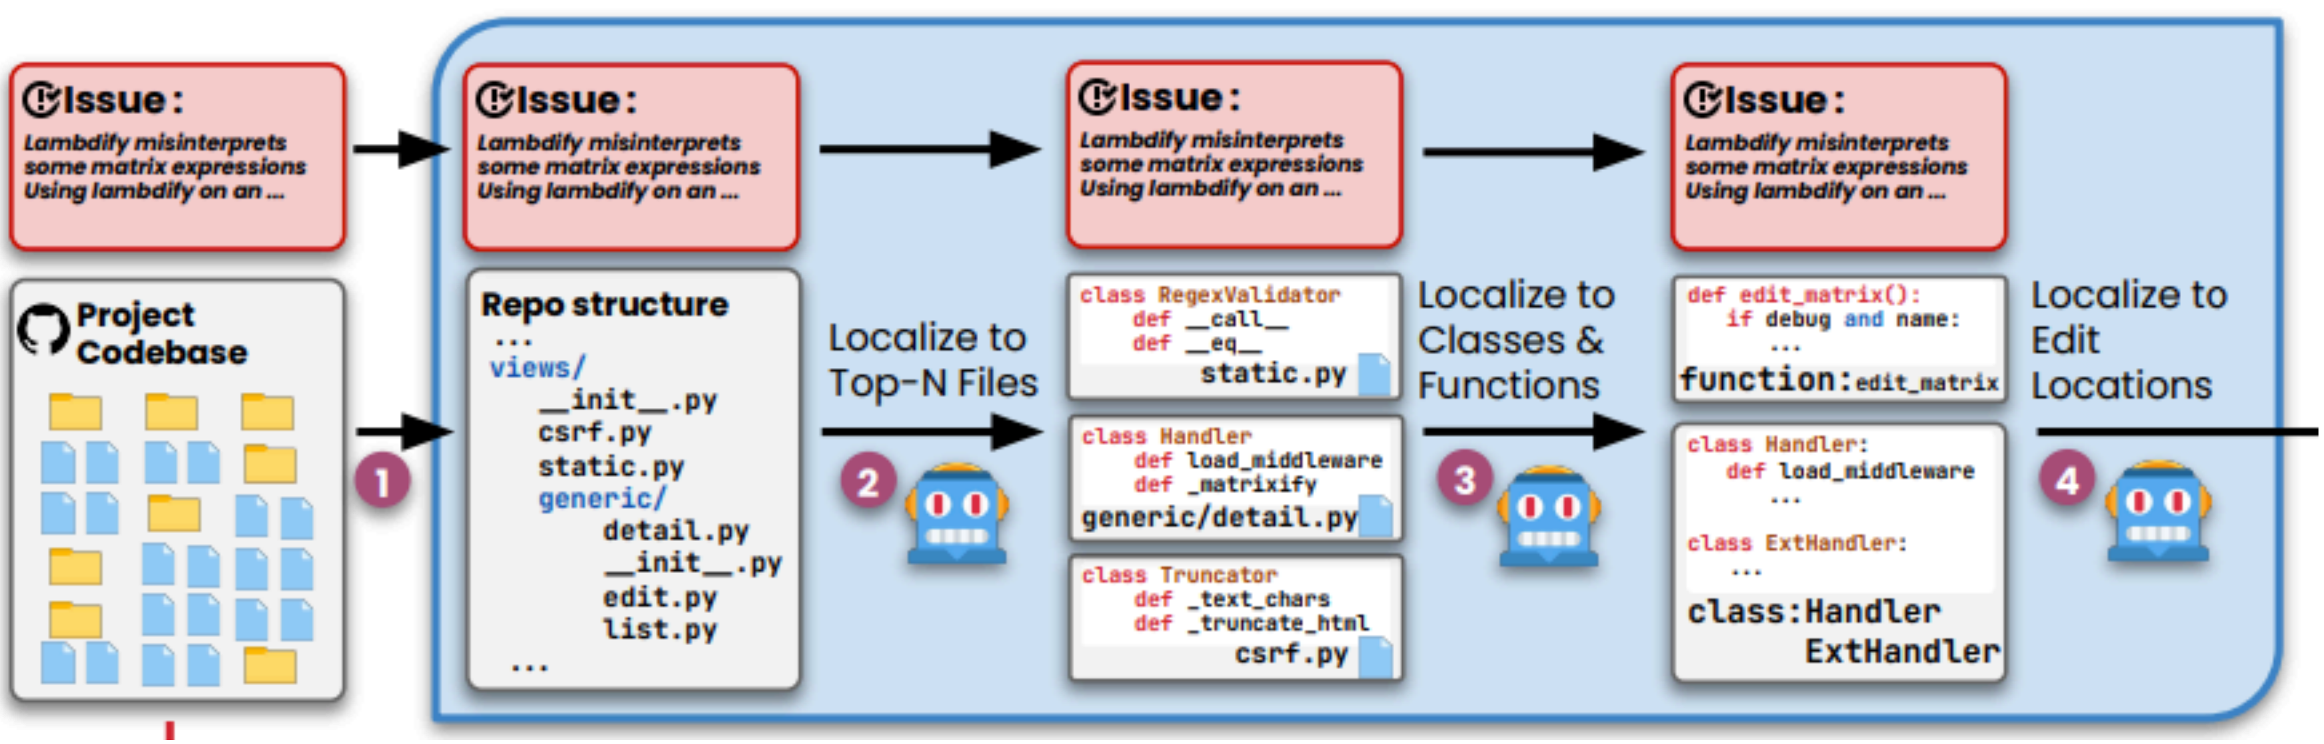

Solution 4: Retrieval-augmented Code Generation

- Retrieve similar code, and fill it in with a retrieval-augmented LM (e.g. CodeRAGBench, Wang+Asai et al. 2024)
- Particularly, in code there is also documentation, which can be retrieved (Zhou et al. 2022)

