



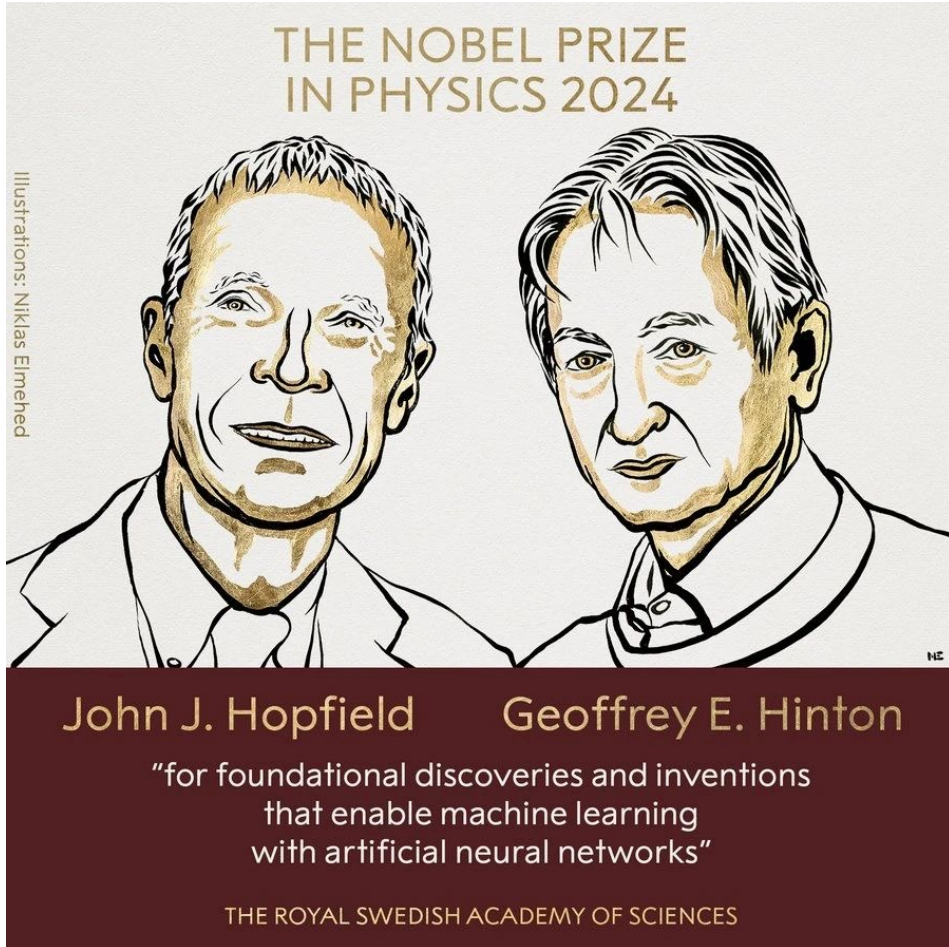
The Brains Behind the AI Turbocharging LLMs for Scientific Discovery!

Xuanjing Huang
Fudan University



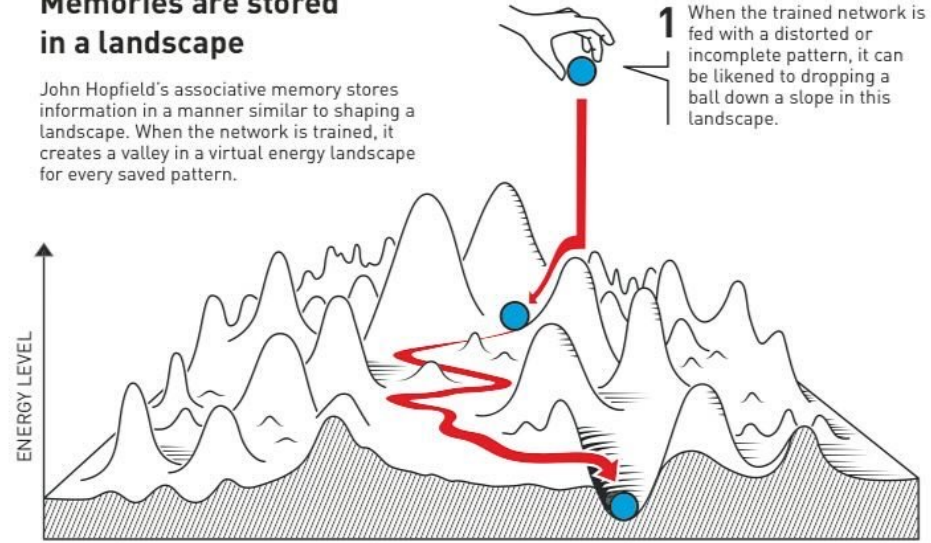
Scientific Discovery in the age of AI

AI & Nobel Physics Prize



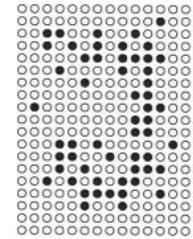
Memories are stored in a landscape

John Hopfield's associative memory stores information in a manner similar to shaping a landscape. When the network is trained, it creates a valley in a virtual energy landscape for every saved pattern.

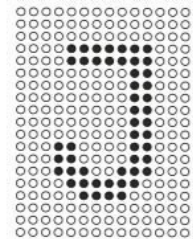


©Johan Jarnestad/The Royal Swedish Academy of Sciences

INPUT PATTERN



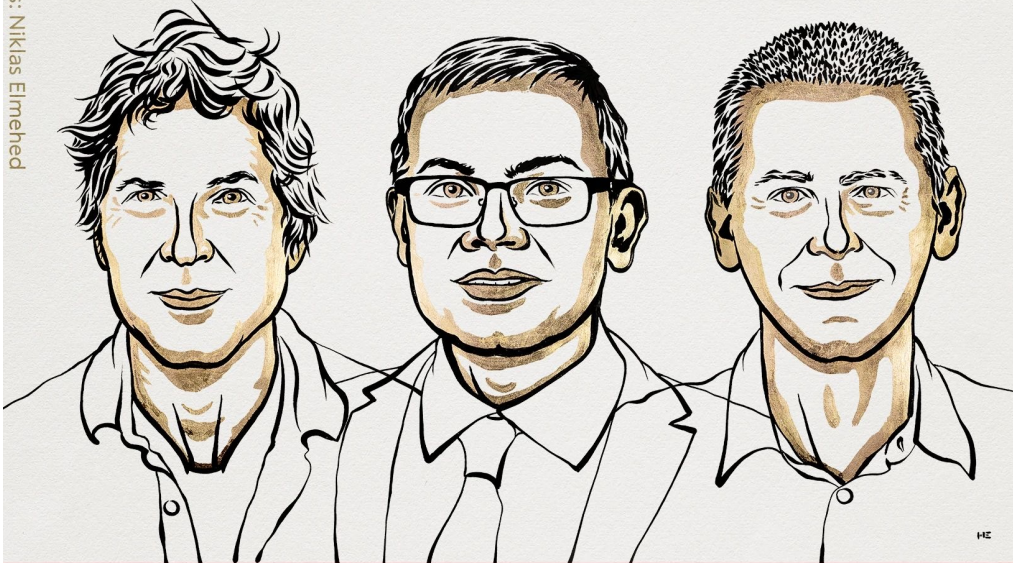
SAVED PATTERN



AI & Nobel Chemistry Prize

Illustrations: Niklas Elmehed

THE NOBEL PRIZE
IN CHEMISTRY 2024



David Baker
"for computational protein design"

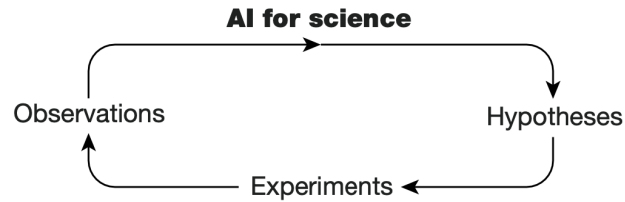
Demis Hassabis
"for protein structure prediction"

John M. Jumper

THE ROYAL SWEDISH ACADEMY OF SCIENCES



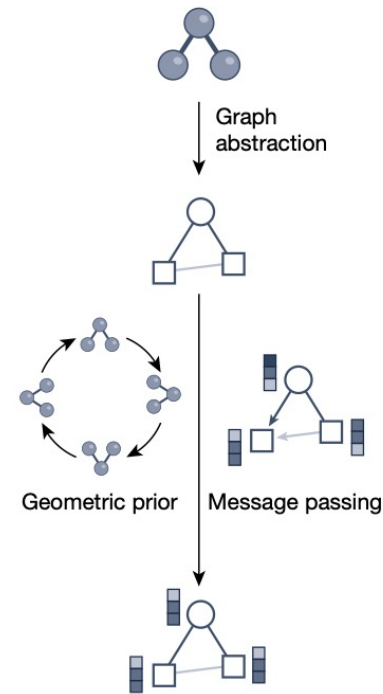
Scientific Discovery in the age of AI - Data Analysis



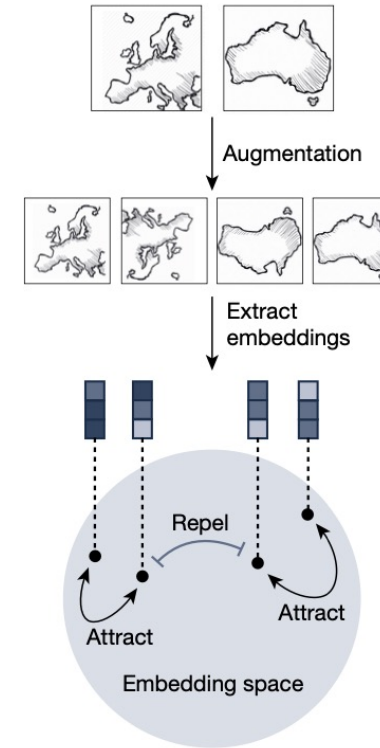
- | | |
|---|---|
| Weather forecasting | Rare event selection in particle collisions |
| Battery design optimization | Language modelling for biomedical sequences |
| Magnetic control of nuclear fusion reactors | High-throughput virtual screening |
| Planning chemical synthesis pathway | Navigation in the hypothesis space |
| Neural solvers of differential equations | Super-resolution 3D live-cell imaging |
| Hydropower station location planning | Symbolic regression |
| Synthetic electronic health record generation | |

Science discovery involves many steps: making **hypotheses**, setting up **experiments**, and collecting and analyzing **data**

