

人工智能导论

智能体 (Agent)

郭兰哲

南京大学 智能科学与技术学院

<https://www.lamda.nju.edu.cn/guolz>

Email: guolz@nju.edu.cn

Recall

- 语言模型：基于概率预测下一个最有可能的词(Next-Token Prediction)
- 现代LLM架构：Transformer架构，核心机制Self-Attention
- ChatGPT：预训练 -> 监督微调 -> RLHF
- LLM推理：Chain-of-Thought
- DeepSeek-R1：监督微调 + 强化学习微调（GRPO）

一些观点

SFT vs RLFT



- 对于形式化推理任务，例如math、code，更容易设计rule-based reward
- **SFT vs RL**
 - 直接利用SFT更多是去拟合数据中的pattern，难以学习到背后的规律；
 - RL则是通过鼓励模型在最大化reward的过程中学习到推理背后的规律，获得的泛化性和推理表现上界更高
 - SFT规范模型输出格式，使得后续的RL可以获得更高收益 [1]
 - 小模型RL不容易出现Long CoT [2]

[1] SFT Memorizes, RL Generalizes: A Comparative Study of Foundation Model Post-training. <https://arxiv.org/pdf/2501.17161>

[2] Demystifying Long Chain-of-Thought Reasoning in LLMs. <https://arxiv.org/pdf/2502.03373>

RL能否带来推理能力提升?

观点: **RL**并不一定真正扩展了大模型的推理能力, 它更多是在提升采样效率、
强化 **base model** 已经会的高奖励推理路径



Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model? <https://arxiv.org/pdf/2504.13837>

LLM是记忆还是推理?

GSM-NoOp

Oliver picks 44 kiwis on Friday. Then he picks 58 kiwis on Saturday. On Sunday, he picks double the number of kiwis he did on Friday, but five of them were a bit smaller than average. How many kiwis does Oliver have?

o1-mini: Oliver picks kiwis over three days as follows:

- Friday: 44 kiwis
- Saturday: 58 kiwis
- Sunday: Double the number he picked on Friday, which is $2 \times 44 = 88$ kiwis

However, on Sunday, 5 of these kiwis were smaller than average. We need to subtract them from the Sunday total: 88 (Sunday's kiwis) $- 5$ (smaller kiwis) $= 83$ kiwis. Now, summing up the kiwis from all three days: 44 (Friday) $+ 58$ (Saturday) $+ 83$ (Sunday) $= 185$ kiwis. Oliver has a total of 185 kiwis.

Llama3-8B: Let's break it down step by step:

Oliver picks 44 kiwis on Friday. He picks 58 kiwis on Saturday. On Sunday, he picks double the number of kiwis he did on Friday, which is $2 \times 44 = 88$ kiwis.

Five of the kiwis picked on Sunday are a bit smaller than average, so we subtract 5 from the total number of kiwis picked on Sunday: $88 - 5 = 83$ kiwis. Now, let's add up the total number of kiwis Oliver has:


44 (Friday) $+ 58$ (Saturday) $+ 83$ (Sunday) $= 185$ kiwis

So, Oliver has 185 kiwis in total.

新趋势：智能体时代



回顾：什么是智能体？



“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*

智能体是一个能够感知环境，进行推理规划，并采取行动与环境进行交互的系统

回顾：什么是智能体？

Knowledge


- 游戏规则
- 经验

Reasoning

- 下一步怎么操作？

Learning

- 玩的越多能力越强



逻辑智能体

基于人工定义的符号与逻辑规则


关键案例

先驱：Shakey the Robot (1969)

第一个能够推理自身行为，并与物理世界交互的通用移动机器人，能够自主规划路径和动作完成“推箱子”任务

优势：严谨的逻辑性、可解释性

局限：缺乏泛化能力，感知能力弱



强化学习智能体

基于强化学习与环境海量交互试错

关键案例

- 围棋：AlphaGo
- 游戏智能体
 - Atari, Open-AI Five, AlphaStar...

优势：封闭环境中掌握超越人类的复杂技能

局限：样本效率低，难以在不同场景迁移



两个角度理解智能体

LLM-first view: 让我们把LLM变成Agent

- 关键：如何在LLM的基础上一步步升级，从chatbot到reasoner再到Agent

Agent-first view: 让我们把LLM放到Agent架构中

- 关键：Agent本身有哪些挑战？LLM的引入可以帮助缓解哪些挑战，又引入了哪些新的挑战？

LLM-First View: 从LLM到Agent,
是如何发展的?

大模型推理的局限



提问：“12345 乘以 67890 等于多少？”



LLM (可能) 回答:
“12345 乘以 67890 等于 838102050。”



这张图片里面有几个红色的球，几个黄色的球

已完成思考

用户现在需要数清楚图片里红色和黄色的球。首先看红色的球，仔细数的话，中间聚集的红色球有 8 个。然后黄色的球，包括周围分散的和中间的，数出来有 11 个。所以要一个一个仔细确认，确保没有重复或遗漏。现在开始数：红色球：从中间的簇里数，能看到 8 个红色的。黄色球：周围分散的加上中间的，一共 11 个。所以结论是红色球有 8 个，黄色球有 11 个。

仔细观察图片，红色的球有 8 个，黄色的球有 11 个。

可靠性不足，模型幻觉

概率近似正确本质，难以保障严谨性

Human + Tool Use

- As humans, we have limited time and memory, feel tired, and have emotions.
- **Human + tool use**
 - Enhanced scalability
 - Improved consistency
 - Greater interpretability
 - Higher capacity and productivity



LLM + Tool Use

- Just like humans, LLMs suffer from the similar limitations. But in the same way,
- **LLMs + tool use**
 - Enhanced scalability
 - Improved consistency
 - Greater interpretability
 - Higher capacity and productivity




Hugging Face



matplotlib

🔍 Search Google or type a URL

LLM + Tool Use



Neuro-Symbolic Learning思想的体现：不再只是依赖单一神经网络，而是通过神经网络大模型 + 外部工具，构建更强大的AI系统

LLM + Tool Use

模型怎么知道有哪些工具，什么时候该调用工具？

纯 Prompt 工程

你可以使用以下工具：

- search(query): 搜索网页
- calculator(expr): 计算数学表达式


请按照 Thought → Action → Observation 的格式回答

模型并没有真正“学会”调用工具，只是在模式匹配，见过类似格式就照着输出

LLM + Tool Use: Prompt

模型怎么知道有哪些工具，什么时候该调用工具？


```
tools = [  
  {  
    "name": "get_weather",  
    "description": "查询指定城市的当前天气情况，包括温度、天气状况和湿度",  
    "input_schema": {  
      "type": "object",  
      "properties": {  
        "city": {  
          "type": "string",  
          "description": "城市名称，例如：上海、北京、广州"  
        },  
        "unit": {  
          "type": "string",  
          "enum": ["celsius", "fahrenheit"],  
          "description": "温度单位，默认摄氏度"  
        }  
      },  
      "required": ["city"]  
    }  
  },  
]
```



ReAct框架：推理→调用→观察→再推理

LLM + Tool Use: Prompt

复杂任务不能一步完成，需要拆解成多个子任务，模型需要决定先做什么、后做什么、哪一步需要什么工具、出错后如何调整



LLM + Tool Use: Prompt

帮我规划一个五天的京都旅行，我喜欢历史文化，预算中等，下个月出发

【Planner 输出】

分析用户需求：

- 目的地：京都
- 时长：5天
- 偏好：历史文化
- 预算：中等
- 时间：下个月

生成执行计划：

- Task 1: 查询下个月京都天气与最佳游览时段
- Task 2: 搜索京都历史文化类核心景点
- Task 3: 查询中等预算酒店选项（距核心景区步行可达）
- Task 4: 查询北京→京都往返机票价格
- Task 5: 根据景点地理位置，规划5天合理游览路线
- Task 6: 整合所有信息，生成完整旅行方案

ReAct v.s. Plan & Solve

ReAct (边想边做):

查天气

↓ 看到结果再决定

搜景点

↓ 看到景点再决定

查酒店

↓

...

Plan & Execute (先想后做):

┌ Planner 全局规划 ┐

Task1 Task2 Task3...

└──────────────────┘

↓

Executor 按序执行

可并行、可回溯、
任务间有明确依赖关系

LLM + Tool Use: SFT

模型怎么知道有哪些工具，什么时候该调用工具？

监督微调

构造训练数据：

用户："上海今天天气？"

标准输出：{
 "tool": "get_weather",
 "parameters": {"city": "上海"}
}

用户："1加1等于几？"

标准输出：直接回答"2"，不调用任何工具

训练数据来源主要有三类：

- 人工标注：人工写出"什么情况调什么工具"的示范数据
- 合成数据：用强模型（GPT-4）自动生成大量工具调用轨迹，再用弱模型学习
- 真实使用日志：收集用户实际使用工具的成功案例

LLM + Tool Use: SFT

The New England Journal of Medicine is a registered trademark of [QA("Who is the publisher of The New England Journal of Medicine?") → Massachusetts Medical Society] the MMS.

Out of 1400 participants, 400 (or [Calculator(400 / 1400) → 0.29] 29%) passed the test.

The name derives from "la tortuga", the Spanish word for [MT("tortuga") → turtle] turtle.

The Brown Act is California's law [WikiSearch("Brown Act") → The Ralph M. Brown Act is an act of the California State Legislature that guarantees the public's right to attend and participate in meetings of local legislative bodies.] that requires legislative bodies, like city councils, to hold their meetings open to the public.

Your task is to add calls to a Question Answering API to a piece of text. The questions should help you get information required to complete the text. You can call the API by writing "[QA(question)]" where "question" is the question you want to ask. Here are some examples of API calls:

Input: Joe Biden was born in Scranton, Pennsylvania.

Output: Joe Biden was born in [QA("Where was Joe Biden born?")] Scranton, [QA("In which state is Scranton?")] Pennsylvania.

Input: Coca-Cola, or Coke, is a carbonated soft drink manufactured by the Coca-Cola Company.

Output: Coca-Cola, or [QA("What other name is Coca-Cola known by?")] Coke, is a carbonated soft drink manufactured by [QA("Who manufactures Coca-Cola?")] the Coca-Cola Company.

Input: x

Output:

LLM + Tool Use: RL

模型怎么知道有哪些工具，什么时候该调用工具？

强化微调


- ✅ 调用了正确工具，参数合法，最终答对 → 高奖励
- ⚠️ 调用了工具但参数错误，答案偏差 → 低奖励
- ❌ 不该调用工具却调了，浪费 token → 负奖励
- ❌ 该调工具却没调，靠胡编蒙混 → 负奖励

目前主流趋势：监督微调+强化微调

Agentic RL

The Landscape of Agentic Reinforcement Learning for LLMs: A Survey

<https://www.alphaxiv.org/abs/2509.02547>



Agentic RL

VERL: <https://github.com/verl-project/verl>

👋 Hi, everyone! verl is a RL training library initiated by ByteDance Seed team and maintained by the verl community.