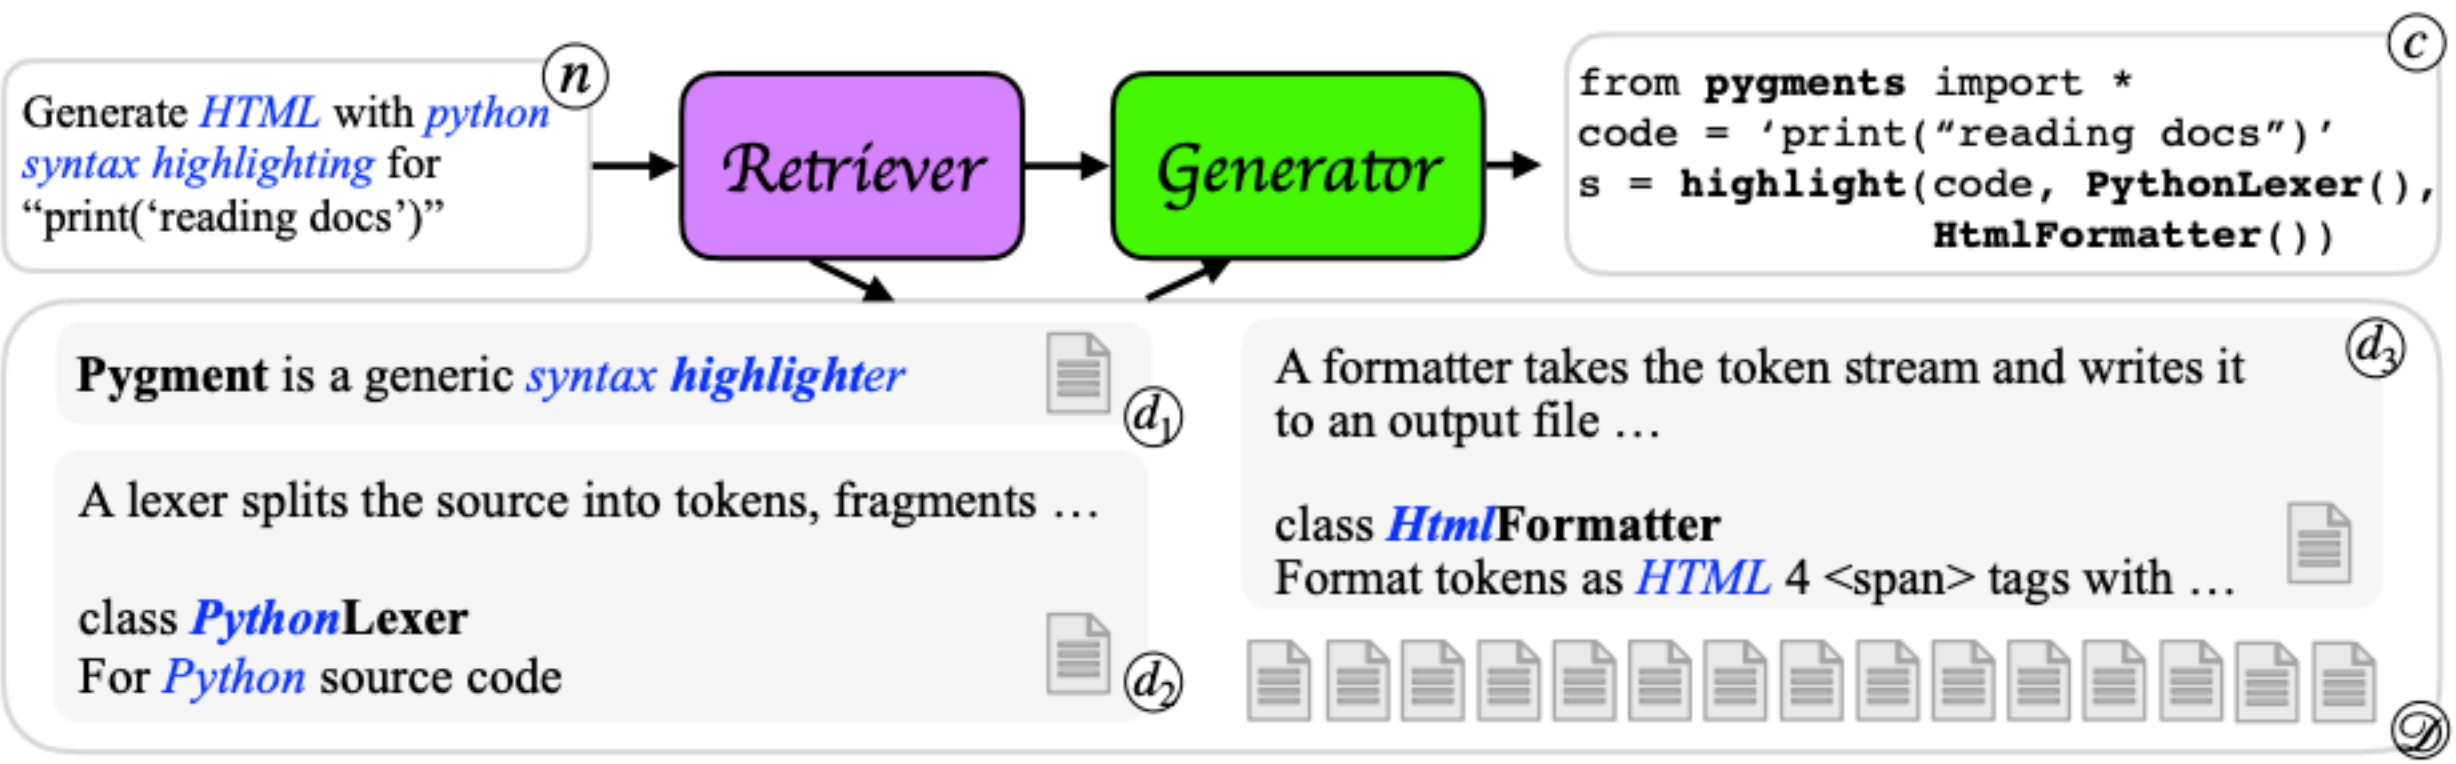

- Unsolved issue: when to perform RAG in agent

Planning and Error Recovery

Hard-coded Task Completion Process

- e.g. Agentless (Xie et al. 2024) has a hard-coded progress of
  - File Localization
  - Function Localization
  - Patch Generation
  - Patch Application

LLM-Generated Plans

- LLM-generated planning step, then one or more executors
- CodeR [Chen et al. 2024]