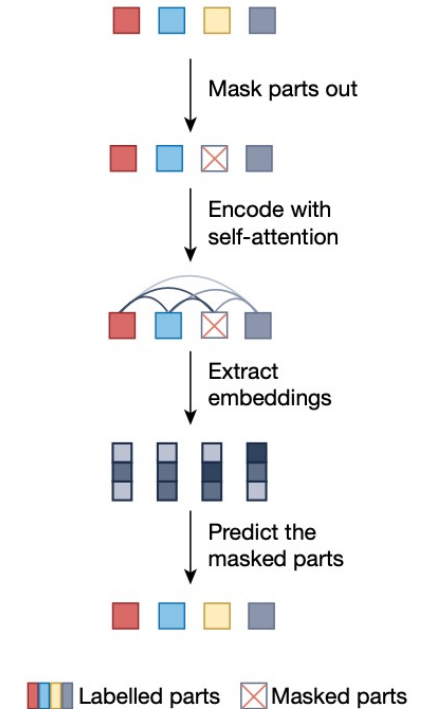
a Geometric deep learning



b Self-supervised learning

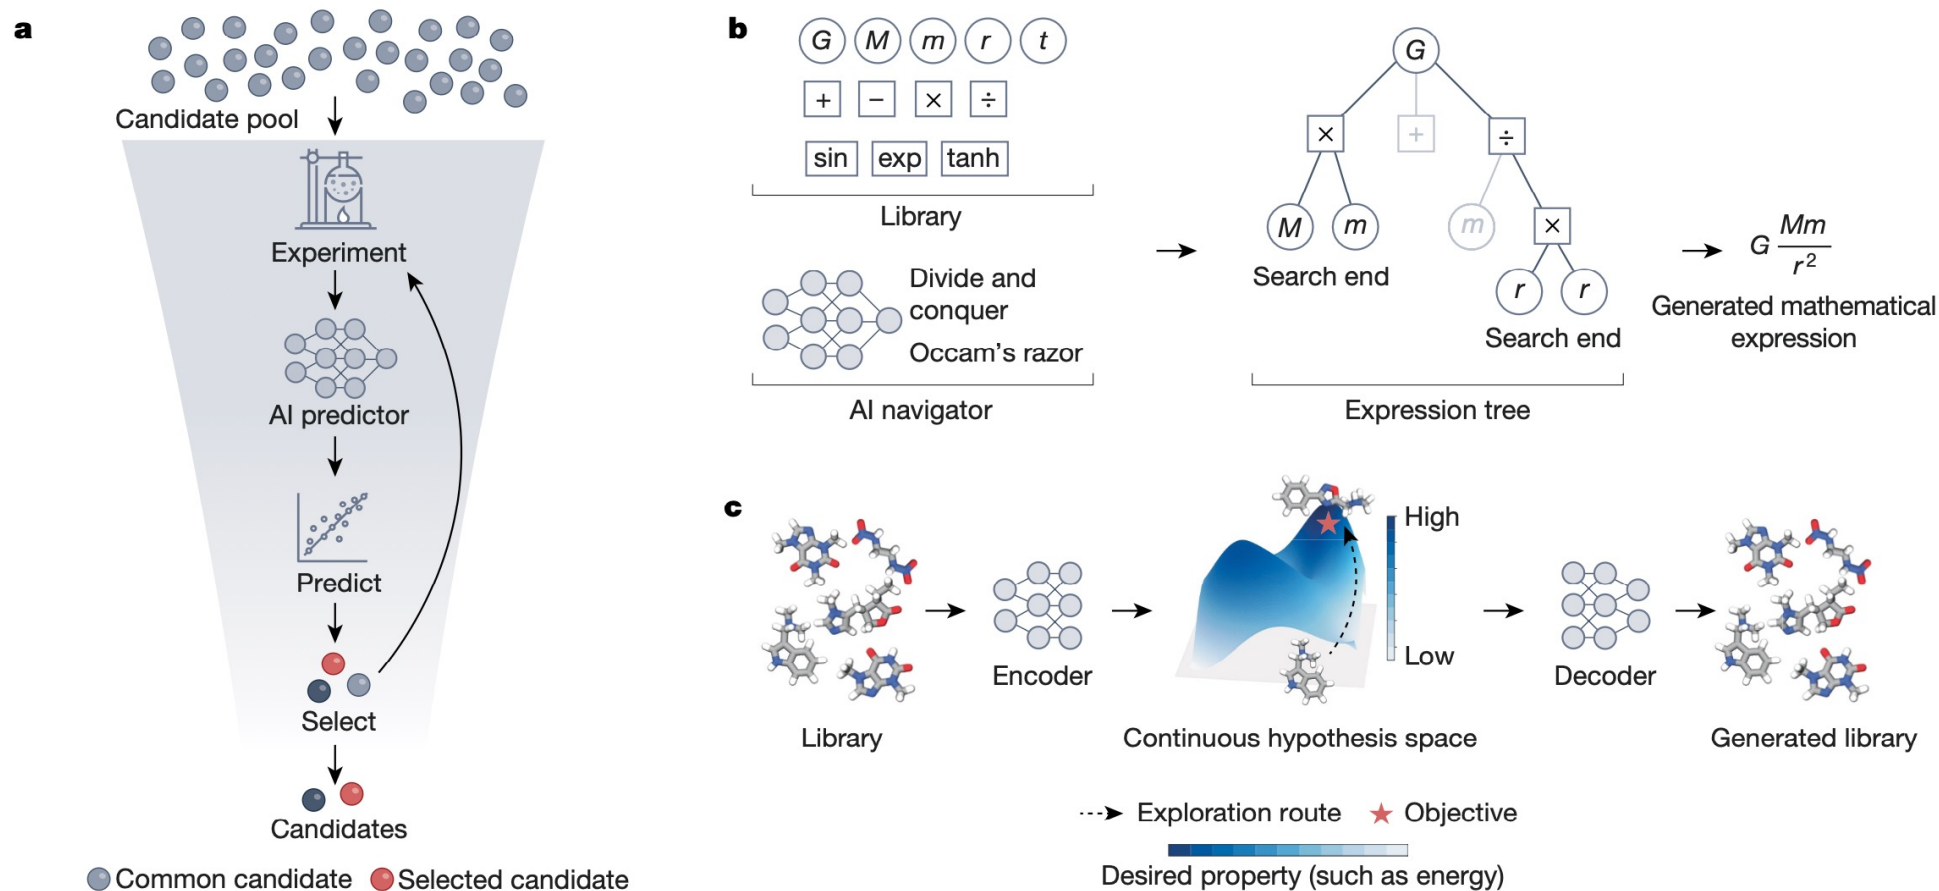


c Masked-language modelling



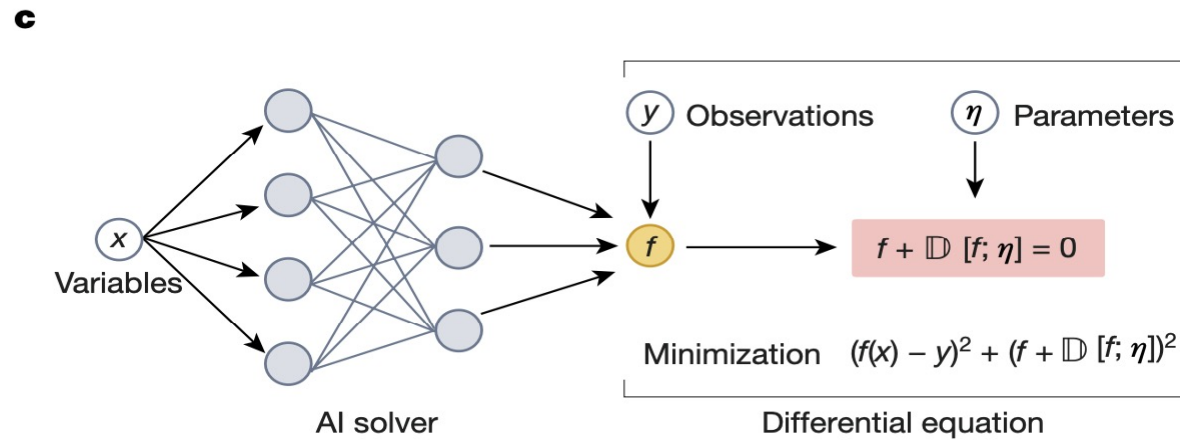
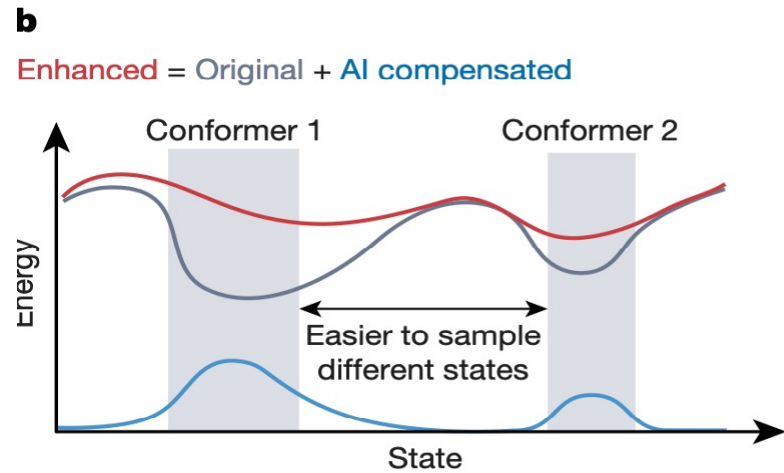
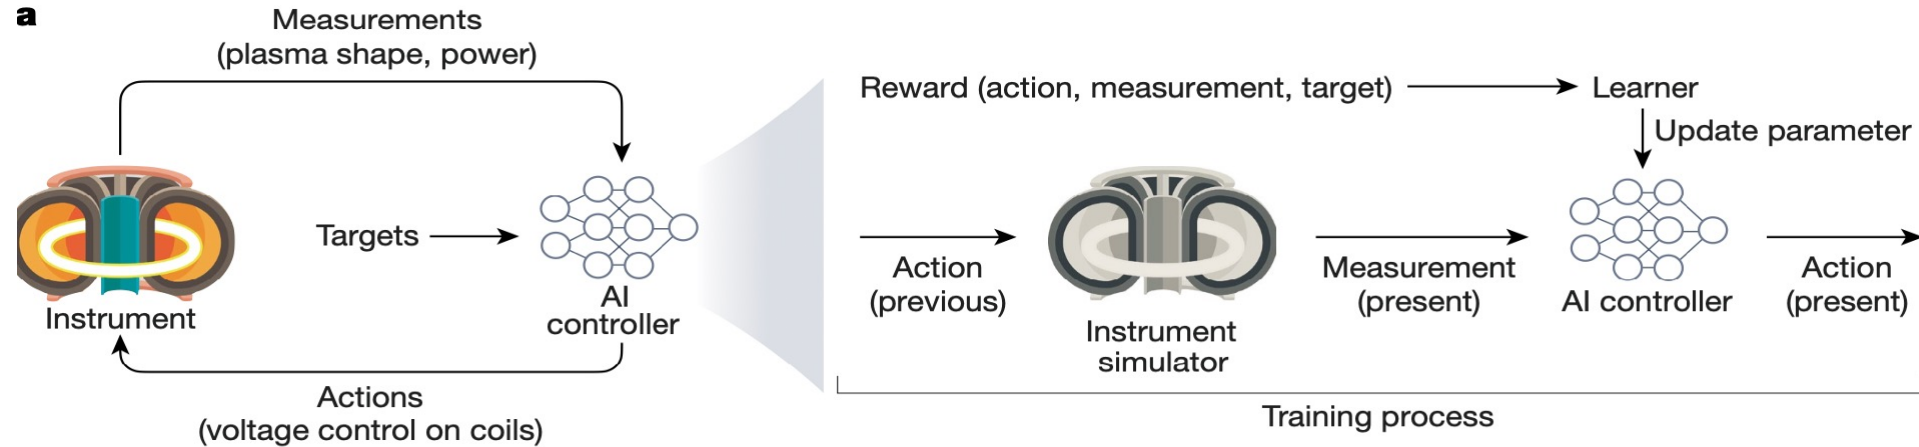
Techniques like **Geometric Deep Learning**, **Self-Supervised Learning**, **Masked Language Modeling** are already used in areas like **weather forecasting**, **battery design**, and **controlling nuclear energy**

Scientific Discovery in the age of AI: Hypothesis Generation



Hypothesis generation of scientific discovery is the starting point for exploring the unknown
 Tools like **AI Predictors** , **AI Navigator** and **AI Differentiator** help us move from a massive pool of possibilities to precisely the target we need

Scientific discovery in the age of AI: Experiments

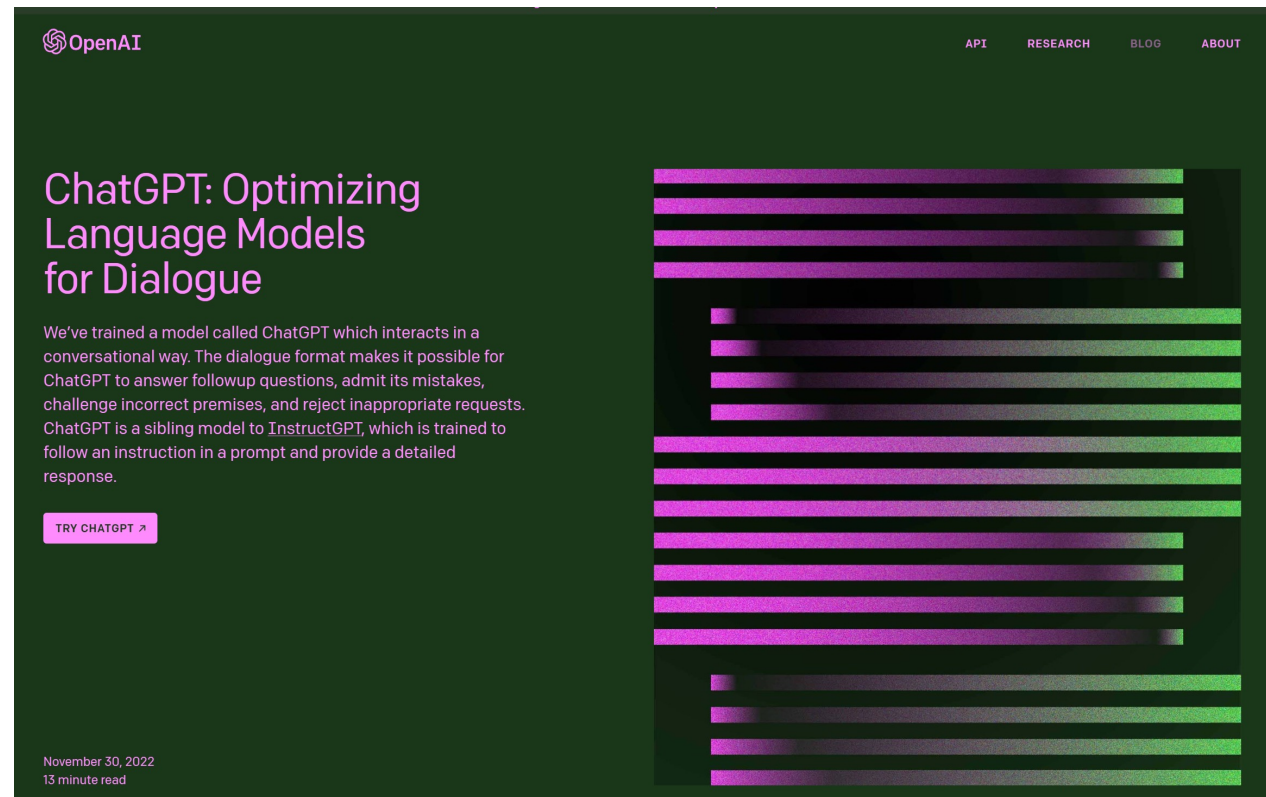
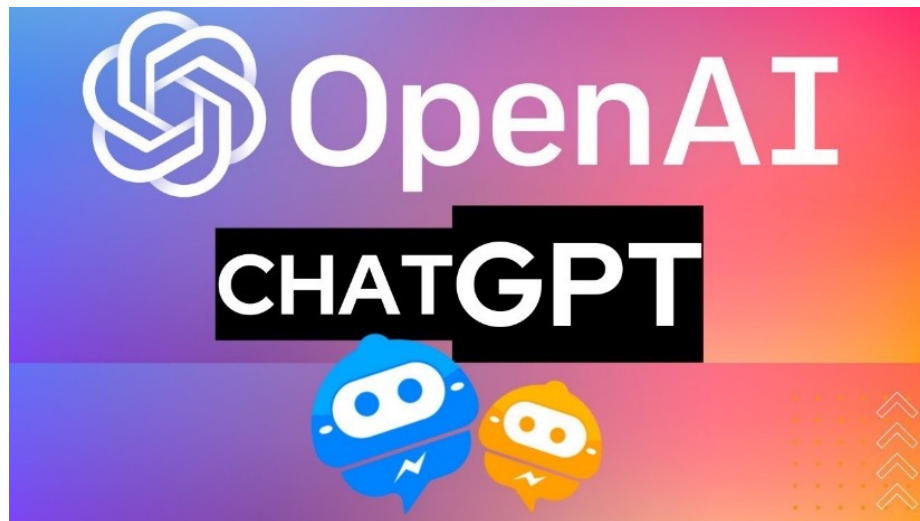