 Ask DeepWiki  Stars 22k  Follow @verl_project  Slack verl  EuroSys  Paper  documentation  微信



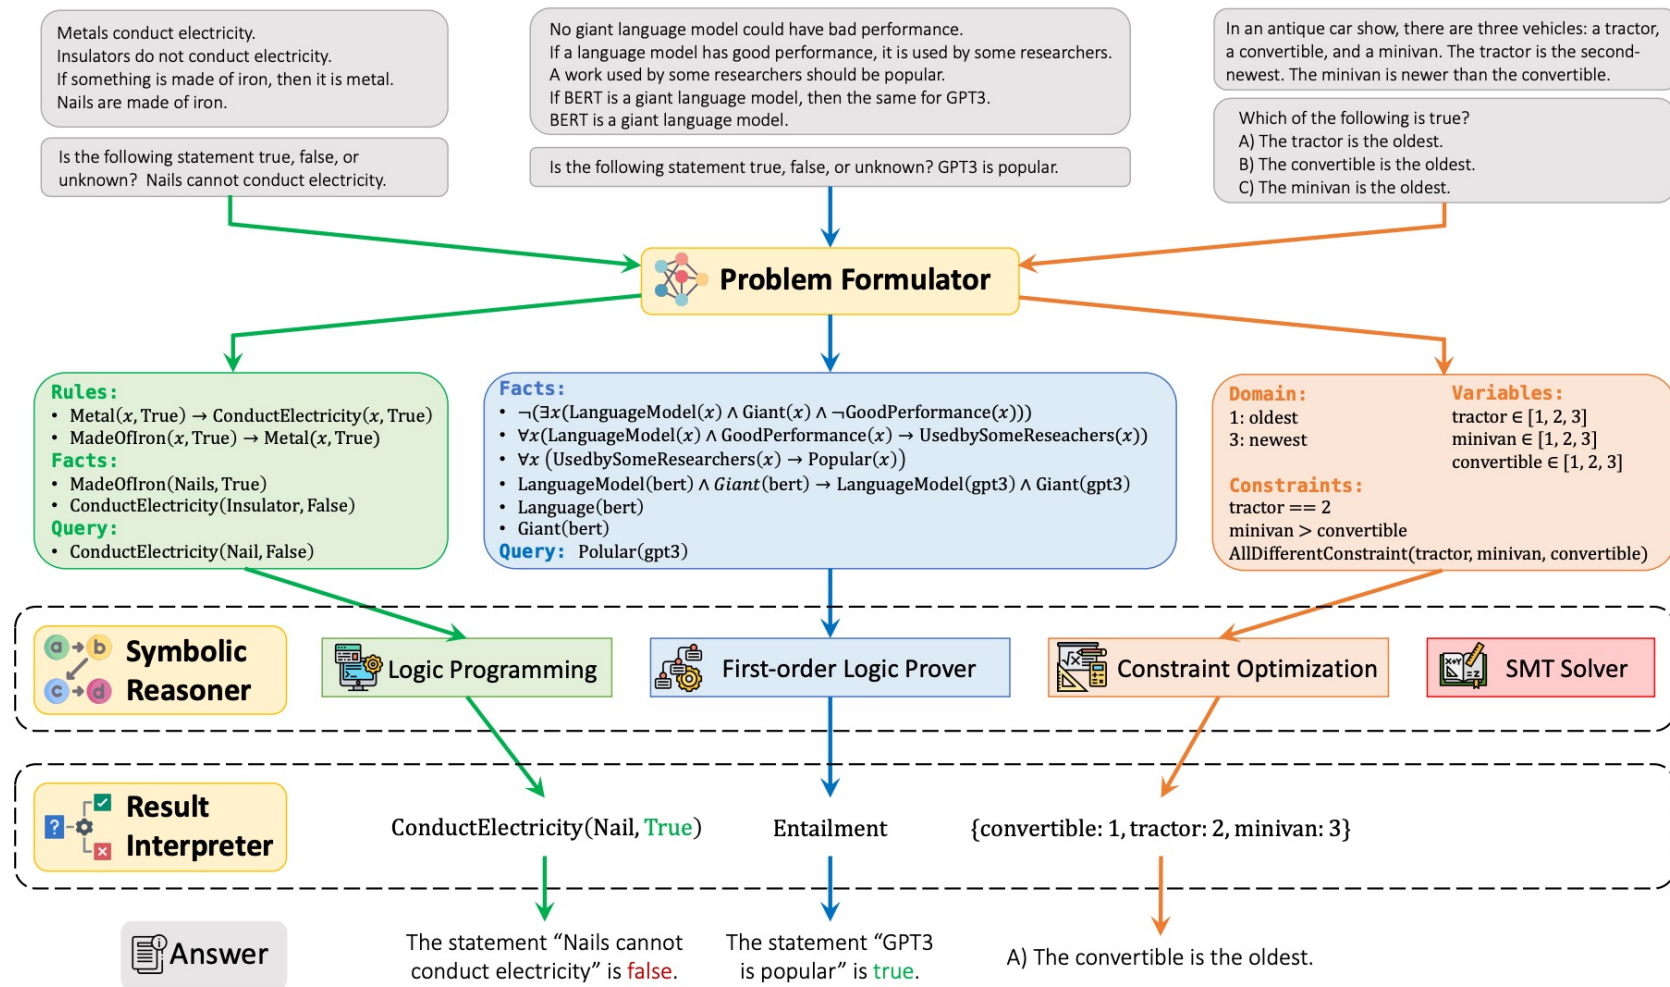
verl: Volcano Engine Reinforcement Learning for LLMs

例子：Logical Reasoning

Convert **Natural Language** into **Formal Language**, then call a **symbolic solver**, such as

- First-order Logic Prover
- SMT Solver
- Constraint Optimizer

...



例子：Math Reasoning

Parse an **natural language** input into a **declarative task specification** (a set of logical constraints), then use a **SAT solver** (e.g., Z3) to solve the problem

Input

Each of five students—Hubert, Lori, Paul, Regina, and Sharon—will visit exactly one of three cities—Montreal, Toronto, or Vancouver, according to the following conditions: Sharon visits a different city than Paul. Hubert visits the same city as Regina. Lori visits Montreal or else Toronto. If Paul visits Vancouver, Hubert visits Vancouver with him. Each student visits one of the cities with at least one of the other four students.

Question: Which one of the following must be true?

(A) If any of the students visits Montreal, Lori visits Montreal. (B) [...]

Chain-of-Thought Prompting (imperative specification)

Specification

We know each student visits one of the cities with at least one of the other four students. We know there are five students and three cities. So there must be three students visiting the one city and two other students visiting another city.

Let's consider option (A).

Assume someone visits Montreal, but Lori does not visit Montreal.

We know Lori visits Montreal or else Toronto. So Lori visits Toronto.

Assume Sharon visits Toronto with Lori.

We know Sharon visits a different city than Paul. So Paul has to visit Montreal.

Hubert and Regina can visit Montreal with Paul with no conflicts. So Lori does not necessarily visit Montreal. This statement is False.

The LLM parses the question, plans the reasoning, and executes it all in the CoT (shown by dashed arrows)


Satisfiability-Aided LM (ours; declarative specification)

Specification

```
students = [Hubert, Lori, Paul, Regina, Sharon]
cities = [Montreal, Toronto, Vancouver]
visits = Function(students, cities)
# Sharon visits a different city than Paul
1 visits(Sharon) != visits(Paul)
# Lori visits Montreal or else Toronto
2 Or(visits(Lori) == Montreal, visits(Lori) == Toronto)
# Each student visits one of the cities with at least one other student
3 ForAll([s1], Exists([s2], And(s2 != s1, visits(s1) == visits(s2))))
.....
# (A)
4 solve(Implies(Exists([s], visits(s) == Montreal), visits(L) == Montreal))
```

The LLM only parses the question to a problem specification in this step

Z3
SAT Solver



→ False

A SAT solver generates and executes a proof plan using automated theorem proving

例子：Math Reasoning

Integrate **natural language reasoning** with the utilization of **external tools (e.g., computation libraries, symbolic solvers)**

Problem: Suppose that the sum of the squares of two complex numbers x and y is 7 and the sum of their cubes is 10. List all possible values for $x + y$, separated by commas.

We are given that $x^2 + y^2 = 7$ and $x^3 + y^3 = 10$. We can factor $x^3 + y^3$ to get $(x + y)(x^2 - xy + y^2)$.

Thus, we can write $10 = (x + y)(7)$. It follows that $x + y$ must divide 10.

The only possibilities are that $x + y$ is 1, 2, 5 or 10. ❌

```
import math
```

```
def solution():
```

```
    x = symbols('x'), y = symbols('y')
    x**2 + y**2 = 7
    x**3 + y**3 = 10
```

```
    x = (-1 + math.sqrt(17)) / 2
    y = (-1 - math.sqrt(17)) / 2
```

```
    # Get all possible values for x+y
    result = [x + y, -x - y]
```

```
>>> SyntaxError: cannot assign to expression. ❌
```

Let's use `SymPy` to calculate and print all possible values for $x + y$,

```
from sympy import symbols, Eq, solve, simplify
```

```
def possible_values():
```

```
    x, y = symbols('x y')
    eq1 = Eq(x**2 + y**2, 7)
    eq2 = Eq(x**3 + y**3, 10)
    solutions = solve((eq1, eq2), (x, y))
    sums = [simplify(sol[0] + sol[1]) for sol in solutions]
    return sums
print(possible_values())
```

```
>>> [-5, -5, 1, 1, 4, 4]
```


Removing duplicates, the possible values for $x + y$ are `\boxed{-5, 1, 4}`. ✅

r_1

a_1

o_1

r_2




Rationale

(a) Rationale-based



Program

(b) Program-based



(c) Tool-integrated Reasoning
(Format used by ToRA)

例子：Math Reasoning

Question

Raymond and Samantha are cousins. Raymond was born 6 years before Samantha. Raymond had a son at the age of 23. If Samantha is now 31, how many years ago was Raymond's son born?

Ground Truth

When Raymond's son was born Samantha was $23 - 6 = 17$ years old. Thus it has been $31 - 17 = 14$ years since Raymond's son was born.
Answer: 14

Chain of Thought

Samantha is $31 - 6 = 25$ years old.
Raymond's son was born $25 - 23 = 2$ years ago.