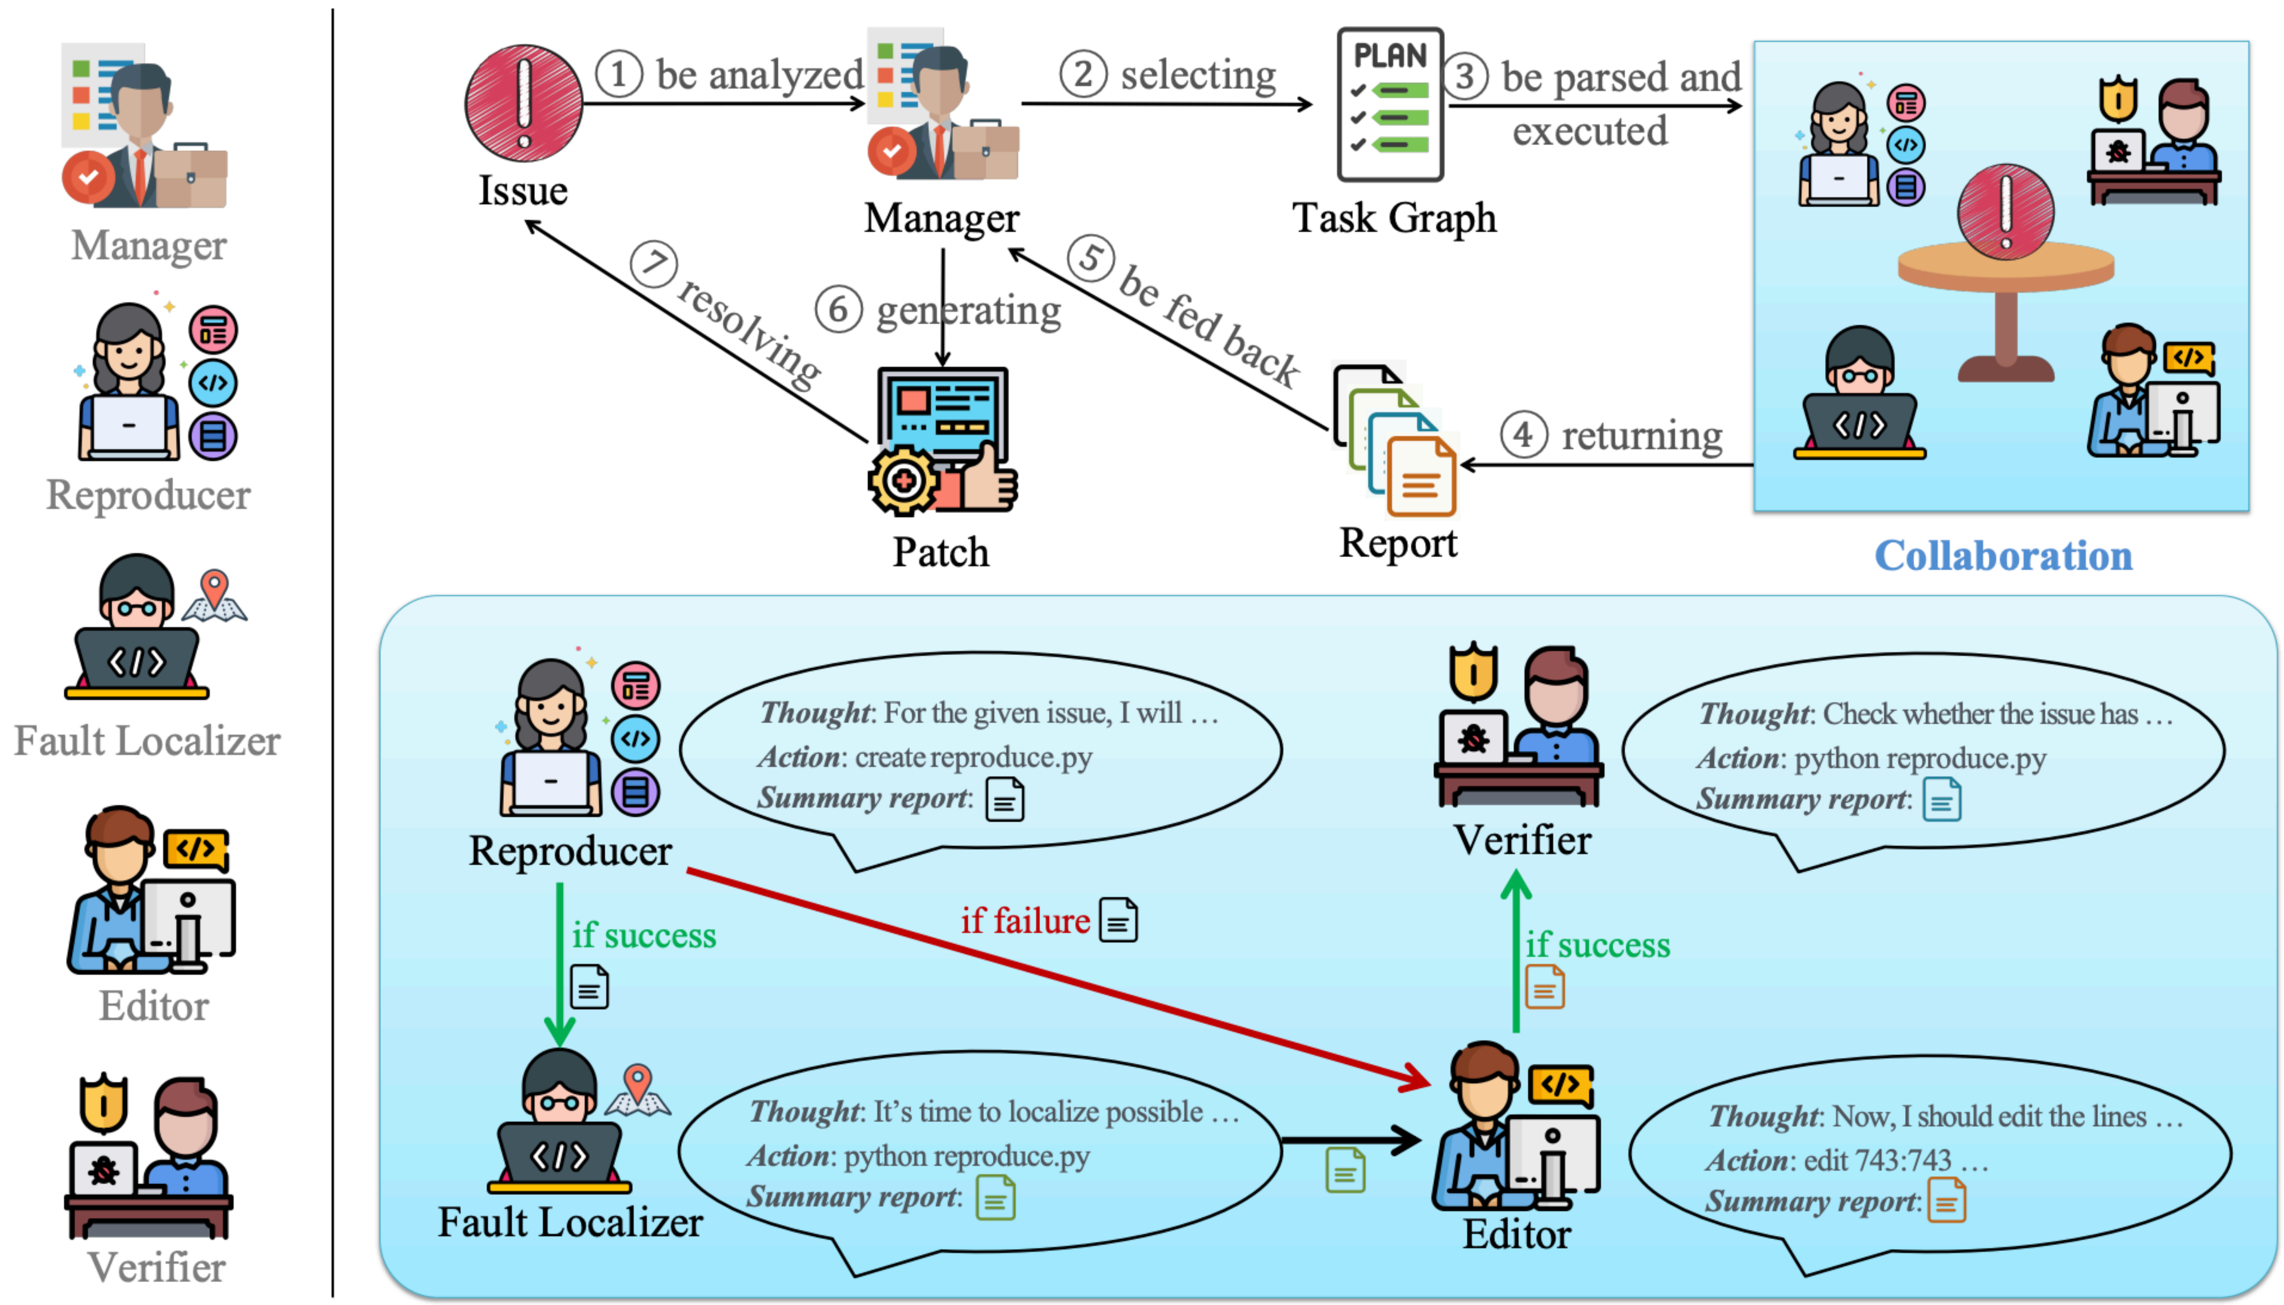

Coding Agents: Takeaway

- Copilots already very useful, code agents getting there
- Current challenges: code LLMs, editing, localization, planning, safety
- Future directions:
    - Agentic training methods
    - Human-in-the-loop
    - Broader software tasks than coding

Questions?