Use AI to manage highly complex, dynamic systems like **controlling nuclear fusion reactors**

Detect Rare Events: In massive computer simulations of complex systems

Use neural networks to **solve partial differential equations**

The Conversational AI That Changed Everything



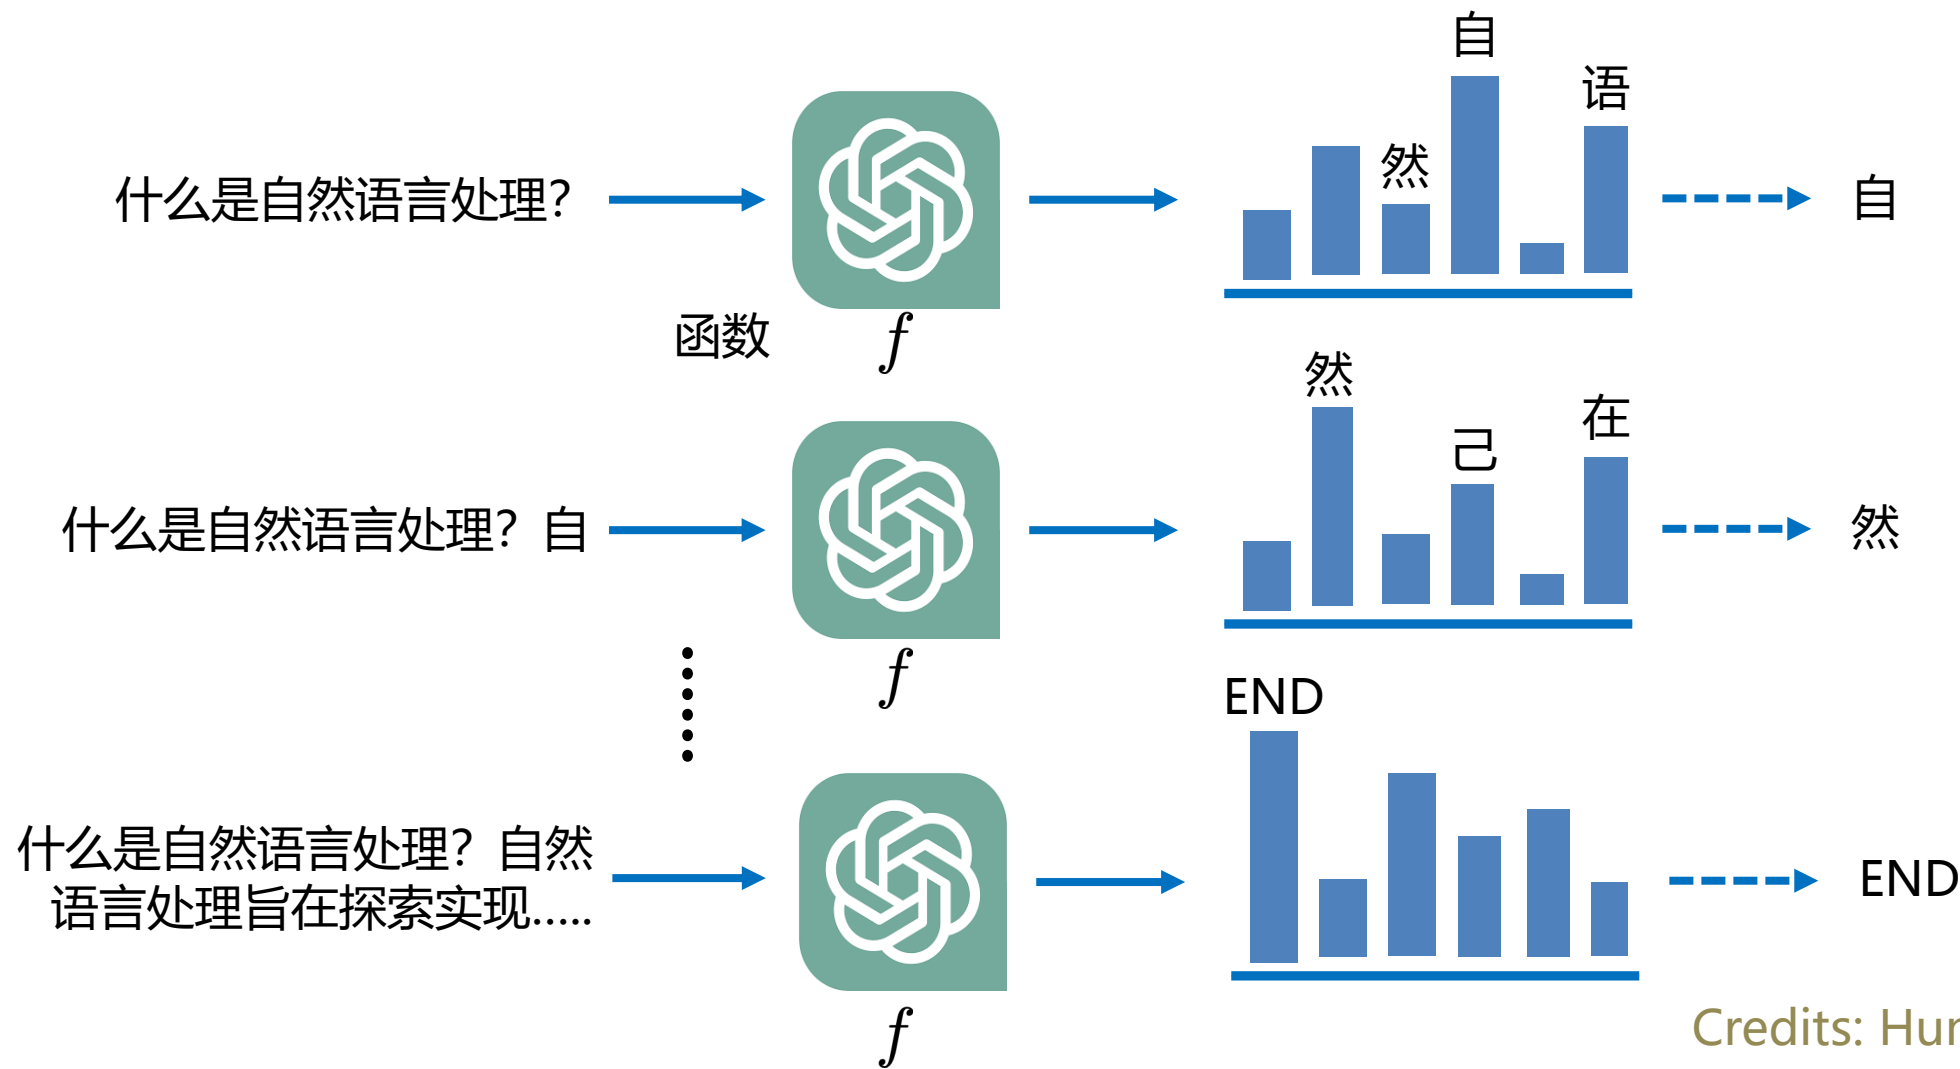
'AI chatbots like ChatGPT will become just as important as the personal computer or the internet!'

Large Language Models

- ▶ **Large Language Models** (LLMs): Machine Learning models with massive parameters and complex computational structures
- ▶ **Emergent Abilities**: Complex, unpredictable capabilities that suddenly appear when model scale (data and parameters) crosses a critical threshold
- ▶ Key Features:
 - ▶ **Massive** parameters and **deep** layers
 - ▶ Offering strong expressive power and high accuracy
 - ▶ Requiring significant **computing resources** for training and inference



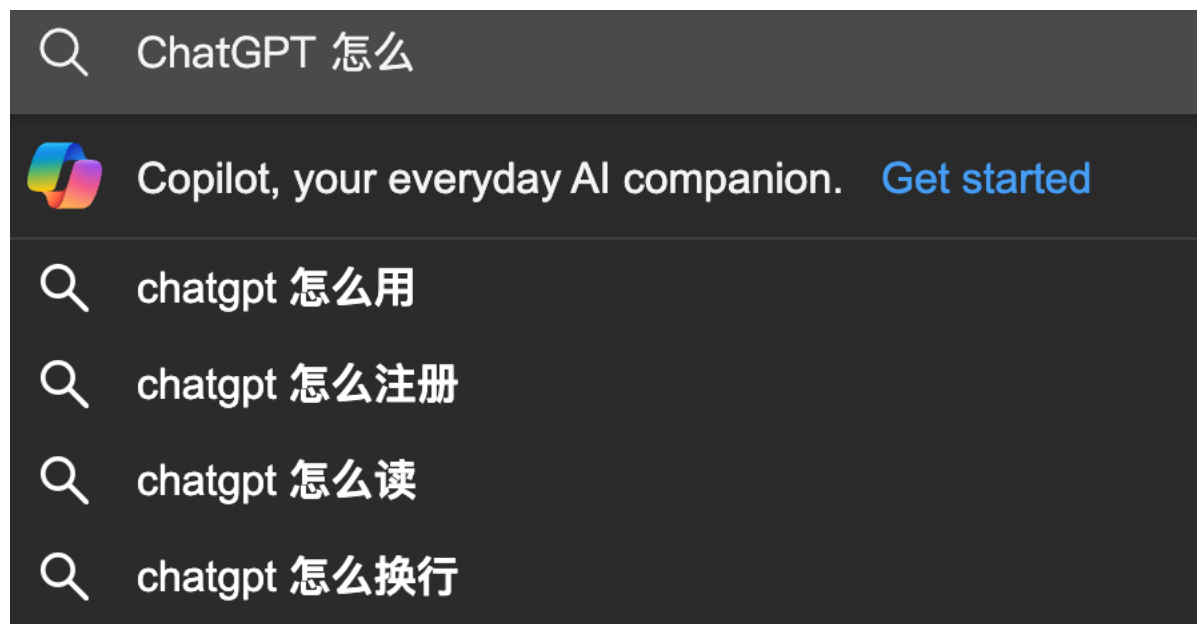
Next Word Prediction



Credits: Hung-Yi Lee

Next Word Prediction

We actually experience this 'next word prediction' task every single day!



Benefit of Next Word Prediction

communicate just like a human!

“把下面这句话翻译成英语：这很好”

“That is good.”

“The capital of China is”

LLM

“Beijing”

Expand a story starting with 'Once there was a student’

“Once there was a student studying for an exam”



LLMs: Your Go-To Assistant for Work & Life

Reading Comprehension & Summarization

Hi, 我是 Kimi~
很高兴遇见你! 你可以随时把网址🔗或者文件📄发给我, 我来帮你看看

<https://arxiv.org/pdf/2402.14963>
请介绍一下这篇文章的内容

这篇文章的标题是《Mirror: A Multiple-perspective Self-Reflection Method for Knowledge-rich Reasoning》, 作者是Hanqi Yan, Qinglin Zhu, Xinyu Wang, Lin Gui, 和 Yulan He, 他们分别来自King's College London, University of Warwick, 和 The Alan Turing Institute. 文章提出了一种名为Mirror的新型自我反思方法, 旨在增强大型语言模型 (LLMs) 在知识丰富问题上的理解与推理能力。

摘要:
文章首先指出, 尽管大型语言模型 (LLMs) 能够迭代地反思自己的输出, 但它们在外部资源的情况下解决知识密集型问题时存在困难。此外, LLMs在自我评估和根据明确负面反馈重新审视预测方面也存在不足。为了解决这些问题, 作者提出了Mirror方法, 这是一种多角度自我反思的方法, 通过一个导航器 (Navigator) 和推理器 (Reasoner) 之间的启发式交互, 引导代理走向多样化且可能可靠的推理路径, 而无须依赖于真实情况。Mirror通过鼓励 (1) 导航器生成的方向多样性, 以及 (2) 推理器生成的响应中策略性引入的扰动之间的一致性来实现这一点。在五个推理数据集上的实验表明, Mirror比其他几种当代自我反思方法具有优越性。此外, 消融研究清楚地表明, 作者的策略缓解了前述挑战。相关代码已在GitHub上发布。

引言:
介绍了大型语言模型 (LLMs) 作为各种任务的重要且灵活的构建块, 并通过迭代修正在许多任务中得到改进, 例如代码生成、算术问题解决和推理。在迭代细化过程中, 评估当前响应并生成有价值反馈的批评模块对提高性能至关重要。

相关工作:
文章回顾了关于通过自我反思概念增强LLMs的研究, 其中LLMs从自动生成的反馈中学习, 以理解和反思自己的输出。这些反馈可能来自LLM本身、单独训练的批评模块或外部来源, 例如Wikipedia或互联网浏览器。

挑战:
作者通过实验来解决在没有真实标签的情况下增强LLMs自我改进能力的挑战, 特别是在知识丰富的推理任务中。

Mirror框架:
详细介绍了Mirror的无监督自我反思框架, 包括问题设置、多角度评估、多样化和有效的搜索空间。

实验:
评估了Mirror在MMLU (一个知识丰富的问答数据集) 和FEVER (一个事实检查数据集) 上的性能, 并与几种基于CoT (Chain-of-Thought) 和自我一致性推理方法进行了比较。

结论:
文章总结了Mirror在知识丰富的推理中的优势, 并讨论了LLMs在事实评估和生成高质量反馈方面的局限性。

限制:
讨论了研究的主要焦点是识别基于生成输出和冻结状态的最佳推理路径, 以及如何充分利用LLMs在复杂推理中的潜力。

伦理声明:
文章提到了使用的两个公开可用数据集: MMLU和FEVER, 并强调了在应用相关技术于现实场景前需要仔细考虑安全含义。

文章最后感谢了英国工程和物理科学研究委员会 (EPSRC) 对这项工作的部分支持。

详细介绍一下Mirror的框架

Mirror框架是一种为大型语言模型 (LLMs) 设计的多角度自我反思方法, 旨在提高模型在知识丰富问题上的推理能力。该框架通过模拟人类辅导过程, 使用一个导航器 (Navigator) 和一个推理器 (Reasoner) 之间的启发式交互, 帮助模型从多个角度进行自我反思。以下是Mirror框架的详细介绍:

- 问题设置:**
 - Mirror将推理过程视为一个马尔可夫决策过程 (MDP), 由状态 (State)、行动 (Action)、转移概率 (Transition Probability)、策略 (Policy)、折扣因子 (Discount Factor) 和奖励函数 (Reward Function) 组成。
 - 在每次反思迭代中, 状态 (st) 是由推理器生成的响应, 包括问题的答案和推理过程; 行动 (at) 是由导航器生成的方向。
- 多角度评估:**
 - 导航器为K个问题生成方向, 这些方向旨在提供多样化的问题解决视角, 导航器和推理器之间的一致性代表了不同视角之间的一致性。
 - 内部一致性 (Intra-consistency) 作为评估标准, 如果超过阈值T0, 则认为推理器应该接受其初始响应; 否则, 将交互一致性 (Inter-consistency) 作为后续迭代中停止标准的指标。

Hi, 我是 Kimi~
很高兴遇见你! 你可以随时把网址🔗或者文件📄发给我, 我来帮你看看

<https://arxiv.org/pdf/2311.09677>
我要在我论文的Related Work部分中引用这篇论文, 请用一句话总结一下它的贡献。

这篇论文提出了一种名为R-Tuning的新方法, 旨在教授大型语言模型识别和拒绝回答超出其知识范围的问题, 通过识别模型参数知识与指令调整数据之间的差距, 并构建基于知识交集的拒绝意识数据集, 从而提高模型在回答已知问题时的准确性, 并避免生成不存在的事实 (hallucination)。

复制 重试一次 分享

Translation

把下面这段话翻译成中文： The concept of artificial intelligence has evolved significantly over the past few decades. Initially, AI was seen as a futuristic idea, something that belonged to the realm of science fiction. However, with the rapid advancements in computing power and data availability, AI has become an integral part of our daily lives. From voice assistants like Siri and Alexa to recommendation systems on Netflix and Amazon, AI is now deeply embedded in the technologies we use every day.



中文翻译：

“人工智能的概念在过去几十年中发生了显著的变化。最初，人工智能被视为一种未来主义的概念，属于科幻小说的范畴。然而，随着计算能力和数据可用性的快速发展，人工智能已经成为我们日常生活中不可或缺的一部分。从像 Siri 和 Alexa 这样的语音助手，到 Netflix 和 Amazon 上的推荐系统，人工智能如今已深深嵌入我们每天使用的技术中。”

这段翻译力求准确、流畅，同时保留了原文的语气和逻辑结构。如果你有其他文本需要翻译，欢迎随时告诉我！

翻译：入乡随俗



When in Rome, do as the Romans do.

翻译：fish in the troubled water



浑水摸鱼。

Text Polishing

你是一个英文学术论文写作专家，以下是一篇学术论文中的一段内容，请先对其进行翻译为英文，并将此部分润色以满足学术标准，提高语法、清晰度和整体可读性。大语言模型已经在各种NLP任务中到达了前所未有的高度，成为实现AGI的不可或缺的途径之一。LLM在大量数据上进行预训练，不仅表现出强大的语言能，而且在多个领域接近人类水平的熟练程度。然而，监督微调通过训练输入输出数据与自然语言中的任务指令相结合，基础的大语言模型在基于以任务为导向的监督数据进行微调后，得到了更好地泛化以及可靠性，并能更好地遵循指令。

提示词

Translated Text:

Large Language Models (LLMs) have reached unprecedented heights in various Natural Language Processing (NLP) tasks, becoming an indispensable pathway to achieving Artificial General Intelligence (AGI). LLMs are pre-trained on vast amounts of data, demonstrating not only powerful linguistic capabilities but also proficiency nearing human levels across multiple domains. However, through supervised fine-tuning, which combines training on input-output data with task instructions in natural language, the foundational LLMs, after fine-tuning on task-oriented supervised data, have achieved better generalization and reliability, and are more capable of following instructions.

Polished Academic Text:

Large Language Models (LLMs) have achieved remarkable success across a spectrum of Natural Language Processing (NLP) tasks, positioning themselves as a critical avenue towards the realization of Artificial General Intelligence (AGI). These models, pre-trained on extensive datasets, have exhibited robust linguistic competencies and have approached human-level proficiency in a multitude of domains. Nevertheless, the process of supervised fine-tuning, which integrates the training of input-output pairs with task directives expressed in natural language, has further enhanced the generalization and reliability of these foundational LLMs. Consequently, they are now more adept at adhering to instructions, a significant advancement in their utility and performance.

介绍梯度下降的通俗易懂的视频，并附上链接



🔍 已搜索到 47 个网页 >

以下是几个通俗易懂的视频资源，帮助你理解梯度下降的基本原理和应用场景：

1. 3Blue1Brown 的《梯度下降法》

- 视频链接: [Gradient Descent, How Neural Networks Learn | Chapter 2, Deep Learning](#)
- 特点: 通过生动的动画和“下山”的比喻，直观地解释了梯度下降的工作原理。适合完全没有数学背景的观众。
- 推荐理由: 3Blue1Brown 的视频以数学可视化著称，即使是文科生也能轻松理解复杂的数学概念 1 8。

2. StatQuest 的《梯度下降法》

- 视频链接: [Gradient Descent, Step-by-Step](#)
- 特点: 用简单的语言和图表，逐步解释梯度下降的每一步操作，并结合实际例子展示其在机器学习中的应用。
- 推荐理由: StatQuest 的视频以清晰和简洁著称，适合初学者快速入门 4 9。

3. 吴恩达的《机器学习课程：梯度下降》



Data Visualization

MV 将上述图形修改为折线图。

以下是一个将上述复杂带方差柱状图修改为带方差的折线图的示例：

```
python Copy code

import numpy as np
import matplotlib.pyplot as plt

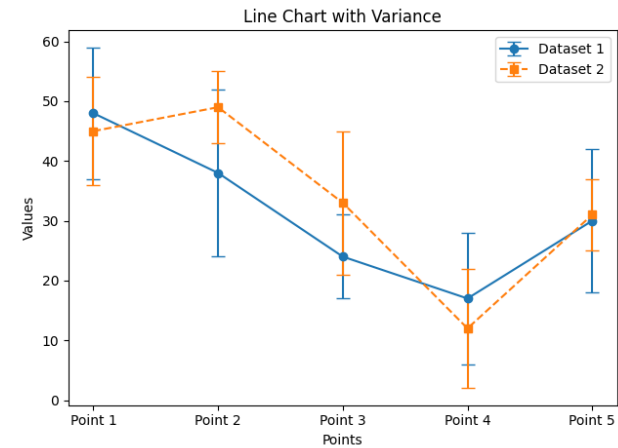
# 随机生成数据和方差
np.random.seed(42)
n_points = 5
data_1 = np.random.randint(10, 50, size=n_points)
std_1 = np.random.randint(5, 15, size=n_points)

data_2 = np.random.randint(10, 50, size=n_points)
std_2 = np.random.randint(5, 15, size=n_points)

# 设置点的标签
labels = [f'Point {i+1}' for i in range(n_points)]

# 设置点的位置
x = np.arange(len(labels))

# 绘制带方差的折线图
fig, ax = plt.subplots()
```



Coding

我有一组时间戳，格式为“Jan 14, 2009 at 5:13 AM”，请问我该如何使用python脚本对这组时间戳进行排序？

要使用Python对这种格式的时间戳进行排序，你可以使用`datetime`模块来解析时间字符串并将其转换为`datetime`对象。之后，你可以轻松地对这些日期对象进行排序。下面是一个简单的脚本示例，展示了如何进行这一操作：

1. 导入必要的模块。
2. 定义时间戳列表。
3. 使用`datetime.strptime`函数将字符串转换为`datetime`对象。
4. 对转换后的日期对象进行排序。
5. 打印排序后的结果。

这里是具体的代码实现：

```
python
from datetime import datetime

# 假设这是你的时间戳列表
timestamps = [
    "Jan 14, 2009 at 5:13 AM",
    "Jan 12, 2009 at 5:13 AM",
    "Jan 13, 2009 at 5:13 AM"
]

# 定义一个函数，用于将时间字符串转换为datetime对象
def parse_timestamp(ts):
    return datetime.strptime(ts, "%b %d, %Y at %I:%M %p")

# 将每个时间字符串转换为datetime对象
dates = [parse_timestamp(ts) for ts in timestamps]

# 对日期进行排序
dates.sort()

# 打印排序后的日期
for date in dates:
    print(date.strftime("%b %d, %Y at %I:%M %p"))
```

这个脚本中，`strptime`函数用于解析日期字符串，`strftime`函数则用于将`datetime`对象格式化回原来的字符串格式，以便打印和查看。这种方式让你能够轻松地处理和排序日期数据。



请帮我把下面代码改成多进程形式

```
from GPTWrapper import GPTWrapper
from sentence_transformers import
SentenceTransformer
import re
import os
import argparse
```

```
input_text = '''Please list 100 diverse time-sensitive
questions that often require the latest knowledge to
answer.
Current date is: 2024-04-15
Please do not use any placeholders!
Please use a variety of sentence structures to pose
questions.
Ensure the diversity of the questions.
Please make sure that the answers to the questions may
change, requiring the retrieval of the latest knowledge
for accurate responses.
Try to ensure that the time sensitivity of the questions is
not obvious.
```

You should follow the following format:

```
<SOS>
Question 1: xxx
Question 2: xxx
...
Question 100: xxx
<EOS>
'''
```



Here's the updated version of your script using Python's `multiprocessing` module to utilize multiple CPU cores. In this version, the embedding generation and similarity checking operations are distributed among processes to increase the performance of these computationally intensive tasks.

```
python
from GPTWrapper import GPTWrapper
from sentence_transformers import SentenceTransformer
from sklearn.metrics.pairwise import cosine_similarity
import re
import os
import argparse
from multiprocessing import Pool, cpu_count, Manager

def get_config():
    parser = argparse.ArgumentParser()
    parser.add_argument('--engine', type=str, default='gpt-3.5-turb')
    parser.add_argument('--embedding_model', type=str, default='fnl')
    parser.add_argument('--generate_number', type=int, default=1000)
    args = parser.parse_args()
    return args

def parse_response(response):
    questions = re.findall(r'Question \d+: (.+?)\n', response)
    return [q.strip() for q in questions]

def add_new_questions(args):
    new_questions, embedding_model, total_questions_proxy = args
    embeddings = embedding_model.encode(new_questions)
    added_questions = list(zip(new_questions, embeddings))
    added_questions = []
    with total_questions_proxy.get_lock(): # Locking to ensure thr
        total_questions = total_questions_proxy.get_obj()
        for question, embedding in new_questions:
            if len(total_questions) == 0:
                total_questions.append((question, embedding))
            else:
```

WisPaper: Searching for Papers



WisPaper

🔍 Scholar Search

⬇️ Download



Search & Screen

You can screen 1000 papers in 5 minutes to get 20 papers you must read

使用大模型的儿童支气管哮喘的生物靶向治疗



Try Our Examples

Could you provide me some examples of works about the use of boosting methods in post-processing? →

List all papers that use autoregressive transformer to generate videos. →

Which papers propose methods to accelerate the generation process in diffusion models? →

NEX: Animation

I need an [animation simulating the motion of a small ball attached to a spring](#). The specific requirements are as follows:

Physical Model: One end of the ball is connected to a spring, with the other end of the spring fixed to a wall.

Adjustable Parameters: The spring stiffness coefficient k and the maximum allowable elongation limit can be adjusted.

Initial Conditions: The simulation starts with the spring stretched n meters from its equilibrium position and then released.

Animation Functionality: The animation should be able to demonstrate the motion under different values of k and different maximum elongation limits by modifying the parameters.