Answer: 2



Prolog Generation

```
:- use_module(library(clpq)).  
  
age_diff(raymond, samantha, 6).  
  
age(raymond, 23).  
age(samantha, 31).  
  
solve(Years_ago) :-  
    age_diff(raymond, samantha, Age_diff),  
    age(raymond, Raymond_age),  
    age(samantha, Samantha_age),  
    {Raymond_son_age = Raymond_age - Age_diff},  
    {Years_ago = Samantha_age - Raymond_son_age}.
```



Prolog
Interpreter

Answer:

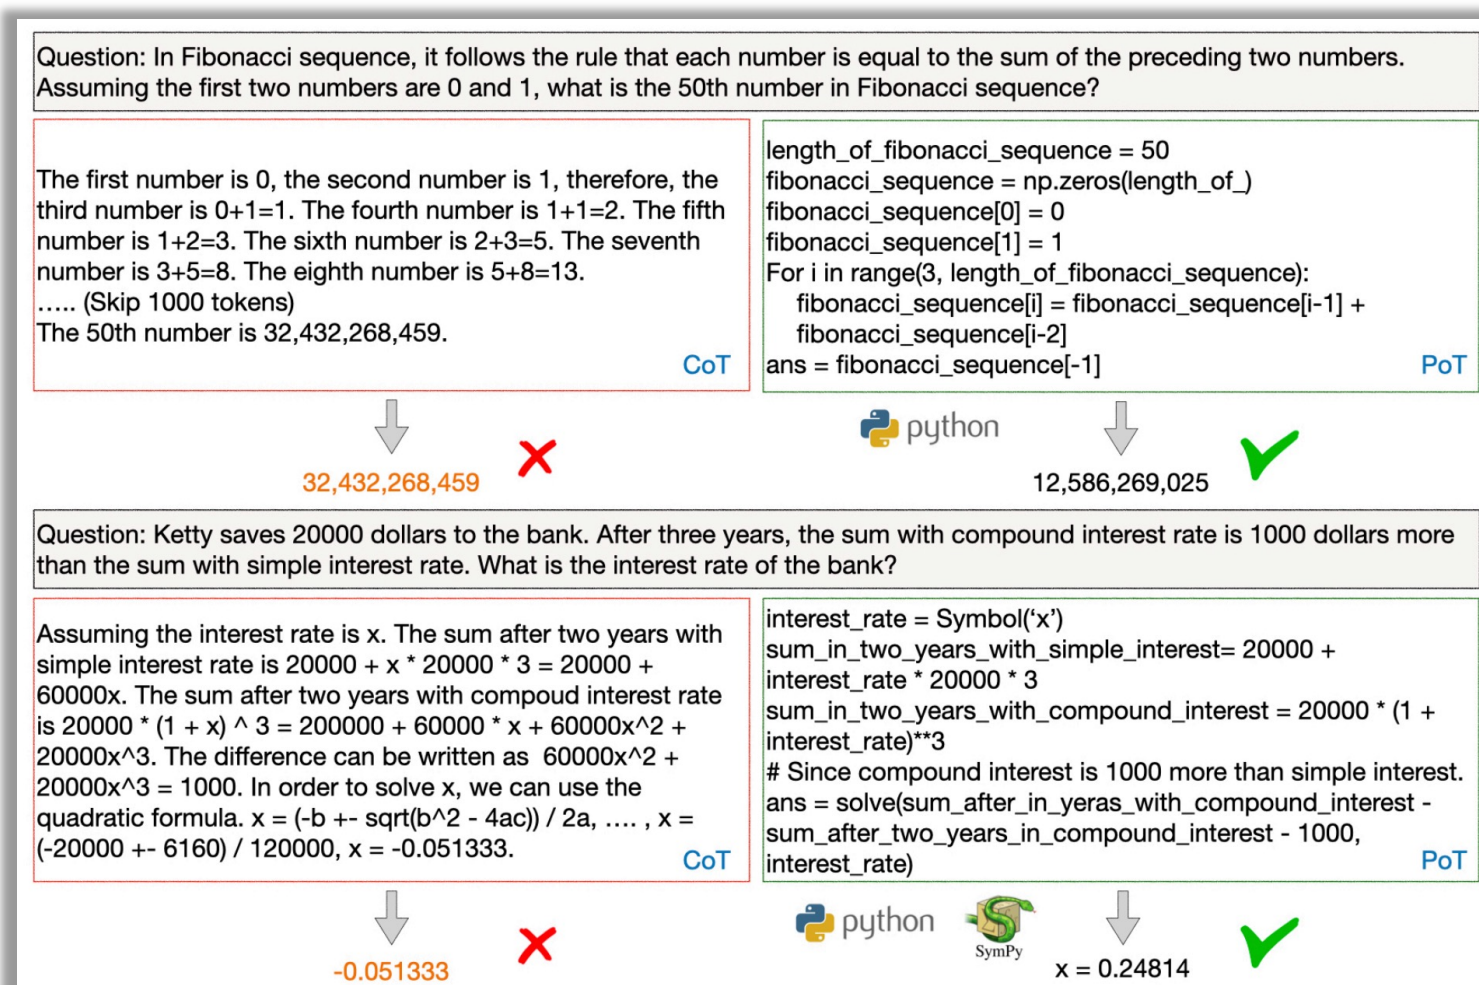
14



例子：Math Reasoning

Program of Thought

Convert the reasoning process expressed in natural language into a **programming language** (e.g., Python, etc.), and then **execute the program interpreter**




[1] [PAL: Program-aided Language Models](#). Gao et al., ICML 2023.

[2] [Program of Thoughts Prompting: Disentangling Computation from Reasoning for Numerical Reasoning Tasks](#). Chen et al., TMLR 2023.

例子：表格推理


- **Parse Stage:** LLM maps input to **Python/SQL** algorithm
- **Execution Stage:** A program interpreter executes the program to derive the answer



例子：多模态推理

Multi-Modal Reasoning

LLMs generate code to manipulate images for visual reasoning



```
# Define the cropping coordinates (x1, y1, x2, y2)
x1, y1, x2, y2 = 1750, 1950, 1850, 2600

# Crop the image
cropped_image = image[y1:y2, x1:x2]

# Apply zoom (resize the cropped image to enhance details)
zoom_factor = 2
zoomed_image = cv2.resize(cropped_image, None,
fx=zoom_factor, fy=zoom_factor,
interpolation=cv2.INTER_CUBIC)

# Save the processed image with a random suffix
random_suffix = random.randint(1000, 9999)
processed_filename = f"224_{random_suffix}.jpg"
processed_path = "/mnt/data/temp_processed_images/" +
processed_filename
cv2.imwrite(processed_path, zoomed_image)
```

例子：多模态推理

LLMs generate a program to compose a diverse set of modules for visual reasoning



Question: Are there both ties and glasses in the picture?

LLM





Image Understanding	Loc OWL-ViT	FaceDet DSFD (pyt)	Seg MaskFormer	Select CLIP-ViT	Classify CLIP-ViT	Vqa ViLT
	Replace Stable Diffusion	ColorPop PIL.convert() cv2.grabCut()	BgBlur PIL.GaussianBlur() cv2.grabCut()	Tag PIL.rectangle() PIL.text()	Emoji AugLy (pyt)	
Image Manipulation	Crop PIL.crop()	CropLeft PIL.crop()	CropRight PIL.crop()	CropAbove PIL.crop()	CropBelow PIL.crop()	
	List GPT3		Arithmetic & Logical	Eval eval()	Count len()	Result dict()
Knowledge Retrieval						

Model Zoo

```
BOX0=Loc(image=IMAGE, object='ties')
ANSWER0=Count(box=BOX0)
BOX1=Loc(image=IMAGE,
objects='glasses')
Answer1=Count(box=BOX1)
ANSWER2=Eval(“‘yes’ if {ANSWER > 0) and
{ANSWER1}>0 else ‘no’”)
RESUL=ANSWER2
```

例子：多模态推理



例子：Embodied Planning

For planning tasks, first **generate the PDDL descriptions**, then use a **PDDL Planner** (e.g., FastDownward) to produce the plan

A Failure Example of GPT-4 in Planning


Problem (P1): You have 5 blocks. One cannot place more than one block on another block. b5 is on top of b3. b4 is on top of b2. b2 is on top of b1. b3 is on top of b4. b1 is on the table. b5 is clear. Your arm is empty.

Your goal is to move the blocks.

b1 should be on top of b2.

b3 should be on top of b5.


b4 should be on top of b1.




GPT-4:

1. Move b5 from b3 to the table.
2. Move b3 from b4 to b5.
3. Move b4 from b2 to the table.
4. Move b2 from b1 to b5.
5. Move b1 from the table to b2.
6. Move b4 from the table to b1.


Module Generated Text Provided Text



LLM-As-Planner



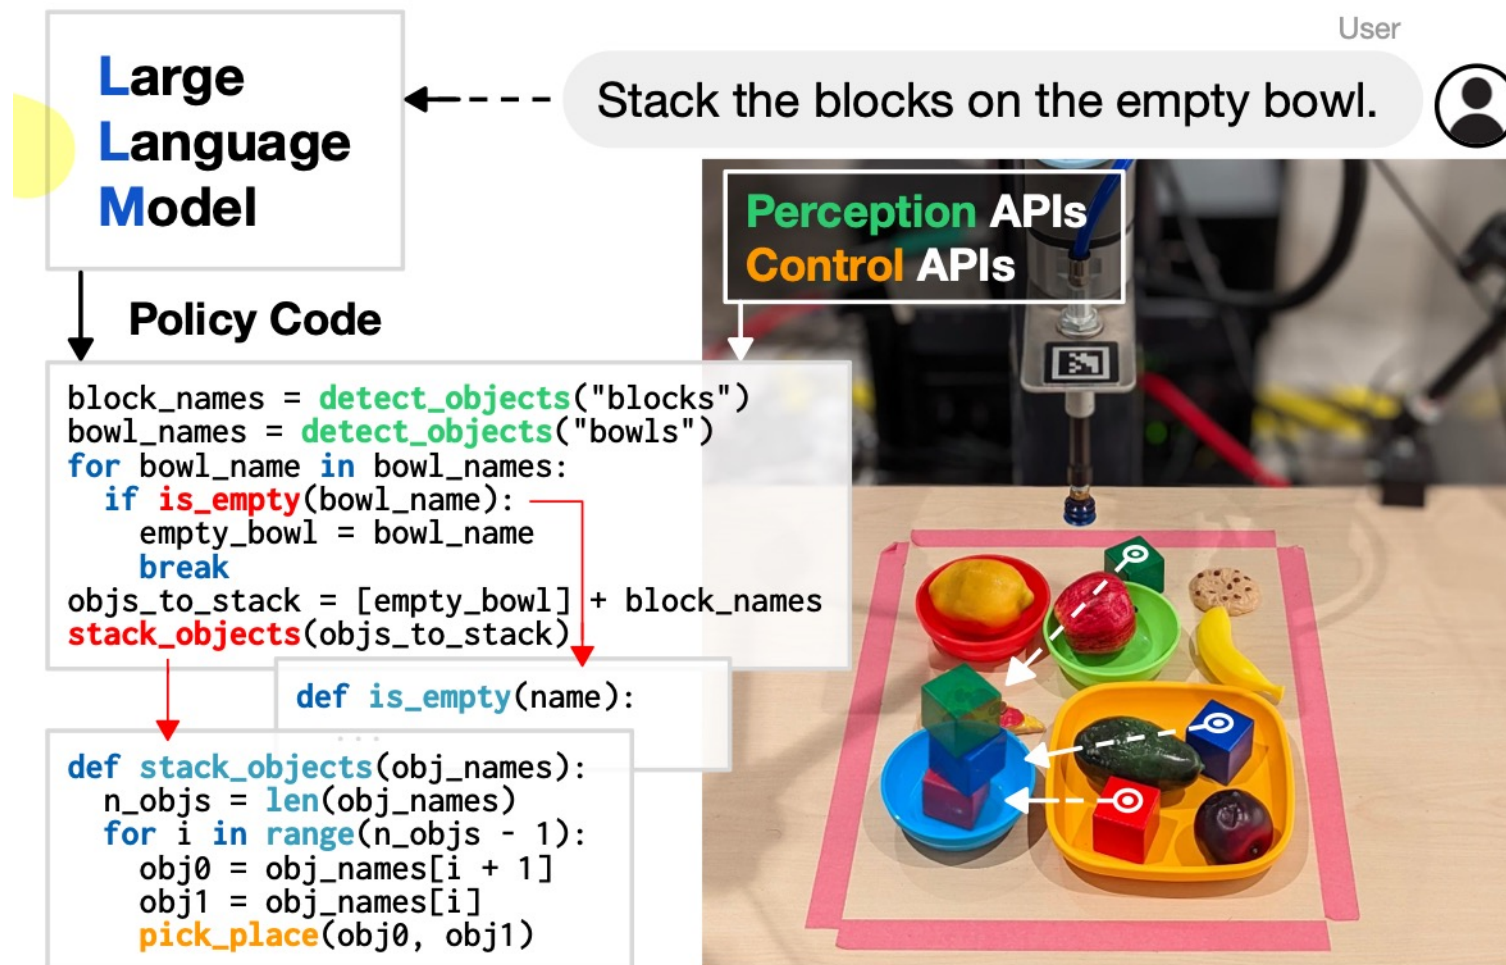
LLM-As-Planner (In-context Learning)



例子：Embodied Planning

Code as Policy

Robots can use LLMs to translate **natural language commands** into **robot policy code** which process **perception outputs**, **parameterize control primitives**, recursively generate code for **undefined functions**, and generalize to new tasks



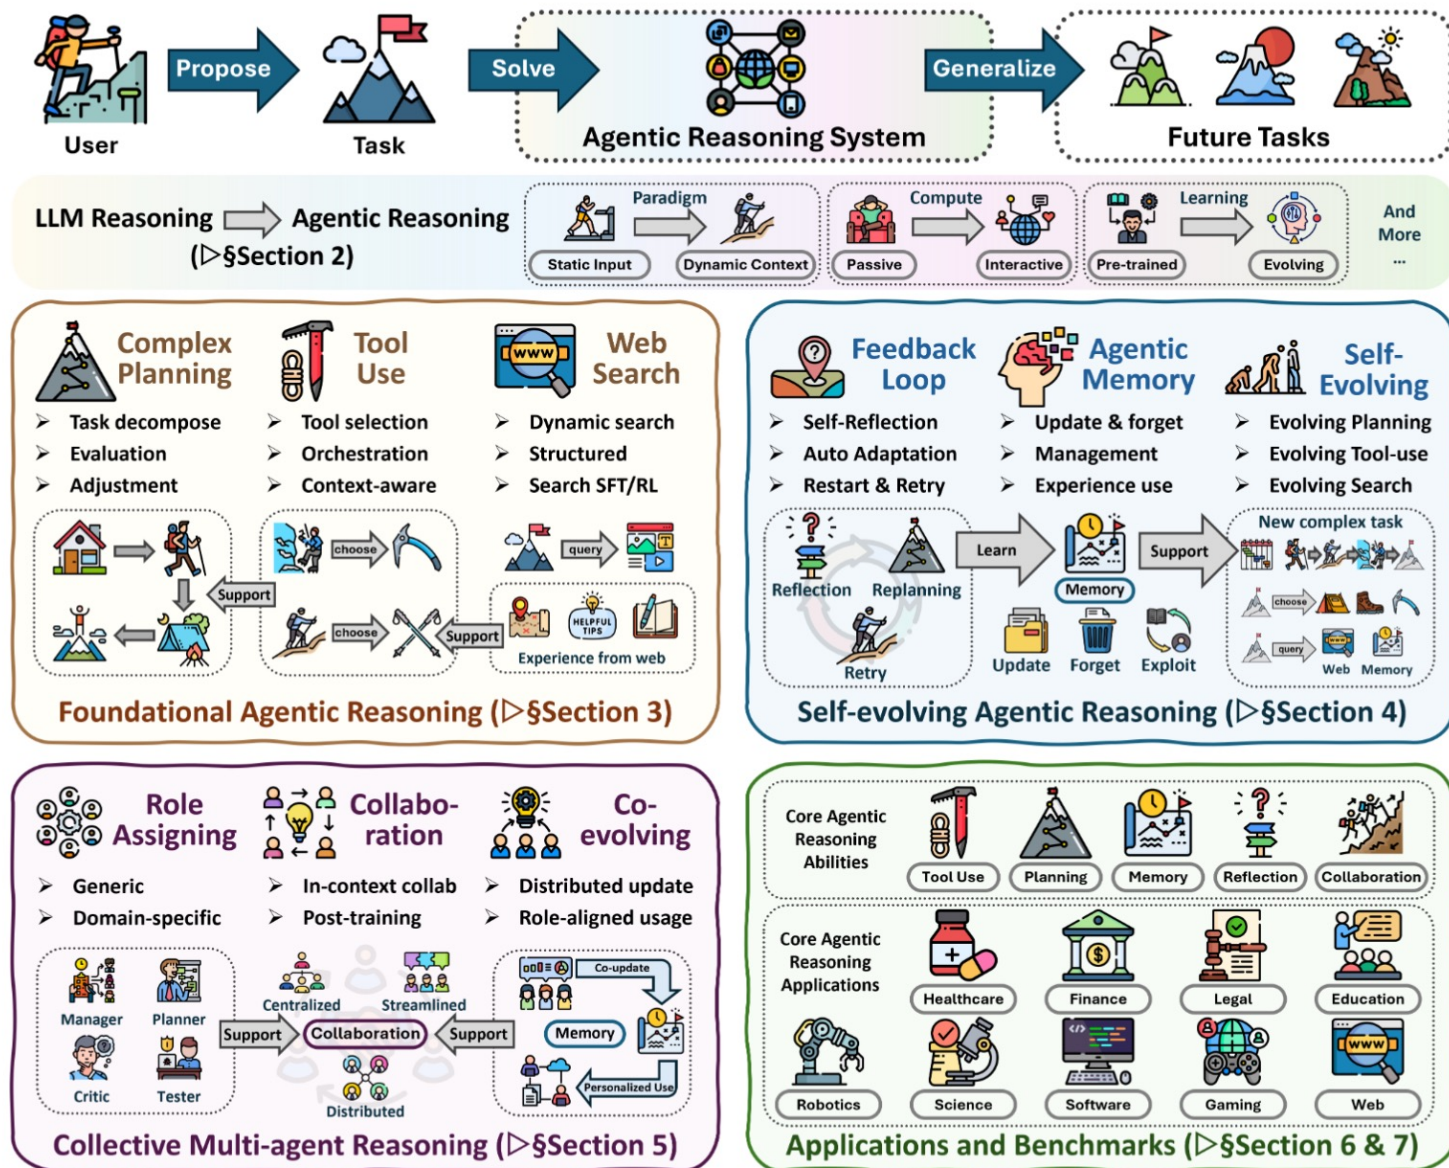
例子：Embodied Planning



Agentic Reasoning

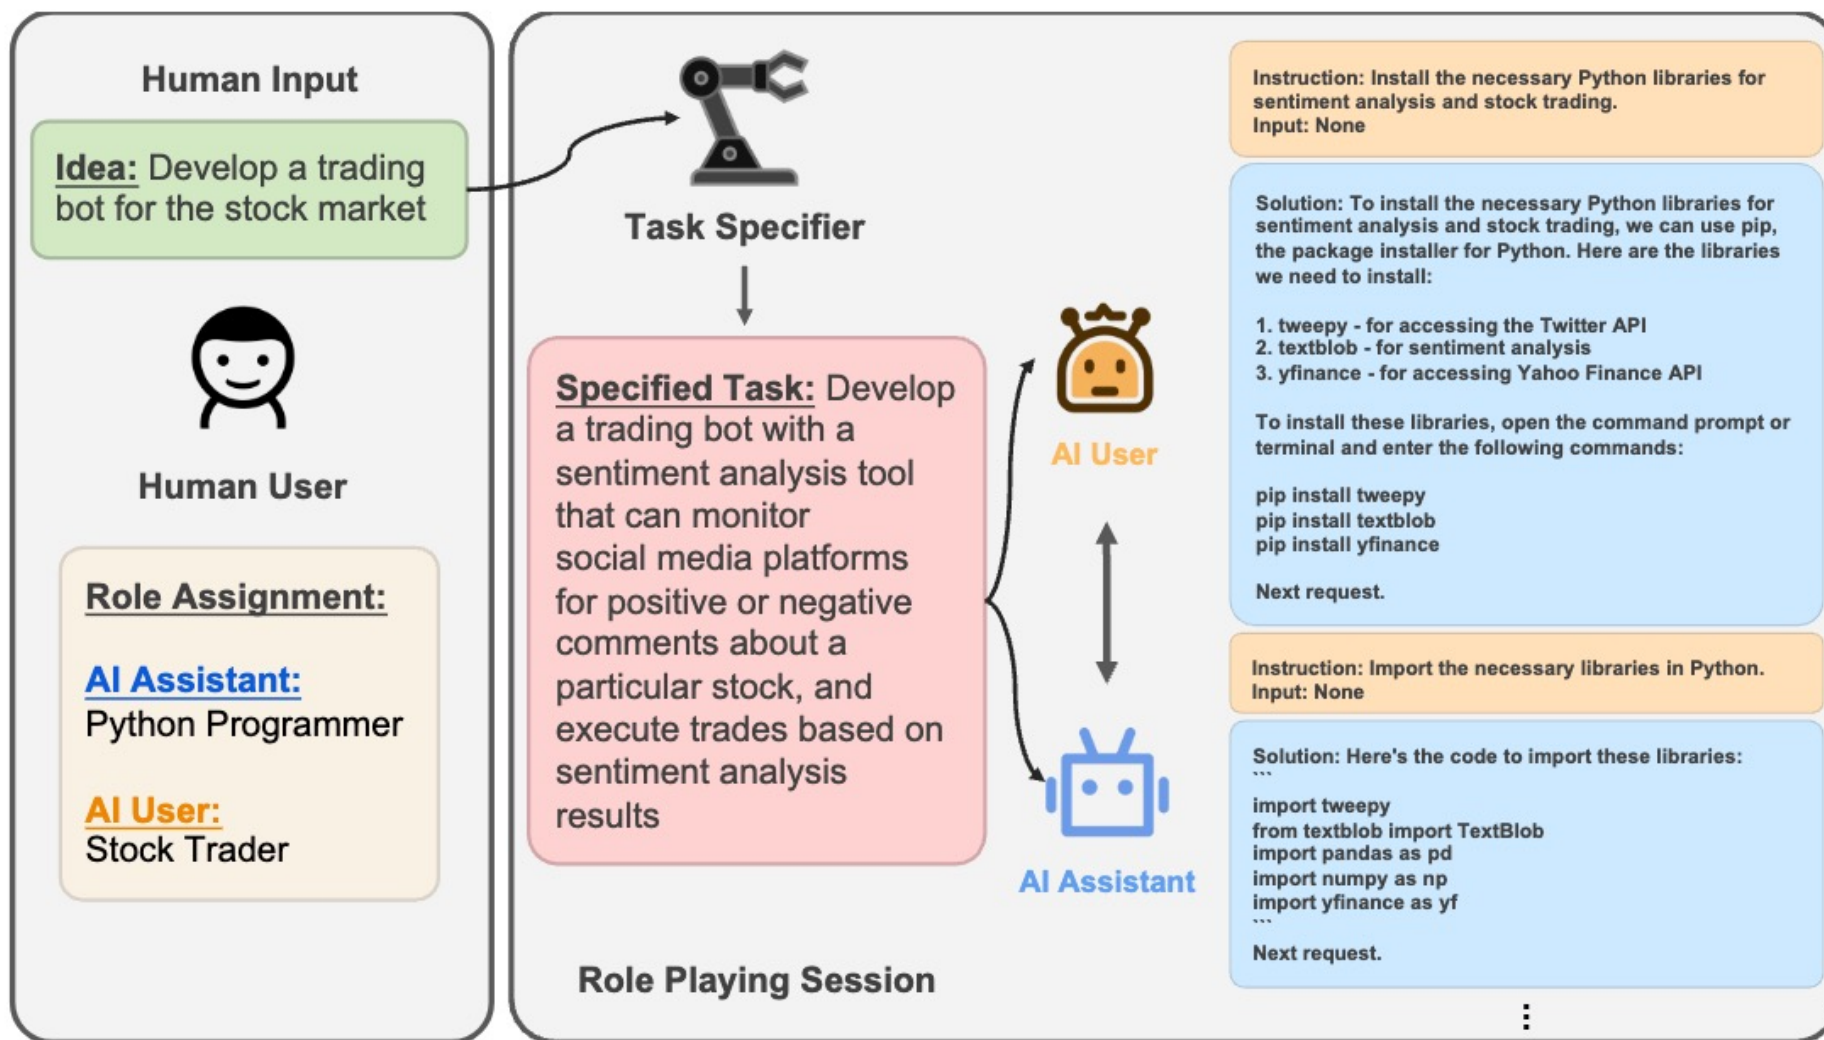
Agentic Reasoning for Large Language Models