简谐运动模拟器

弹簧振子的周期运动与能量守恒

实时状态
已复位

$F = -100.0 \text{ N}$

0.0 0.6 1.2 1.8 2.4 3.0
位移 (m)

位移-时间图 速度-时间图

加速度-时间图 能量-时间图

弹簧刚度 K 50.0
振子质量 M 2.0 KG
初始位移 X_0 2.0 M
最大伸长限度 3.0 M
重力加速度 G 9.8 M/S²
阻尼系数 Γ 0.00
时间倍率 1.0x

▶ 播放 || 暂停 ↺ 复位

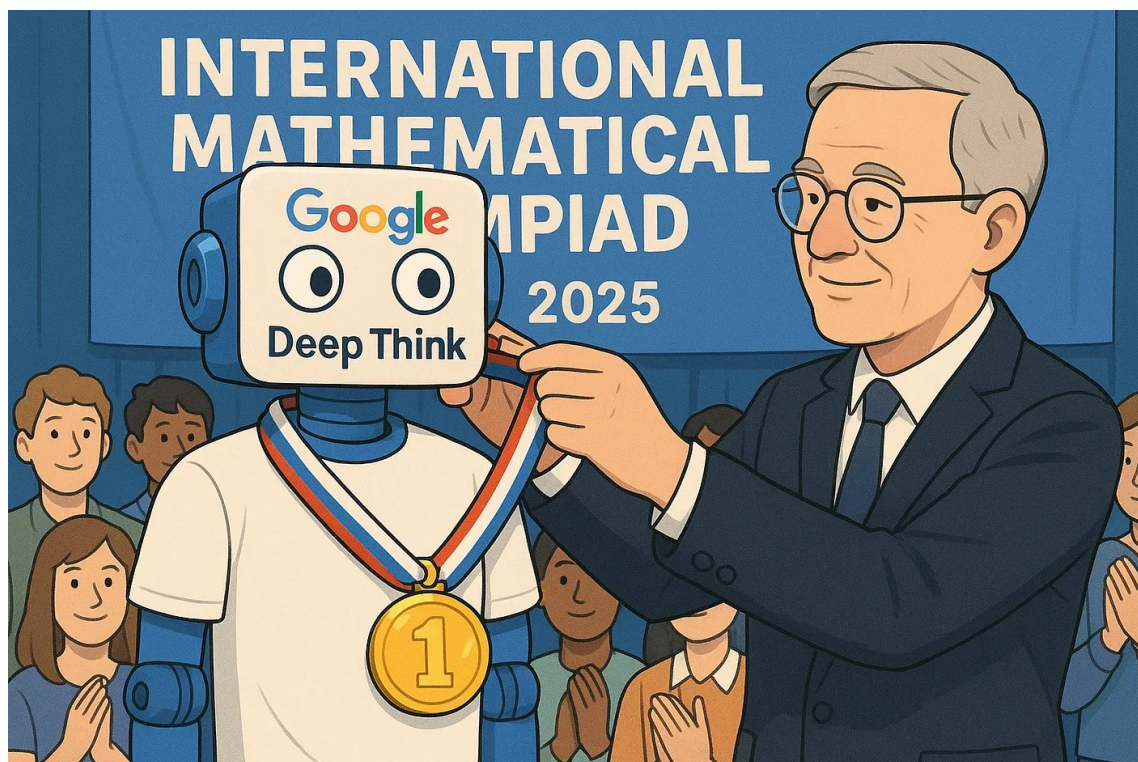
$\omega = \sqrt{k/m} = 5.00 \text{ rad/s}$
 $T = 2\pi/\omega = 1.26 \text{ s}$

位移 X 2.000 m 速度 V 0.000 m/s 加速度 A -50.000 m/s² 动能 E_K 0.000 J 势能 E_P 100.000 J 机械能 E 100.000 J



Evolution Paths of LLMs' Reasoning Capabilities

Google Deep Think @IMO



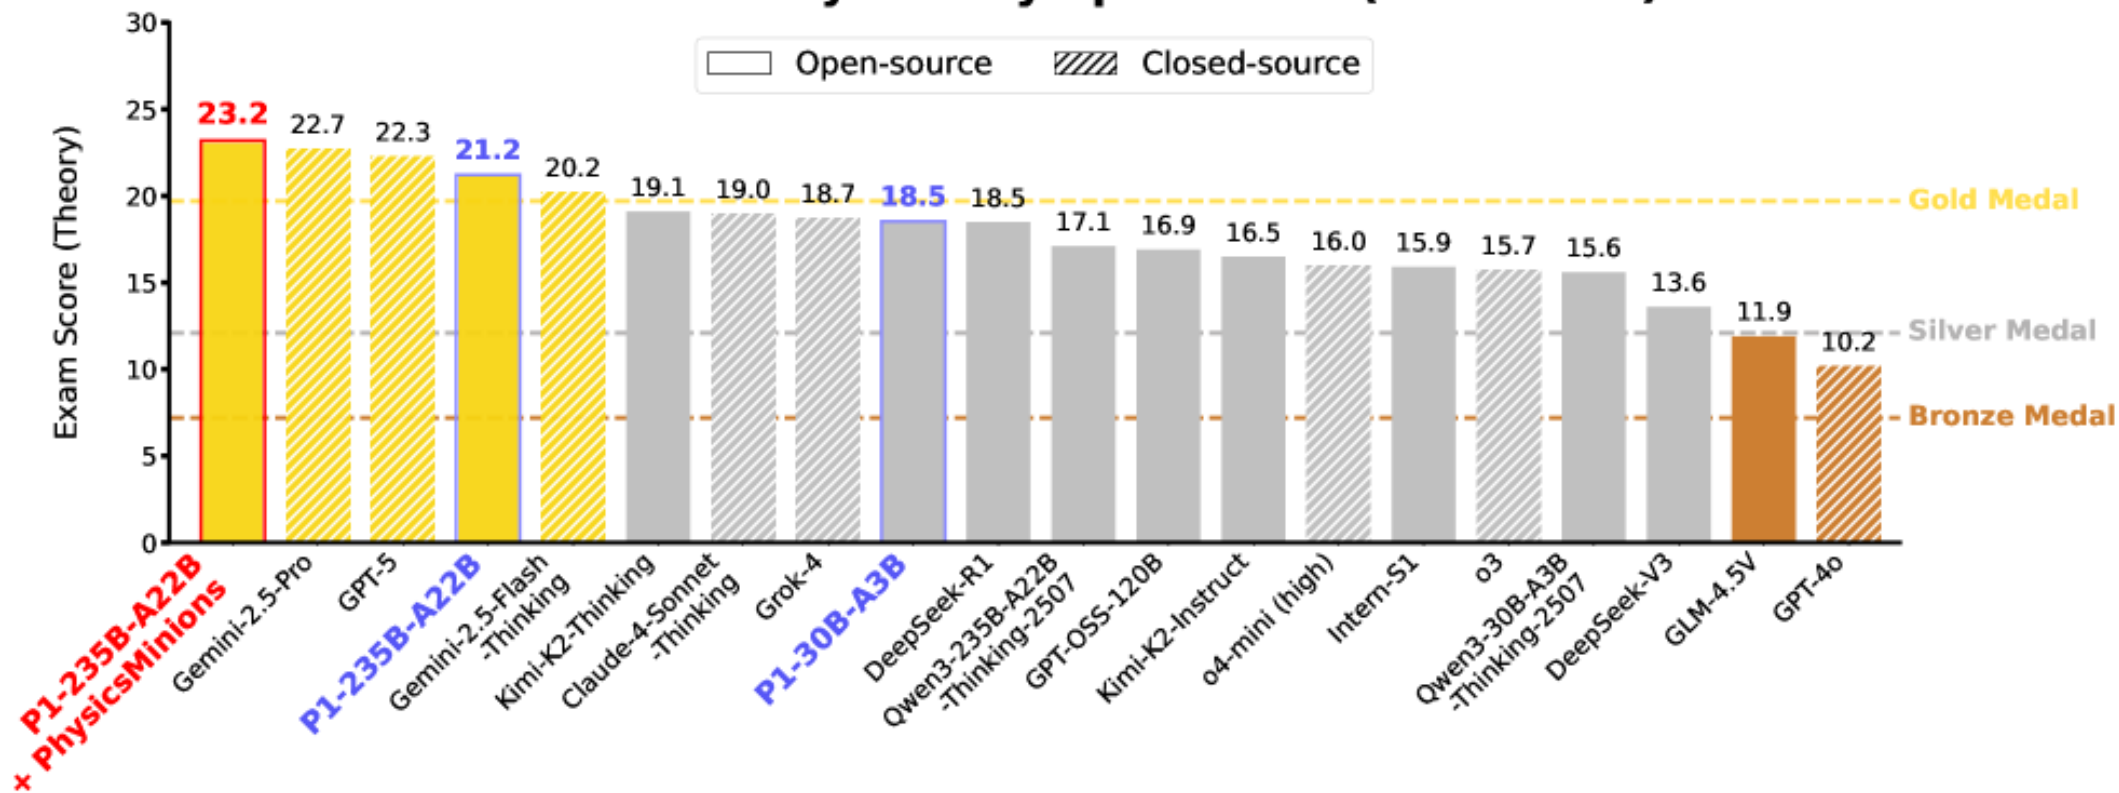
Metric	IMO 2024 System (AlphaProof/AlphaGeometry 2)	IMO 2025 System (Gemini Deep Think)
Medal Standard	Silver ⁶	Gold ¹
Score	28 / 42 ⁶	35 / 42 ⁷
Problems Solved	4 / 6 ¹⁰	5 / 6 ⁴
Methodology	Neuro-symbolic hybrid systems ⁶	General-purpose Large Language Model ⁷
Language Interface	Formal Languages (e.g., Lean) ⁶	End-to-End Natural Language ⁴
Human Intervention	Required for problem translation ⁴	None ("fully in natural language") ¹⁰
Computation Time	2-3 days ⁶	Within 4.5-hour contest limit ¹

The International Mathematical Olympiad (**IMO**) is a global challenge for the world's most talented young minds

IMO was considered the '**Grand Challenge**' for AI research

Intern-P1 @lphO

International Physics Olympiad 2025 (IPhO 2025)



Multiple AI models (both open-source and closed-source) have achieved **gold-medal-level** performance, proving AI's capability to match top-level physics competition standards..

National College Entrance Examination (Gaokao)

7家大模型挑战高考数学新一卷

模型	客观题 (文本) 得分	图像题 (单选) 得分	解答题得分	总分
Gemini 2.5 pro	68	0	77	145
Doubao-1.5-thinking-vision-pro	68	0	76	144
DeepSeek R1	68	/	76	144
o3	65	0	75	140
Qwen3-235b	68	/	71	139
hunyuan-t1-latest	68	/	68	136
文心X1 Turbo	68	/	66	134

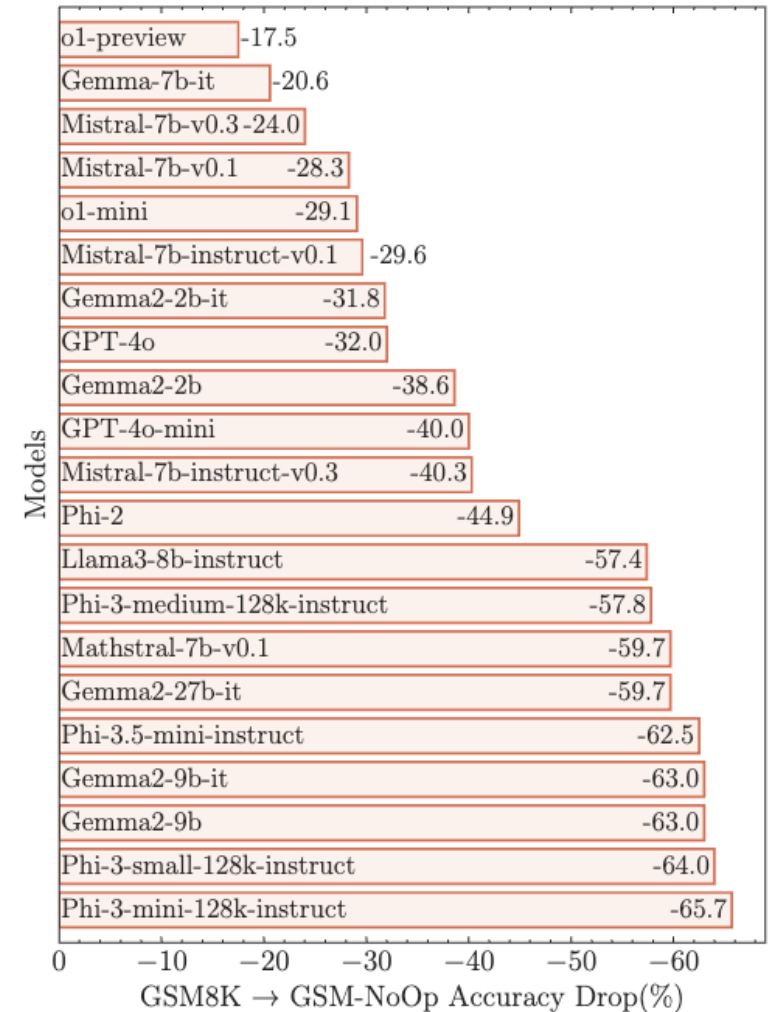
注：本次评测分为客观题（文本）、第6题图像题（单选）和解答题，客观题（文本）总计68分，图像题总共（单选）5分，解答题77分，总分150分

Primary School Math Problem?

Q0: Oliver picked 44 kiwifruits on Friday, then 58 on Saturday. On Sunday, he picked twice as many kiwifruits as he did on Friday. How many kiwifruits does Oliver have in total?

Q1: Oliver picked 44 kiwifruits on Friday and 58 on Saturday. On Sunday, he picked twice as many as on Friday, **though 5 of them were smaller than average.** How many kiwifruits does Oliver have in total?

GPT-o1: '... We need to subtract them from the Sunday total: 88 (Sunday kiwis) – 5 (small kiwis) = 83 kiwis... 44+58+83 equals to...



Chain of Thought (CoT)

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. ✅

Zero-Shot CoT

【1st prompt】 Reasoning Extraction

Q: On average Joe throws 25 punches per minute. A fight lasts 5 rounds of 3 minutes. How many punches did he throw?

A: Let's think step by step.



LLM



In one minute, Joe throws 25 punches.
In three minutes, Joe throws $3 * 25 = 75$ punches.
In five rounds, Joe throws $5 * 75 = 375$ punches.

【2nd prompt】 Answer Extraction

Q: On average Joe throws 25 punches per minute. A fight lasts 5 rounds of 3 ...

A: Let's think step by step.

In one minute, Joe throws 25 punches. ... In five rounds, Joe throws $5 * 75 = 375$ punches. .