<https://www.alphaxiv.org/abs/2601.12538>

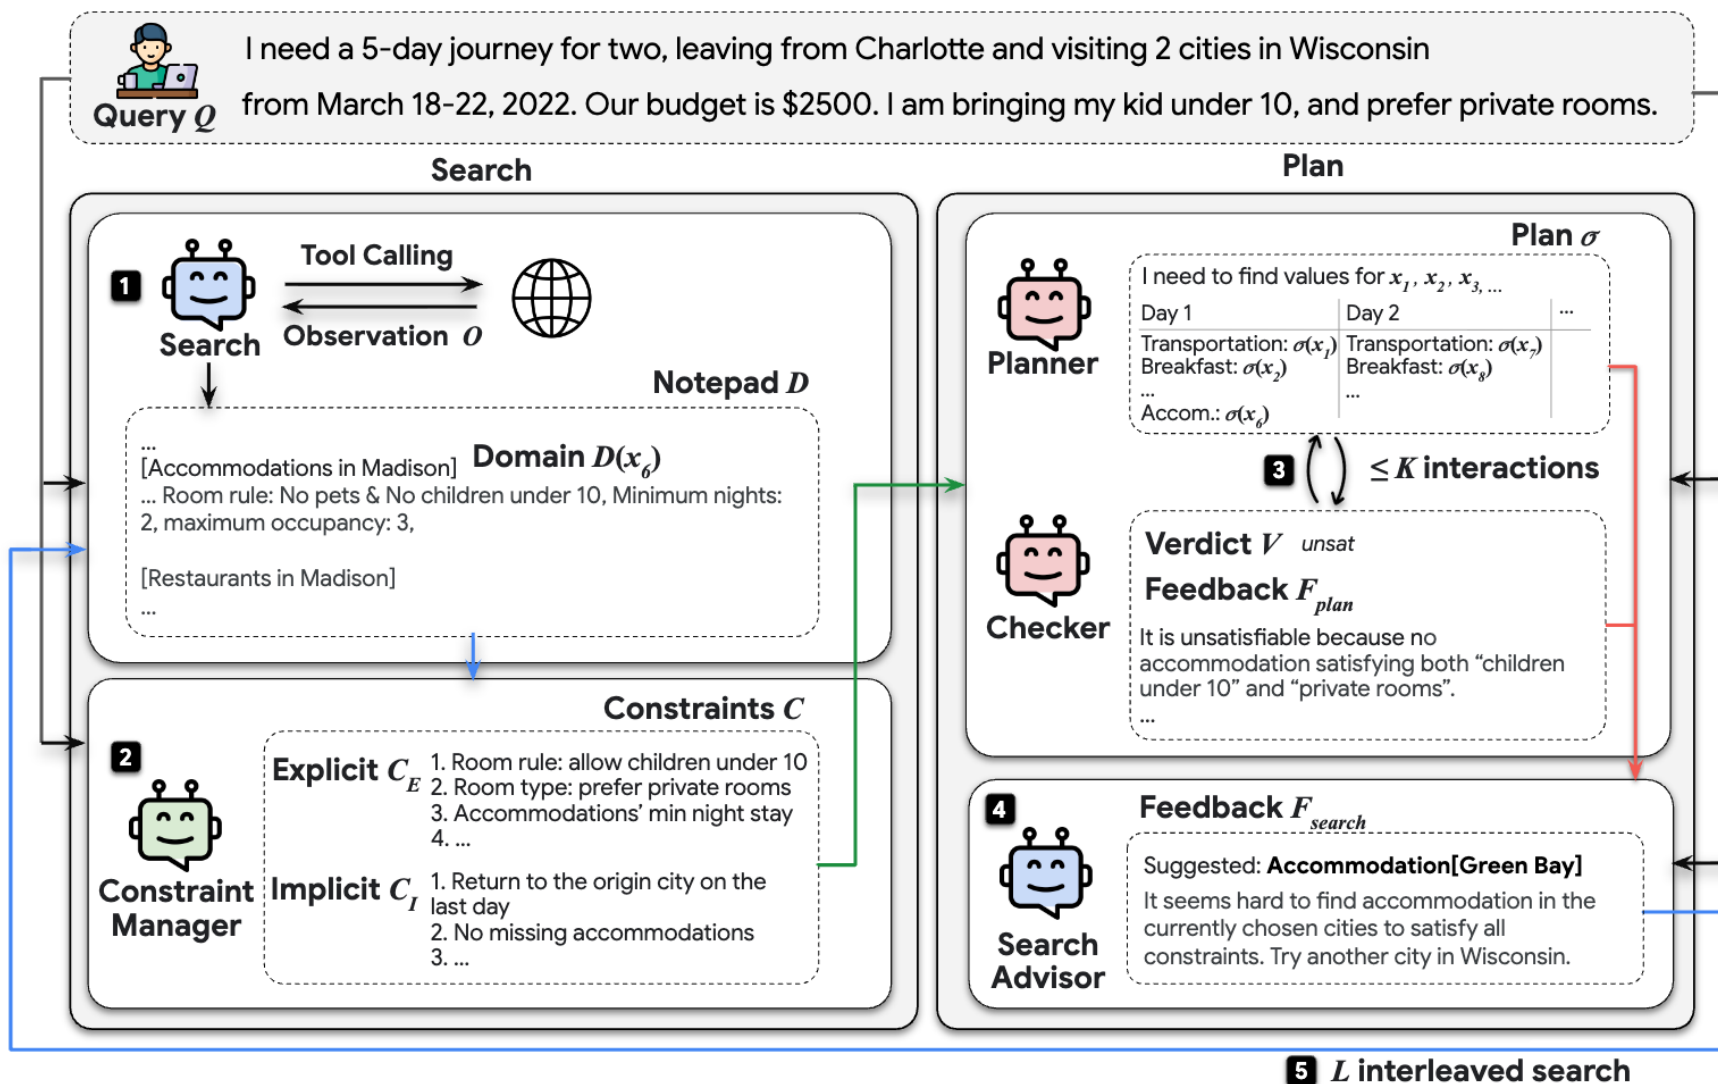


一个Agent不一定能解决问题

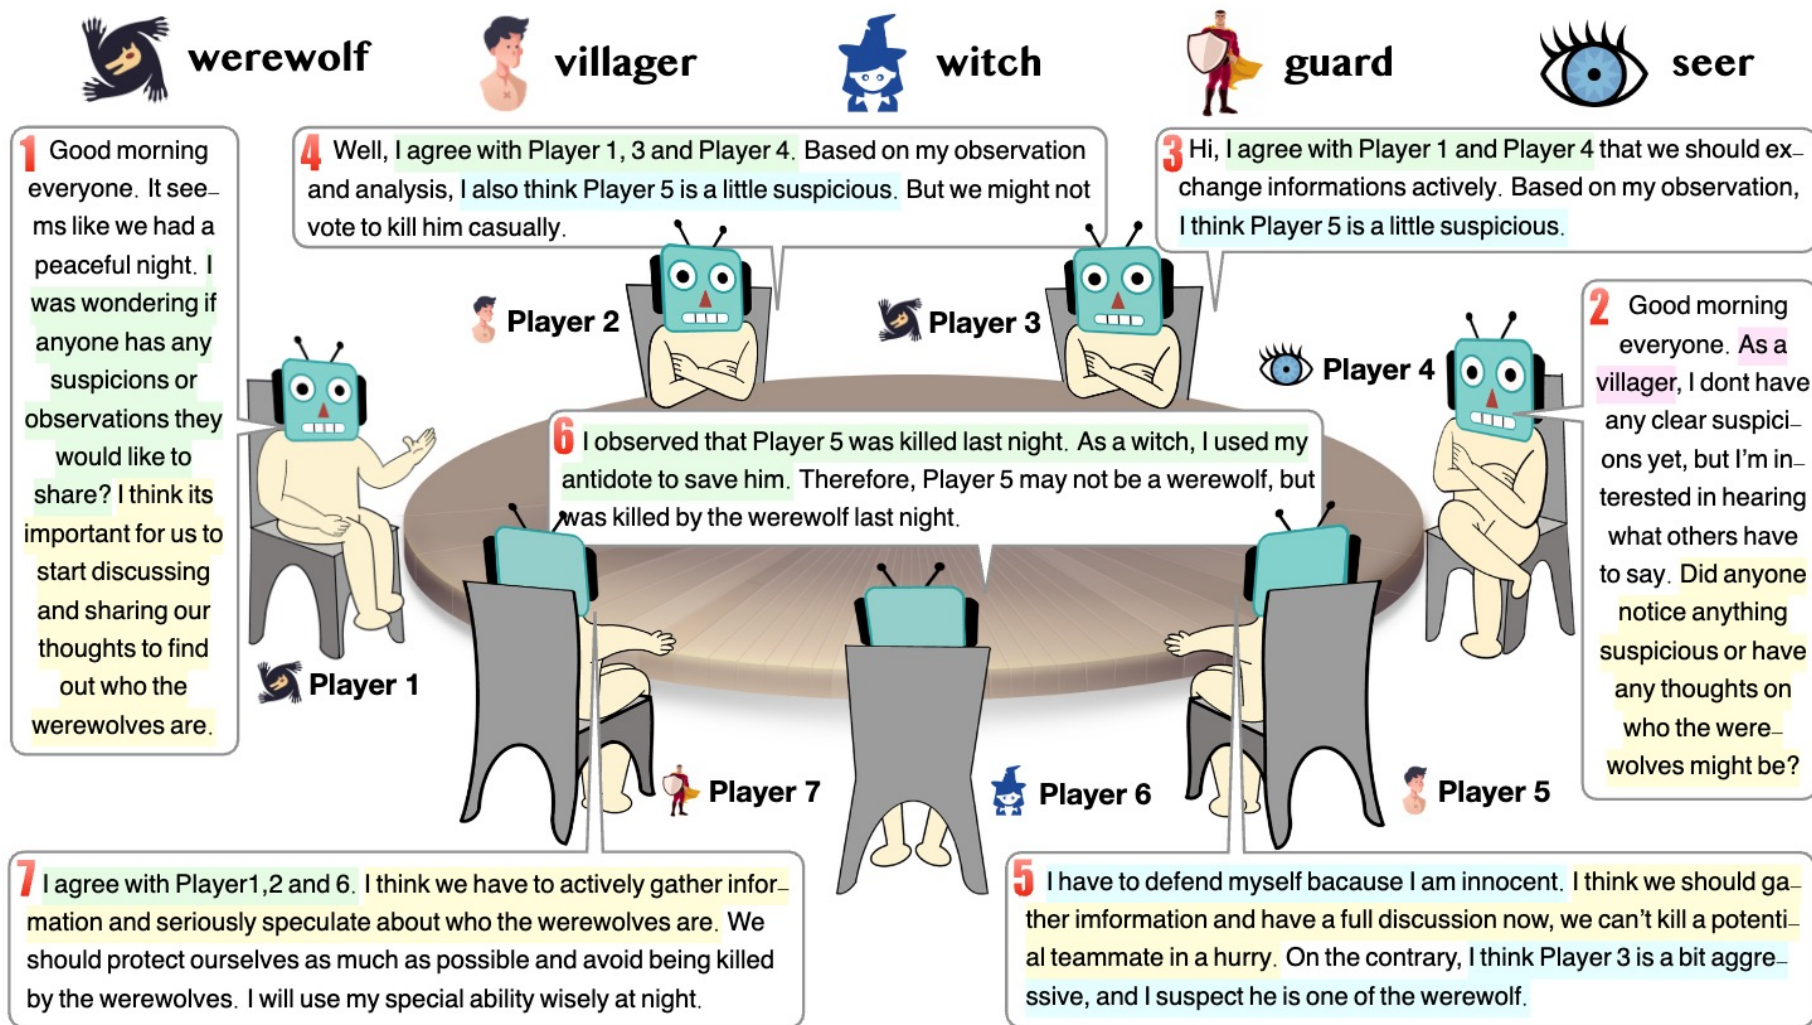
从单一Agent到Multi-Agent



从单一Agent到Multi-Agent




从单一Agent到Multi-Agent




Memory

Memory is everything. Without it
we are nothing.


— — Eric Kandel




Memory



Memory



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



SPUDCOMICS.COM

© 2010 LONNIE EASTERLING

THE TRAGEDY OF A THREE SECOND MEMORY

Memory

类比人类记忆，Agent的记忆分类

Type by content	Definition
Working Memory (工作记忆)	当前任务上下文 (Context)
Episodic Memory (情景记忆)	具体事件的记忆 (Experience)
Semantic Memory (语义记忆)	存储知识 (Knowledge)
Procedural Memory (程序记忆)	存储技能 (Skill)

- 情景记忆

长期·可提取

带时间戳的个人经历，“我在哪里、何时经历了什么”。
昨天的会议内容；上次去某地旅行的经历
- 语义记忆

长期·去时间化

事实与概念知识，脱离了习得时的情境，可直接提取。
巴黎是法国首都；水的化学式是 H₂O
- 程序记忆

长期·内隐

技能与操作流程，难以言表，靠身体执行而非意识提取。
骑自行车；打字；演奏乐器

Skill

Skill 是从历史经验中提炼出的、可复用的任务执行单元，封装了
"在什么情况下、执行什么操作、达到什么效果"的完整知识

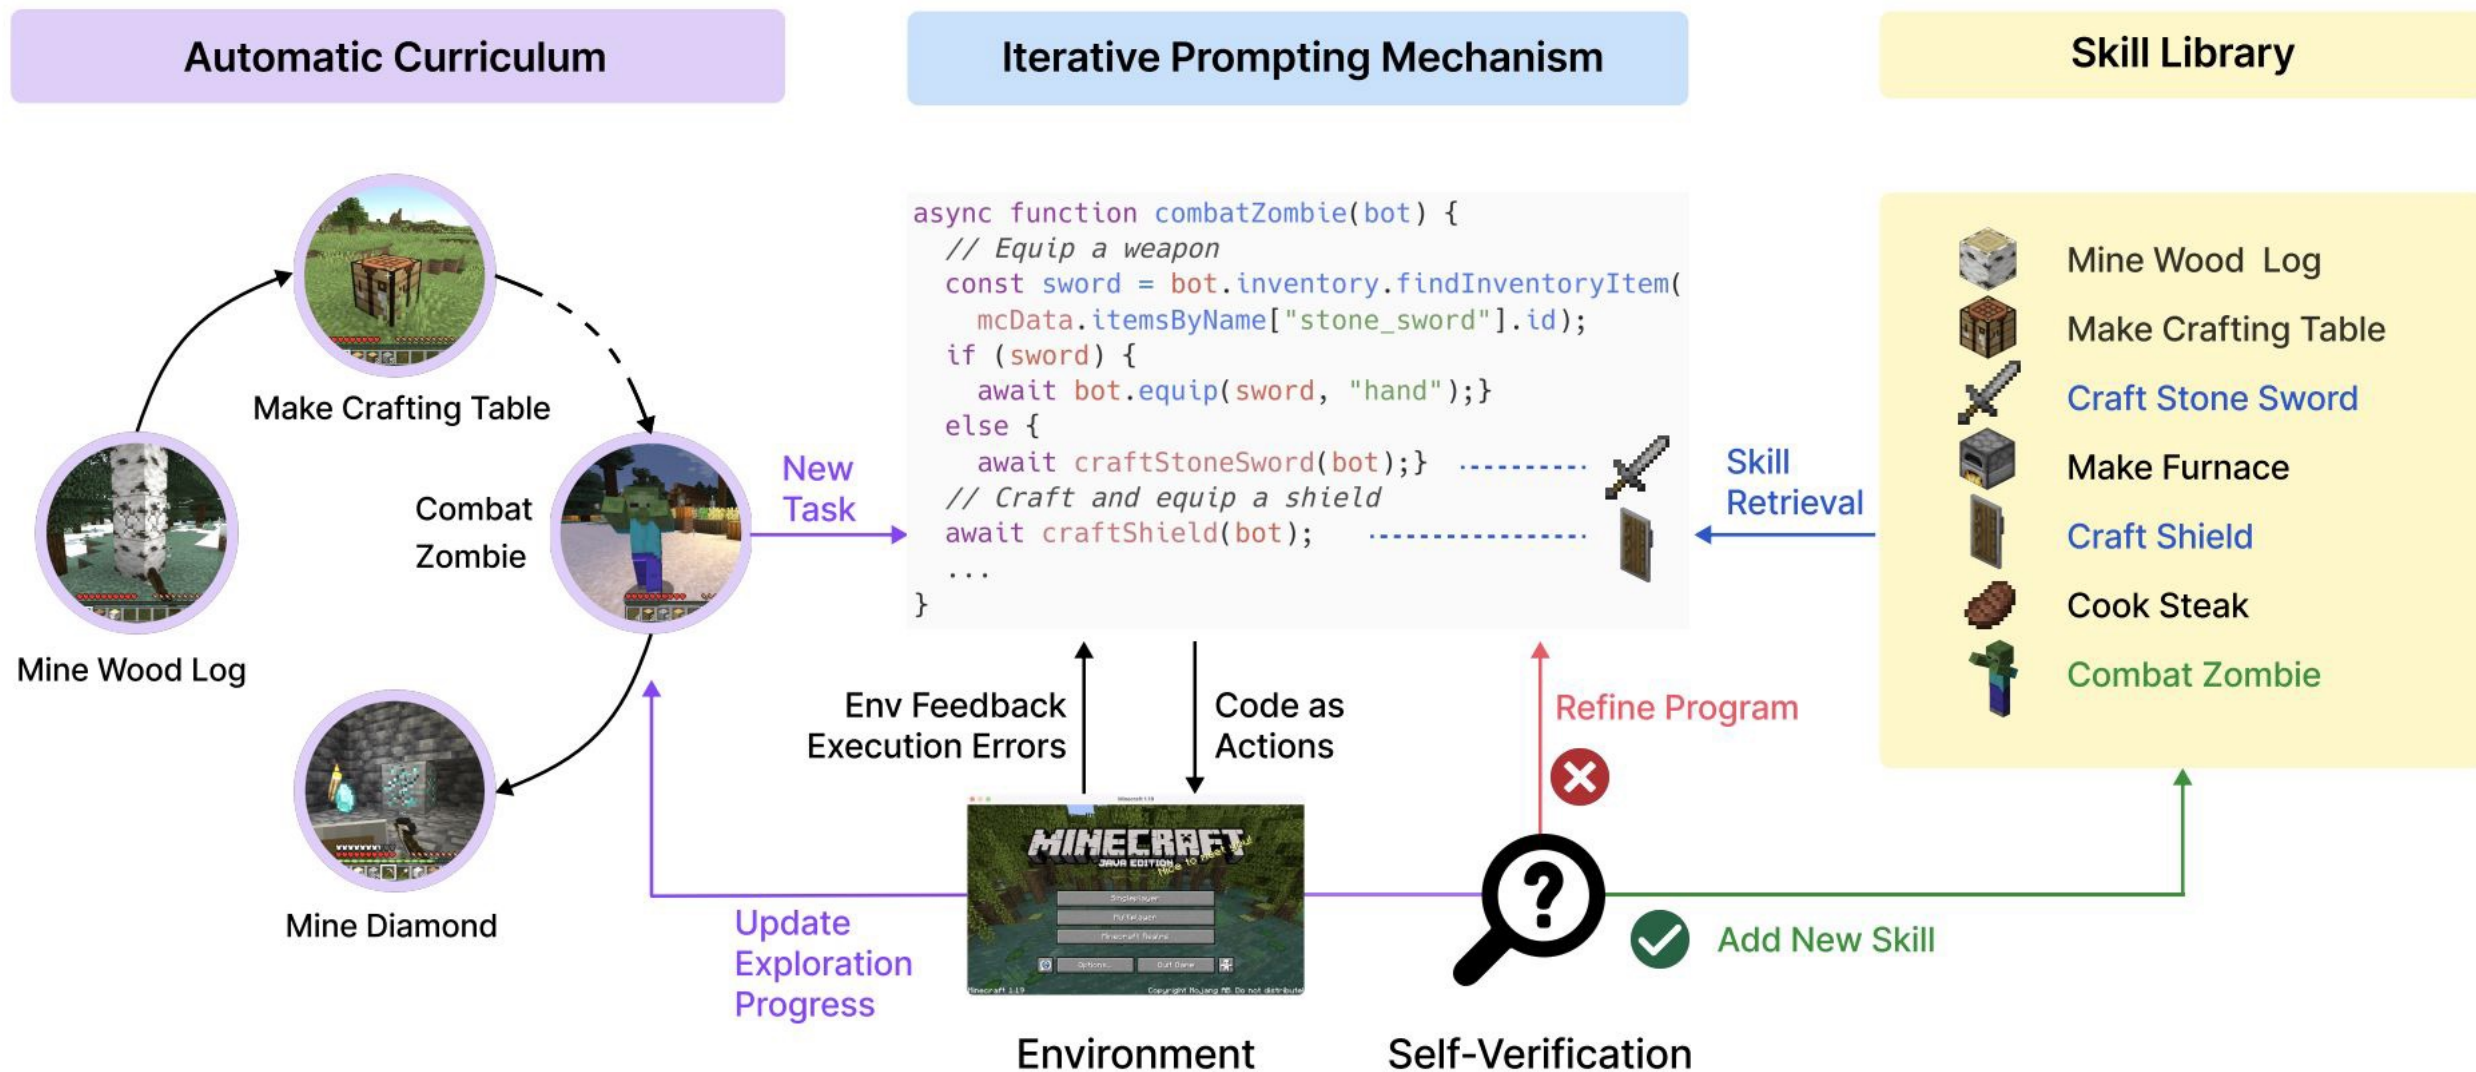
一个Skill

触发条件：什么情况下使用该技能

执行过程：具体怎么操作

预期效果：成功的标志是什么

Skill



Skill

Coding Agent

Skill: 读取CSV并清洗数据

触发条件: 用户上传CSV文件, 要求数据分析

执行过程:

1. 用 pandas 读取文件
2. 检测并处理缺失值
3. 检测并删除重复行
4. 输出清洗后的数据概览

预期效果: 返回干净的 DataFrame + 清洗日志

Skill: 生成数据可视化报告

触发条件: 用户要求"画图"或"可视化"

执行过程:

1. 判断数据类型 (时序/分类/数值)
2. 选择合适图表类型
3. 用 matplotlib 生成图表
4. 自动添加标题、坐标轴标签、图例

预期效果: 返回清晰美观的图表文件

Skill

Research Agent

Skill: 快速定位领域核心论文

触发条件: 用户提出某个研究方向，需要了解背景

执行过程:

1. 用领域关键词搜索 arxiv 和 Google Scholar
2. 按引用数筛选出 top 10 论文
3. 提取每篇论文的：问题、方法、结论、局限性
4. 识别引用关系，找出奠基性工作
5. 按时间线梳理技术演进脉络

预期效果: 输出“领域地图” - 关键论文 + 演进时间线

Skill 生态

全部 自媒体 金融 法律 互联网 科研 教育 健康医疗 生活服务 通用工具 其它

技能包 数据集 开发辅助 办公效率 数据分析 影音娱乐

行业技能包 | 官方严选优质技能组合

查看更多 >



A股投资分析技能包

基于深度量化策略与多维数据雷达，全方位洞察A股板块热度、...