Therefore, the answer (arabic numerals) is

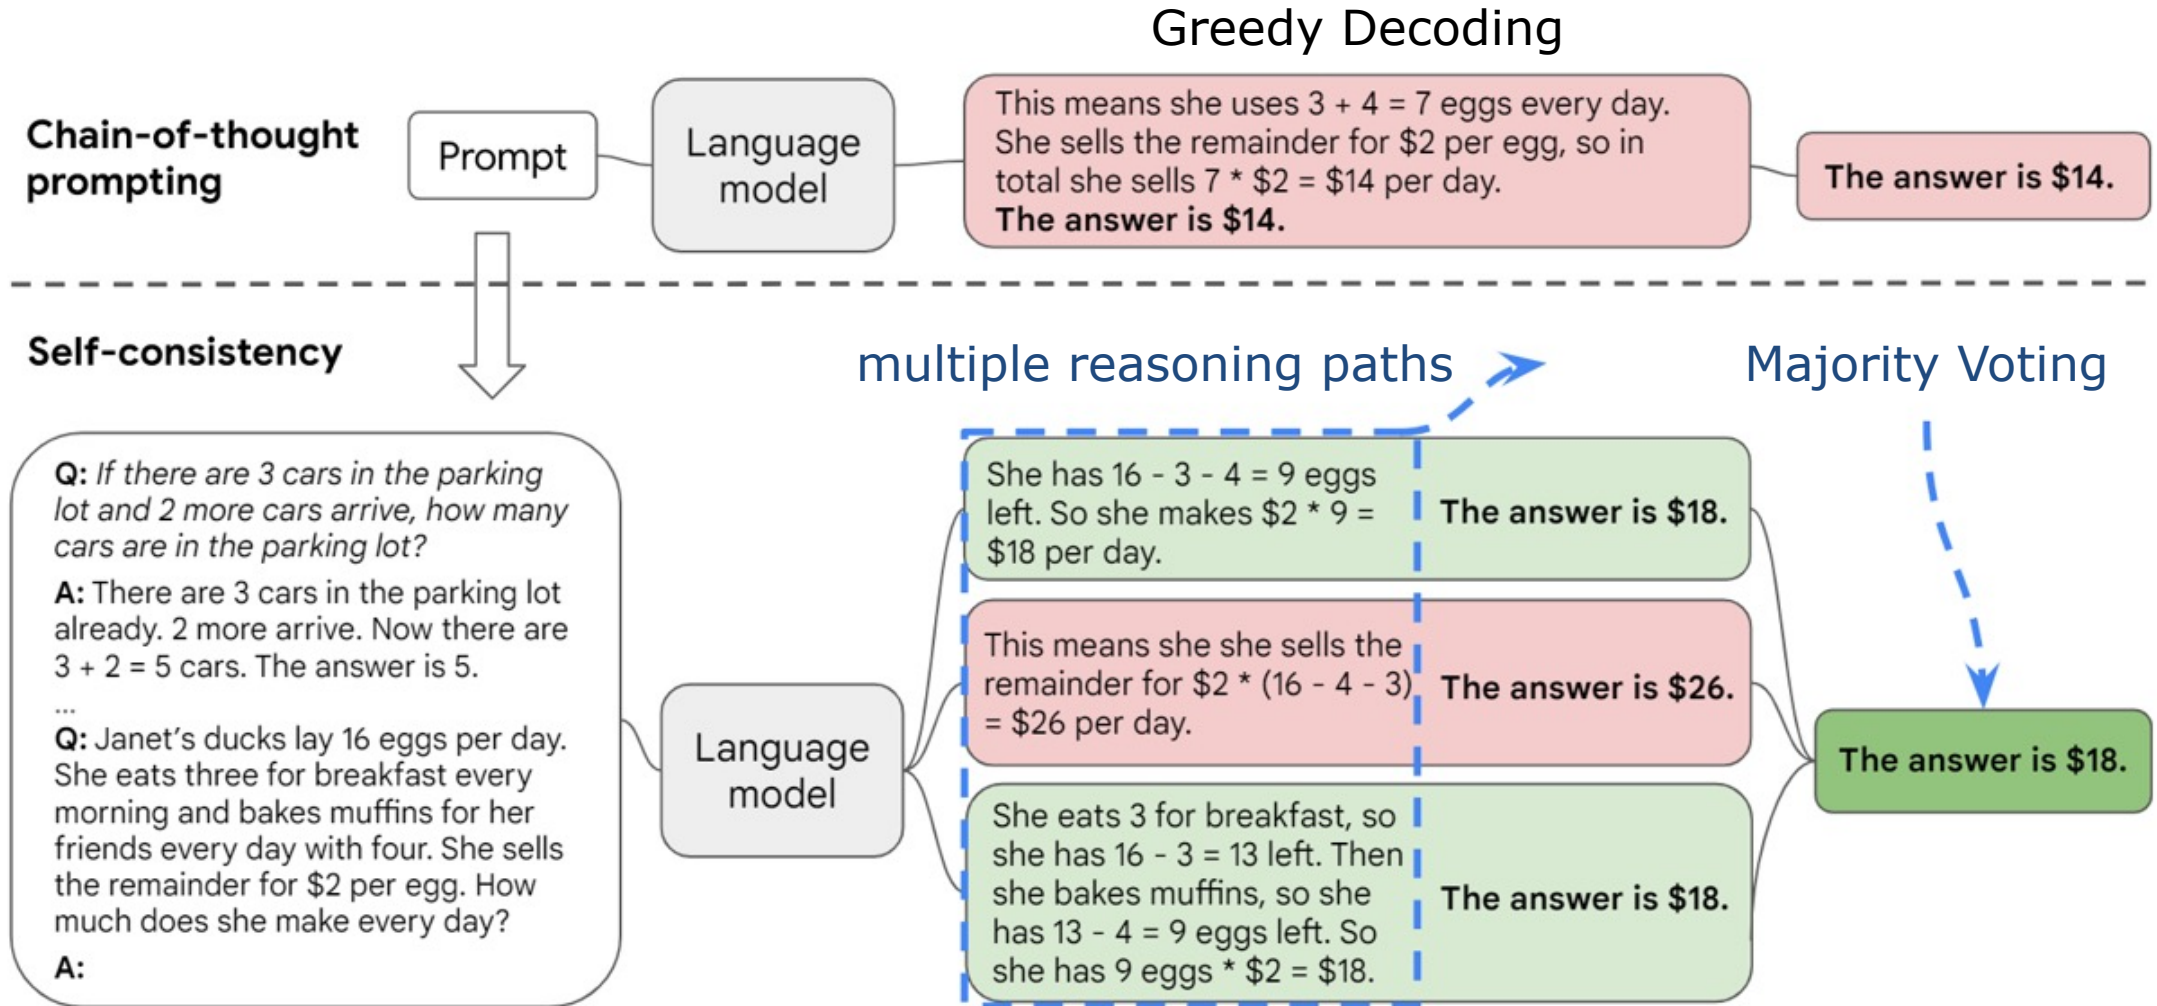


LLM



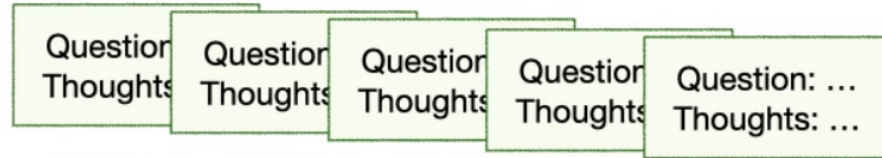
375.

Self-Consistency Improves CoT Reasoning



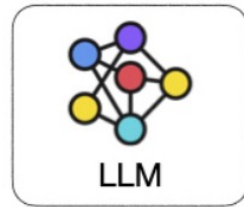
Program-of-Thought (PoT)

Examples



Train 1 departs at 6 AM from city A and arrive in city B in 11 AM. Train 2 departs at 7:30AM from city B and arrive in City A at 11 AM. If both trains depart from their station simultaneously, when will they meet?

Few Shot



Prompt

```
train1_travel_time = 11 - 6
# 7:30 AM refers to 7.5
train2_travel_time = 11 - 7.5
distance_of_city_a_and_b = Symbol("D", positive=True)
train1_speed = distance_of_city / train1_travel_time
train2_speed = distance_of_city / train2_travel_time
ans = distance_of_city / (train1_speed + train2_speed)
```

Question: Toulouse has twice as many sheep as Charleston. Charleston has 4 times as many sheep as Seattle. How many sheep do Toulouse, Charleston, and Seattle have together if Seattle has 20 sheep?

Answer this question by implementing a solver() function.
def solver():
 # Let's write a Python program step by step, and then return the answer
 # Firstly, we need define the following variable:

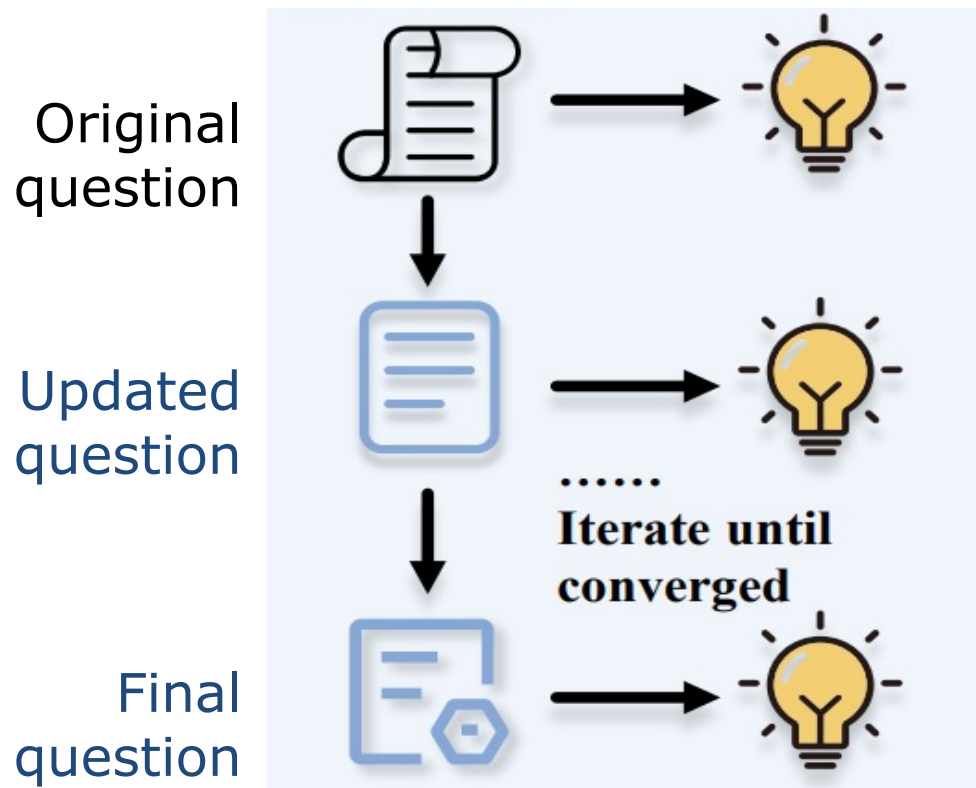


Prompt Zero Shot

```
Seattle = 20
Charleston = Seattle * 4
Toulouse = Charleston * 2
# Then, we need to calculate the sum of the three
sum = Seattle + Charleston + Toulouse
# Finally, we need to return the answer
return sum
```

Refining the Problem Itself

► Self-Polish



[Original Problem₀] : Kylie makes 10 beaded necklaces on Monday and 2 beaded necklaces on Tuesday. Then Kylie makes 5 beaded bracelets on Wednesday. 20 beads are needed to make one beaded necklace. 10 beads are needed to make one beaded bracelet. Ada bought 2000 tomatoes from the grocery store. How many beads does Kylie use in total to make her jewelry?

[Answer₀] : 120.



Begin, Problem Refine (Remove Irrelevant Information)

[Refined Problem₁] : Kylie makes 12 beaded necklaces, 5 beaded bracelets. Each beaded necklace needs 20 beads. Each beaded bracelet requires 10 beads. How many beads does Kylie use in total to make her jewelry?

[Answer₁] : 155.



A₁ != A₀, Continue Refine (Reorder Conditions)

[Refined Problem₂] : Kylie makes 12 beaded necklaces, and each beaded necklaces needs 20 beads. She also makes 5 beaded bracelets, and each beaded bracelet needs 10 beads. How many beads does Kylie use in total to make her jewelry?

[Answer₂] : 290.



A₂ != A₁, Continue Refine (Summary Local Conditions)

[Refined Problem₃] : Kylie requires 240 beads to make beaded necklaces. She also requires 50 beads to make beaded bracelets. How many beads does Kylie use in total to make her jewelry?

[Answer₃] : 290.



A₃ == A₂, Return A₃

Uncertainty-aware Adaptive Guidance

Uncertainty Identification

Q: Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers' market?

A: [1] 16 eggs are laid per day. [2] Janet eats 3 for breakfast every day. [3] $16 - 3 = 13$ eggs are left. [4] ~~13 eggs are baked into muffins.~~

Adaptive Reasoning Adjustment

Q_D: Alfred is storing a tonne of maize each month for the next 2 years. If 5 tonnes ..., how many tonnes of maize does he have at the end of the 2 years.

A_D: Before his maize is stolen, the maize he stores at the end of 2 years is $12 * 2 = 24$ tonnes of maize. After his maize is stolen, he is left with ...

Q: Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers' market?

A: [1] 16 eggs are laid per day. [2] Janet eats 3 for breakfast every day. [3] $16 - 3 = 13$ eggs are left. [4] Janet bakes muffins with 4 eggs. [5] $13 - 4 = 9$ eggs are left. [6] Janet sells 9 eggs at \$2 per egg. [7] So she makes \$18 per day.

Demonstration Clustering

Q_{D1}: Natalia sold clips to 48 of her friends in April, and then she sold half as many clips...

A_{D1}: Natalia sold $48/2 = 24$ clips in May. Natalia sold $48 + 24 = 72$ clips altogether in April ...

Q_{D2}: Weng earns \$12 an hour for babysitting. Yesterday, she just did 50 minutes of ...

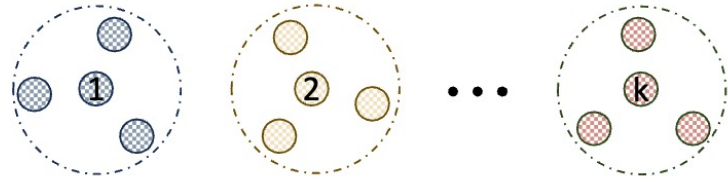
A_{D2}: Weng earns $12/60 = 0.2$ per minute. Working 50 minutes, she earned 0.2×50 ...