技能包 | 3.2w

一键添加

包含以下技能

板块热度分析 科技股分析 量化股票分析 红利股分析 查看更多



自媒体创作技能包

提供集小红书与公众号图文创作、多平台改写及违禁词查询于一...

技能包 | 3.9w

一键添加

包含以下技能

小红书笔记创作 公众号文案创作 今日头条爆款操盘手 查看更多



全能法务助理技能包

专为法务从业者、律师与企业合规团队打造的智能法务工具箱。...

技能包 | 1.1w

一键添加

包含以下技能

法律咨询报告 诉讼大师 法律类案检索 法律备忘录 查看更多



调研分析技能包

融数据挖掘、商业咨询与市场洞察于一体的超级分析技能包，一...

技能包 | 1.7w

一键添加

包含以下技能

超级分析师 行业趋势深度调研 竞品/公司调研报告 查看更多

Skill 生态

🔍 科研严选技能



arXiv 论文数据搜索

数据集 免费

arxiv 论文搜索技能

🔥 1.2w人在用

添加

💬 训练平台研究员，检索「Frontier Tuning」技术论文。



Semantic Scholar论文搜索

数据集 免费

学术论文知识图谱查询技能，精准检索关联研究内容与脉络。

🔥 1.1w人在用

添加

💬 筛选生物医学RAG幻觉评估研究



OpenAlex 学术数据库

数据集 免费

OpenAlex 学术数据库查询技能

🔥 1.3w人在用

添加

💬 科研助理，筛多模态医学报告高被引研究。



Crossref 文献引用溯源

数据集 免费

学术文献元数据查询技能，助力快速检索文献关键信息

🔥 1w人在用

添加

💬 研究生，查询「大学生一旦过了13周」学业压力论文DOI。



「毕导」的学术雷达

免费

「毕导」为科研人员量身定制，追踪研究领域最新论文，生...

🔥 2.6w人在用

添加

💬 生物信息PI，盯空间转录组分析方法的学术热点与选题建议。



PubMed 生物医学文献搜索

数据集 免费

PubMed 医学文献查询技能，助力医学领域文献检索与研究...

🔥 6.2K人在用

添加

💬 检索GLP-1用于MASH临床文献



bioRxiv 论文搜索

数据集 免费

bioRxiv 论文检索技能

🔥 1.9K人在用

添加

💬 发育生物学博士，查类器官与胚胎轴建立的生物医学预印本。



OpenCitations 开放引文库

数据集 免费

论文引用关系数据查询技能

🔥 9.7K人在用

添加

💬 心理老师，梳理「大学生一旦过了13周」焦虑研究的高被引论文。



文献引用格式转换

免费

论文引用格式智能识别与互转工具，支持 11 种主流学术 / 文...

🔥 6.3K人在用


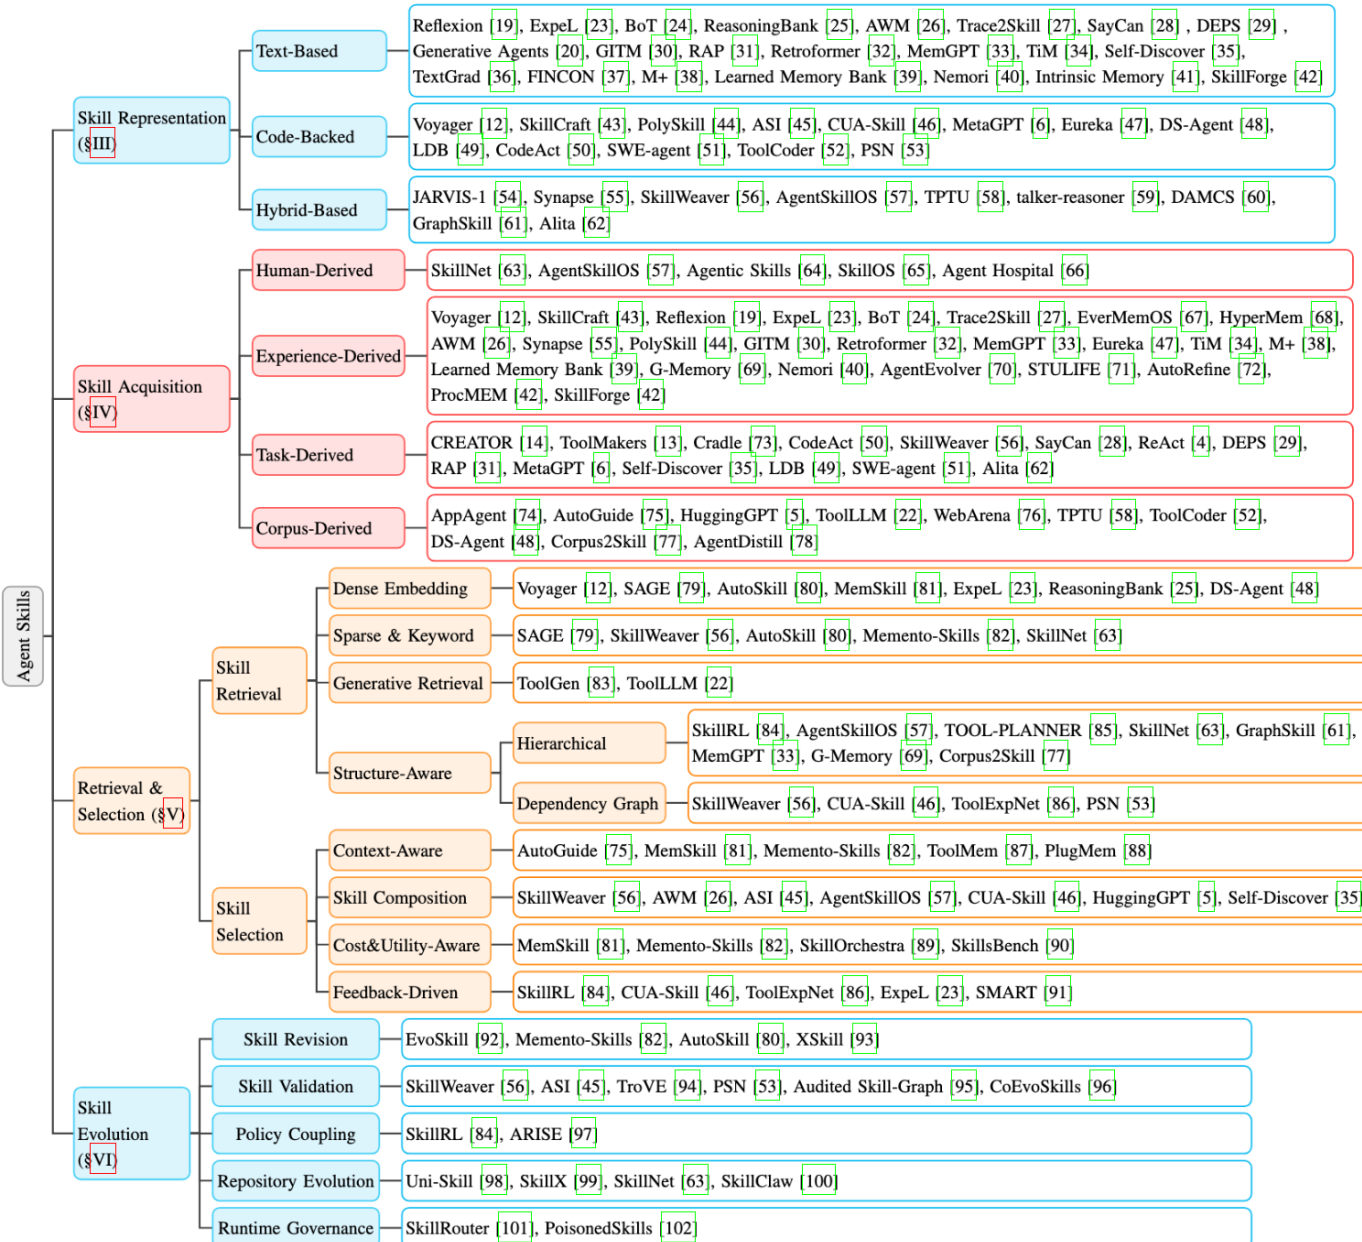
添加

💬 识别「Cochrane」导出记录，改成Vancouver投稿格式。

前沿科学问题

- Skill如何表示：文本？代码？逻辑语言？
- Skill如何生成：专家写？自动获取？
- Skill如何检索：从Skill库中找到和当前任务相关的Skill
- Skill如何演进：自主进化智能体
-

前沿科学问题



A Comprehensive Survey on Agent Skills: Taxonomy, Techniques, and Applications

https://www.alphaxiv.org/abs/2605.07358?utm_source=chatgpt.com

Evolution of AI Agents: LLM-First View



Chatbot

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. ✓

Reasoner




Tool User



Agent

核心能力	自然语言对话	分步思考、思维链	搜索、计算、调用API	规划、执行、观察、反馈、记忆
主要局限	只能回答、不能执行	仍局限于文本空间	需要外部工具编排	长程任务、物理世界？


Evolution of AI Agents: Agent-First View




Symbolic Agent

感知能力	缺乏符号接地能力		
推理能力	严谨符号逻辑推理		
适应能力	依赖预先定义的符号空间		

Evolution of AI Agents: Agent-First View




Symbolic Agent




Neural Agent

感知能力	缺乏符号接地能力	神经网络的感知能力	
推理能力	严谨符号逻辑推理	缺乏推理能力	
适应能力	依赖预先定义的符号空间	样本效率低，难以在不同场景迁移	


Evolution of AI Agents: Agent-First View



Symbolic Agent



Neural Agent



LLM Agent


感知能力	缺乏符号接地能力	神经网络的感知能力	高，几乎任意模态任意数据
推理能力	严谨符号逻辑推理	缺乏推理能力	基于语言的思维链推理
适应能力	依赖预先定义的符号空间	样本效率低，难以在不同场景迁移	大模型自身的通用性

Applications

- Digital World
 - Coding Agents
 - Mobile Agents
 - Web/App Agents
 - Computer Agents
- Physical World
 - Gaming Agents
 - Robotics Agents