Relevance 
Originality 

•
•

Q_{Dk}: At 30, Anika is $4/3$ the age of Maddie. What would be their average age in 15 years?

A_{Dk}: If Anika is 30 now, in 15 years, she'll be $30 + 15 = 45$ years old. At 30, Anika is $4/3$...



the model **monitors the degree of uncertainty** in every step of its reasoning in real-time
Once the model detects that the uncertainty level crosses a certain limit, it **traces back** to the source of the doubt

Reinforcement Learning

Question: Venn earns \$12 per hour as a babysitter. Yesterday, she worked 50 minutes. How much money did she make?

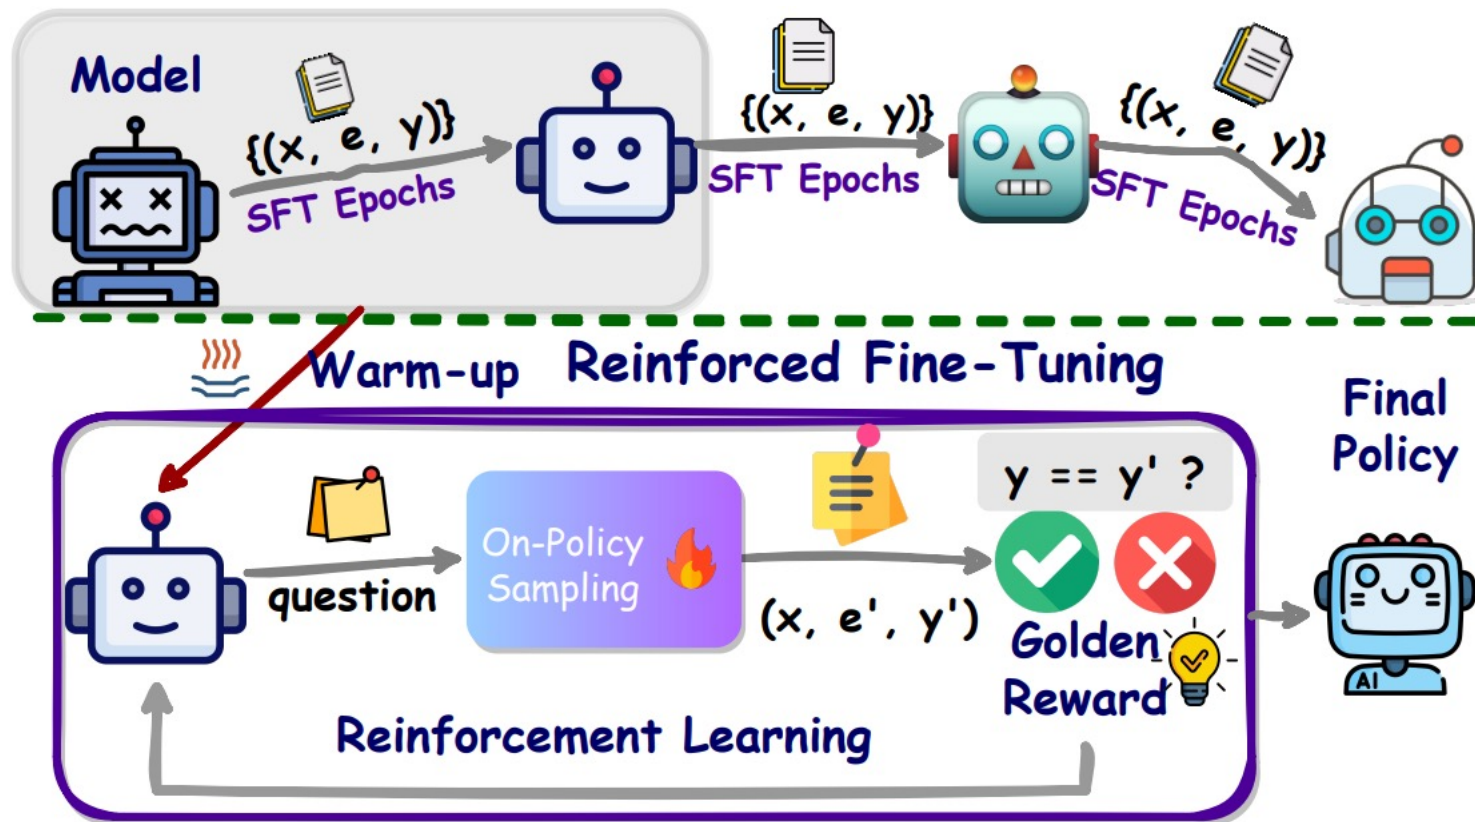
Chain of Thought (e):

We first calculate her hourly rate and then multiply it by the working hours.

Step 1: Convert 50 minutes to hours. Since 1 hour = 60 minutes, 50 minutes = $50 \div 60 = 5/6$ of an hour.

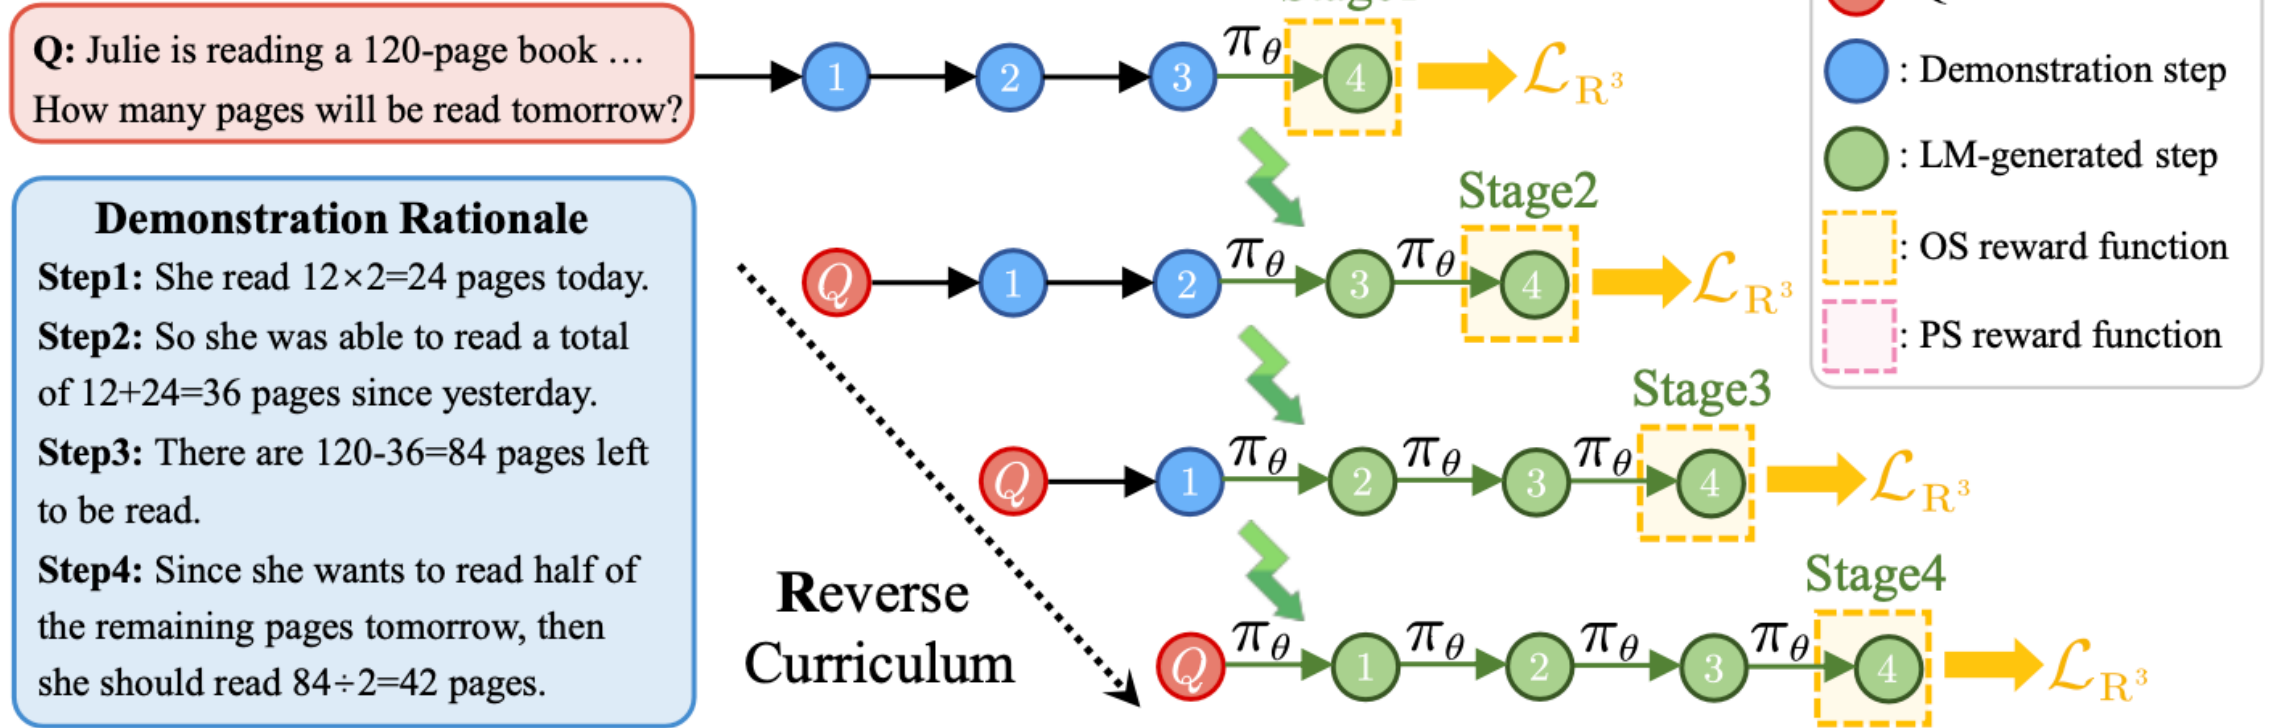
Step 2: Calculate earnings by multiplying the hourly rate by the working hours: $\$12/\text{hour} \times 5/6 \text{ hour} = \10 .

Therefore, Venn earned \$10 for 50 minutes of babysitting. The answer is 10.



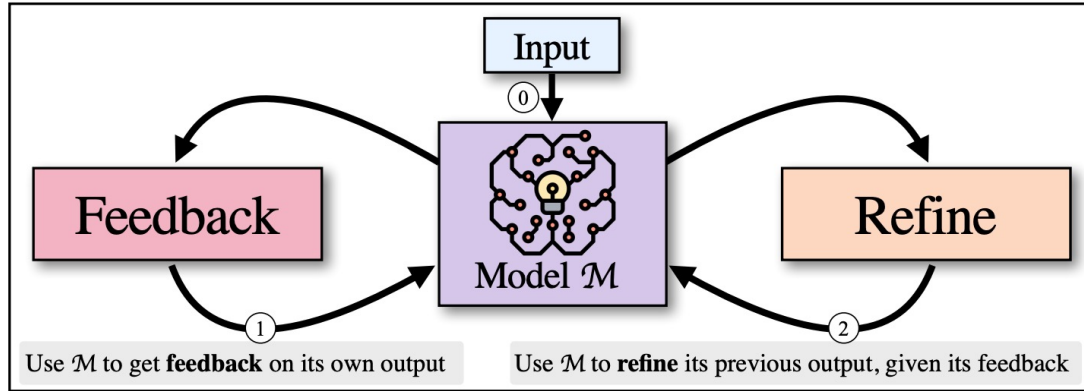
Process Supervision

R^3 : Learning Reasoning through Reverse Curriculum RL



Process Supervision focuses on every single detail to ensure the reasoning chain is complete and accurate. Every correct calculation is recognized and rewarded, which motivates the model to follow the right logical order.

Self-Refine



Given an input (0):

SELF-REFINE starts by generating an output and passing it back to the same model M to get feedback (1). The feedback is passed back to M , which refines the previously generated output (2). Steps (1) and (2) iterate until a stopping condition is met.

Questions:

1. Could you **review** your previous answer step by step to ensure no details are omitted?
2. Would you mind **breaking down** your last response to check for any oversights?
3. I would greatly appreciate it if you could **re-examine** your recent answer, analyze it step by step, and see if anything is missing.
4. Could you thoroughly **analyze** your previous response to ensure all aspects are covered?
5. Please dissect your earlier answer step by step and **verify** that all points have been addressed.

Stepwise Self-Refine

Question: In a 8 fluid ounce bottle of Vitamin Water, there are 125 calories. How many calories would be contained in a 12 fluid ounce bottle?

Cot Response

Let's break it down step by step:

1. **Calculate the calories per ounce:**
 - In a 8 fluid ounce bottle, there are 125 calories.
 - Therefore, the calories per ounce are:
$$\frac{125}{8} = 15.625 \text{ calories per ounce}$$
2. **Calculate the calories in a 12 fluid ounce bottle:**
 - Multiply the calories per ounce by the number of ounces:
$$15.625 \times 12 = 187.5 \text{ calories}$$
3. **Final answer:**
 - The answer is 187.5.

Self-talk Response

Let's think deeply about the question.

We need to determine how many calories would be contained in a 12 fluid ounce bottle of Vitamin Water. First, let's calculate the number of calories in a 12 fluid ounce bottle.