Coding Agents

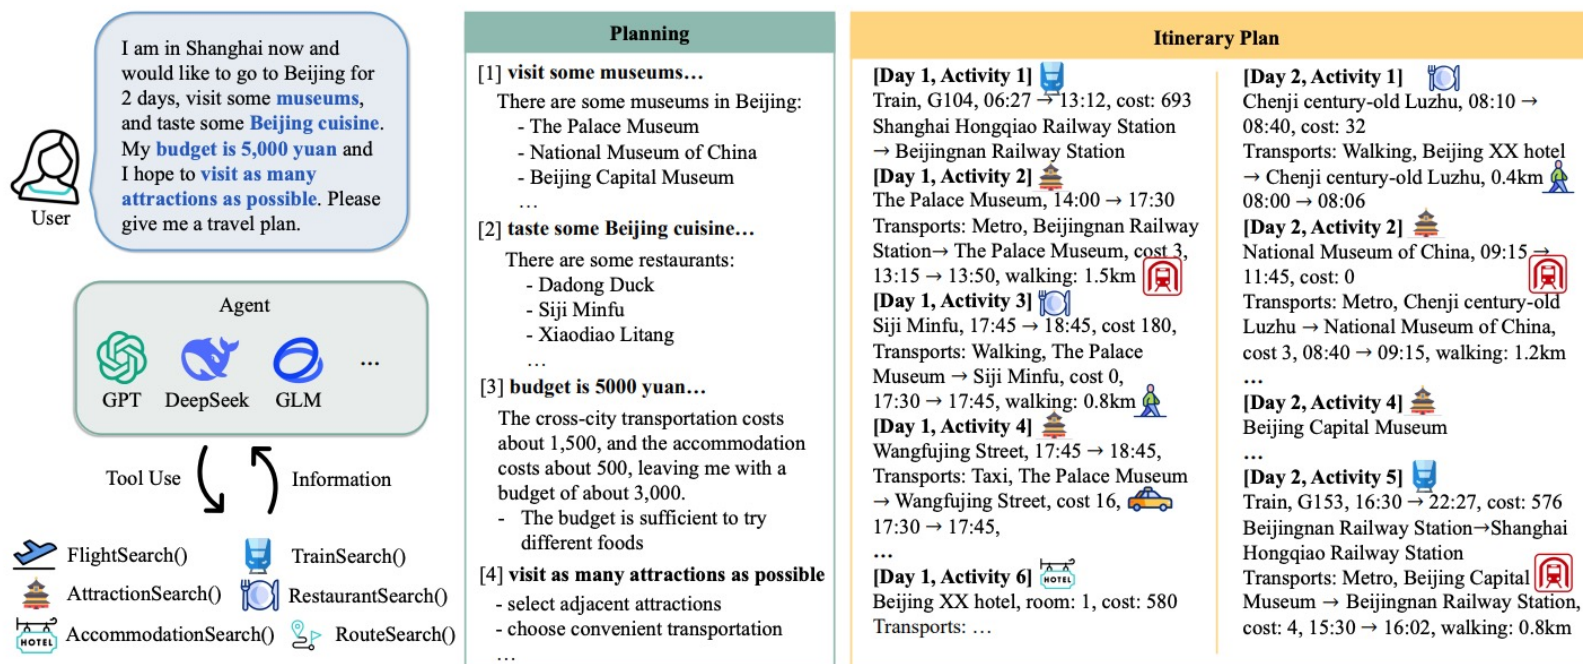
- 环境(Environments): Project code repo, file systems, IDEs,
- 观察(Observation): Code files, outputs, docs, errors, commit history,
- 动作(Action Space): Coding edits, File search/view, test updates,



Web/Tool-Use Agents

Web/App Use

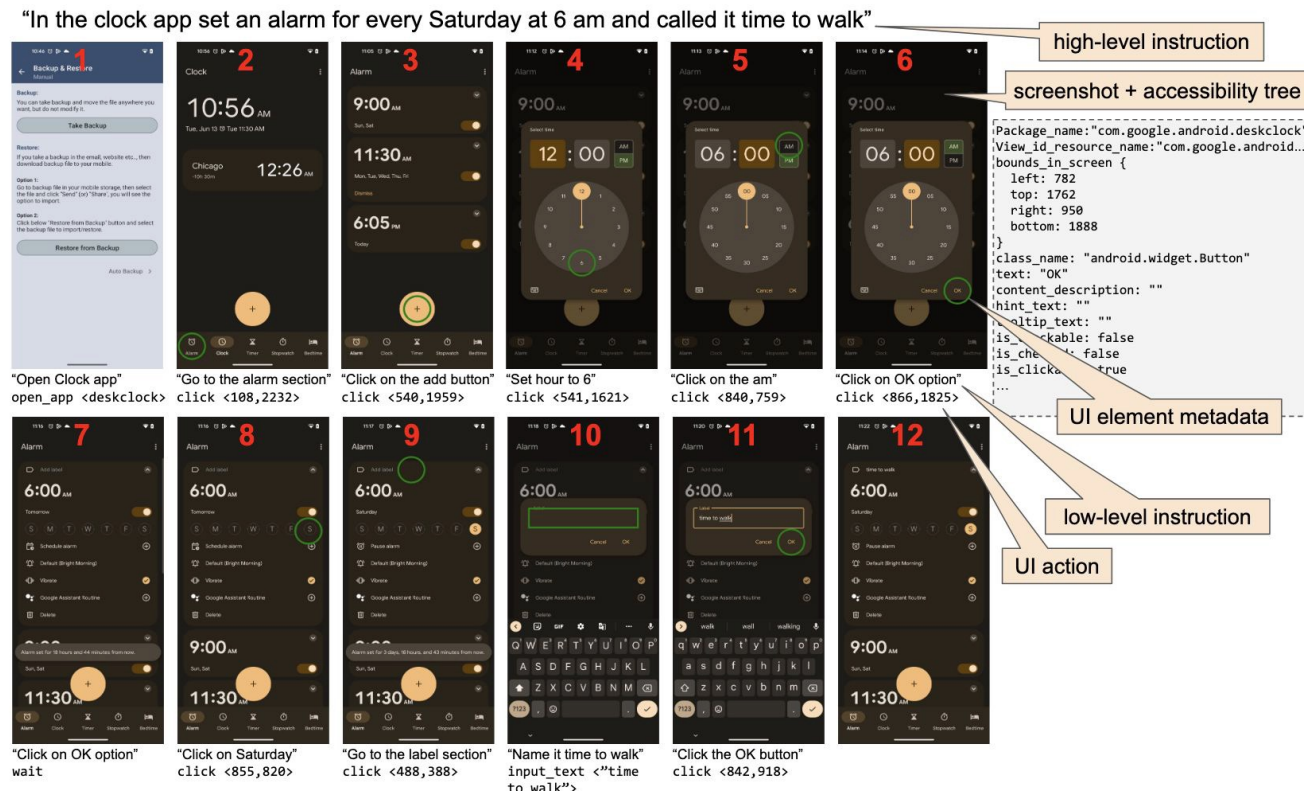
- Environment: web browsers/apps
- Observation space: screenshots, DOM trees, HTML, historical actions...
- Action space: browser/app controls (e.g., click, type, scroll, drag, hover...)



Mobile Use Agents

Mobile Use

- Environment: mobile device systems
- Observation space: screenshots, a11y trees, HTML, historical actions...
- Action space: mobile controls (e.g., tap, type, swipe...)



Education Agent

教学智能体



AI写作助手

@梁春龙

通过鲁迅的写作风格，进行文章的编写

1 0



计划书评分智能体

@黄裕强

计划书评分智能体

147 0



计划书点评智能体

@黄裕强

创意组、创业组、公益组计划书点评智能体

64 0




涉海口语陪练

@官方

六大海事场景，练就涉海英语口语功底

54 2



文书生成助手

@官方

能够准确识别用户需求，完成复杂的文书生成工作。

590 3



操作系统原理课程智能体

@官方

面向《操作系统原理》课程，能够帮助学生通过知识点问答、学习路径规划，实现自主课程学习。

28 0



计算机组成原理与系统结...

@官方

面向《计算机组成原理与系统结构》课程，能够帮助学生通过知识点问答、学习路径规划，实现自主课...

27 0



程序设计课程智能体

@官方

面向《C语言程序设计》课程，能够提供知识问答、代码分析、作业解析等能力。

334 4



海法通译

@官方

智能海法翻译助手，精准传递法理要义

145 9



儿童绘画点评

@官方

上传你独一无二的绘画作品，为你解读画作的创意与想象力，助力发现大师潜力。

17K 1



创意探究课题生成

@官方

基于多领域知识，智能生成新颖课题，帮助教师生成创意探究案例。

398 11




学习搭子

@官方

一位温暖耐心的学习搭子，用鼓励和陪伴点亮你的学习之路。

2K 19



Research Agents

<https://github.com/wanshuiyin/Auto-claude-code-research-in-sleep>

```
>
(base) yangruofeng@yangruoengdeAir ~ % aris
```



A	R	I	S
Auto Research in Sleep			
adversarial multi-agent			
Claude x GPT-5.4 x high			



```
Executor      Anthropic · claude-opus-4-6
Reviewer      gpt-5.4
Permissions    danger-full-access
Directory      /Users/yangruofeng
Session        session-1[REDACTED]
```

Type **/help** for commands · `/model` or `/reviewer` to switch

```
> 你好你是谁你可以干什么
```


你好！我是 **Claude Opus 4.6** (model ID: ``claude-opus-4-6``)，由 Anthropic 开发。我运行在 **ARIS** (**Auto Research in Sleep**) 研究自动化系统中。

我能做什么？ 🧐 科研全流程自动化

Financial Agents

TradingAgents: Multi-Agents LLM Financial Trading Framework

<https://github.com/tauricresearch/tradingagents>



Generative Agents


Generative Agents: Interactive Simulacra of Human Behavior

<https://arxiv.org/pdf/2304.03442>



Generative Agents

<https://www.bilibili.com/video/BV1xamQBuEgS?t=0.7>



Generative Agents

<https://github.com/666ghj/MiroFish>



The screenshot displays the MiroFish web application. The main area features a dense, colorful network graph with nodes and edges. A large, stylized text overlay in the center reads "预测万物，群体智能" (Predict everything, collective intelligence) in blue and "开源项目MiroFish" (Open-source project MiroFish) in red. Below this, another large text overlay reads "预测红楼梦未完结局" (Predict the ending of the Red Chamber Dream) in yellow. On the right, a "Node Details" panel for a character named "宝玉" (Baoyu) is visible, showing fields for Name, UUID, Created, Properties, and Summary. The interface also includes a top navigation bar with "图谱" (Graph), "双栏" (Dual Panel), and "工作台" (Workspace) tabs, and a bottom status bar indicating "Step 1/5 图谱构建" (Step 1/5 Graph Construction) and "Ready".

预测万物，群体智能

开源项目MiroFish

预测红楼梦未完结局

相当于是时序GraphRAG发挥的作用

Game Agents

Digital Games

- Environment: game worlds/levels...
- Observation space: screenshots of game states, inventory, location...
- Action space: game controls (e.g., drop, move, attack, resource management...)



Robotics Agents

Robotics for physical interaction

- Environment: physical world spaces
- Observation space: visual input, sensor readings, physical states, proprioception...
- Action space: motor controls (e.g., move, grasp, manipulate...)