We know that there are 125 calories in an 8 fluid ounce bottle. To find the number of calories in a 12 fluid ounce bottle, we multiply the number of ounces by the number of calories per ounce:

$$12 \times 125 = 1500$$

Wait, let's check this calculation. Oops, that's not quite right. The calculation should be based on the ratio of ounces, not multiplying by 125 directly. The correct approach is to find the number of calories per ounce by dividing the total calories by the number of ounces in the bottle:

$$\frac{125}{8} = 15.625$$

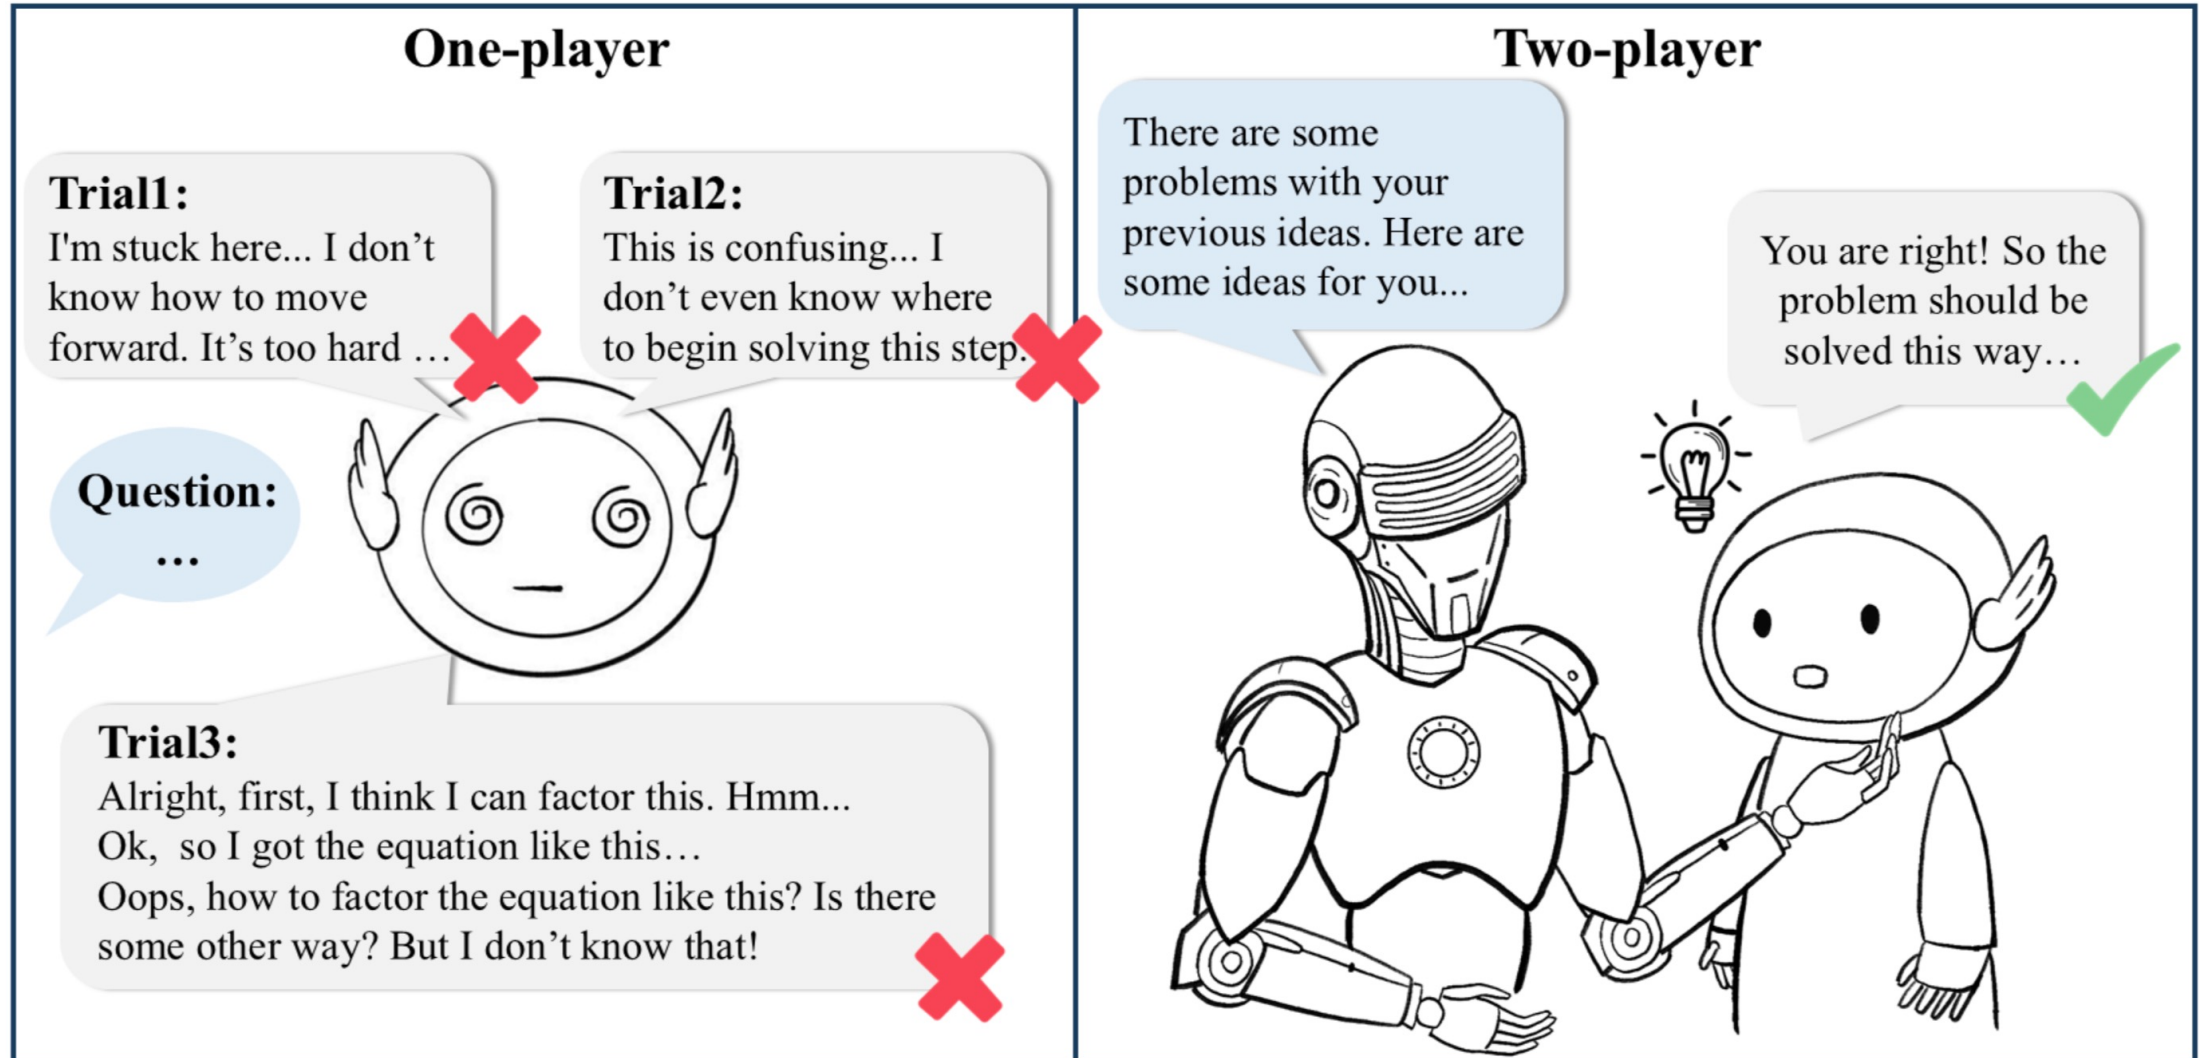
Now, let's calculate the calories in a 12 fluid ounce bottle using the correct ratio:

$$12 \times 15.625 = 187.5$$

Looks good! The calculation is now based on the correct ratio, and the result is accurate.

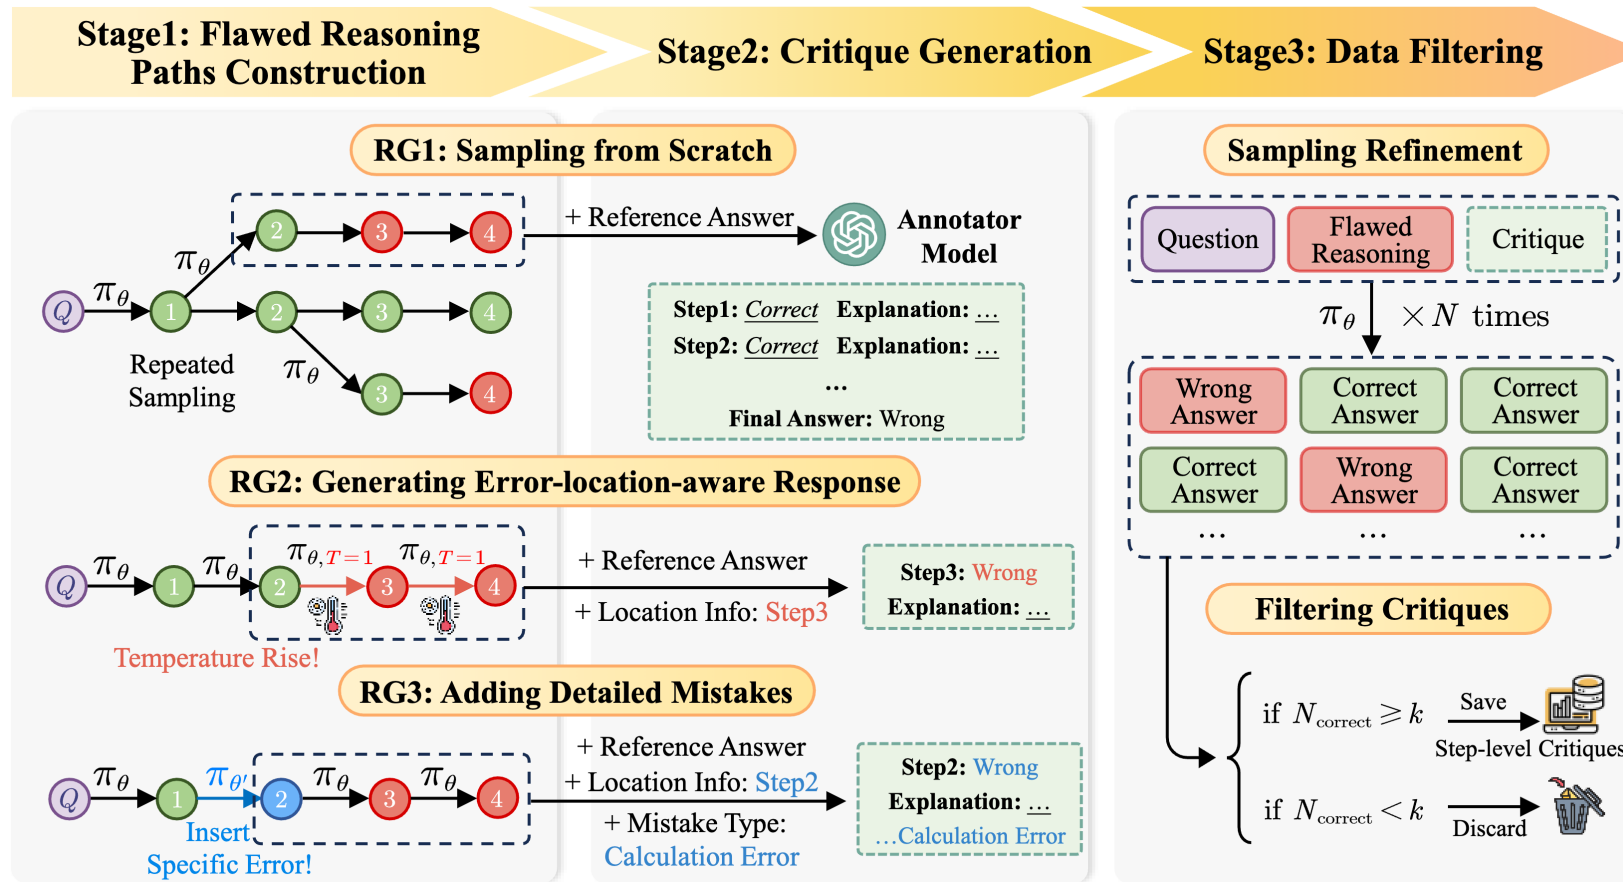
Answer: The answer is 187.5.

Enhancing LLM Reasoning via Critique Models



Enhancing LLM Reasoning via Critique Models

AutoMathCritique: Three-Step Strategy for Constructing Step-Level Critique and Feedback Data



Enhancing LLM Reasoning via Critique Models

Question: For how many positive integer values of x is the sum $x^2 + 4x + 4$ less than 20?

Original Response

Let's break it down step by step:

1. We have that $x^2 + 4x + 4 < 20$.

2. We can factor this as $(x + 2)^2 < 16$.

...

5. There are 5 positive integer values of x that satisfy this inequality. The answer is 5.

Feedback from Critique Model

Step1: We have that $x^2 + 4x + 4 < 20$.

Correctness of the step: Correct

Explanation: The given inequality is correctly written as $x^2 + 4x + 4 < 20$.

Step2: We can factor this as $(x + 2)^2 < 16$.

Correctness of the step: Wrong

Explanation: The quadratic expression $x^2 + 4x + 4$ correctly factors into $(x + 2)^2$, but the inequality is written incorrectly. The correct inequality should be $(x + 2)^2 < 20$. This error leads to an incorrect constraint in later steps.

...

Correctness of the final answer: Wrong

Explanation: The solution mistakenly concludes that there are 5 positive integer values of x that satisfy the inequality, but the correct count is 2.

Refined Response

Let's break it down step by step:

We start with the inequality $x^2 + 4x + 4 < 20$.

Step 1: Factor the Expression. The expression $x^2 + 4x + 4$ can be factored as: $(x + 2)^2 < 20$

...

Step 5: Identify Positive Integer Solutions. The positive integer solutions for x are: $x = 1$ and $x = 2$.

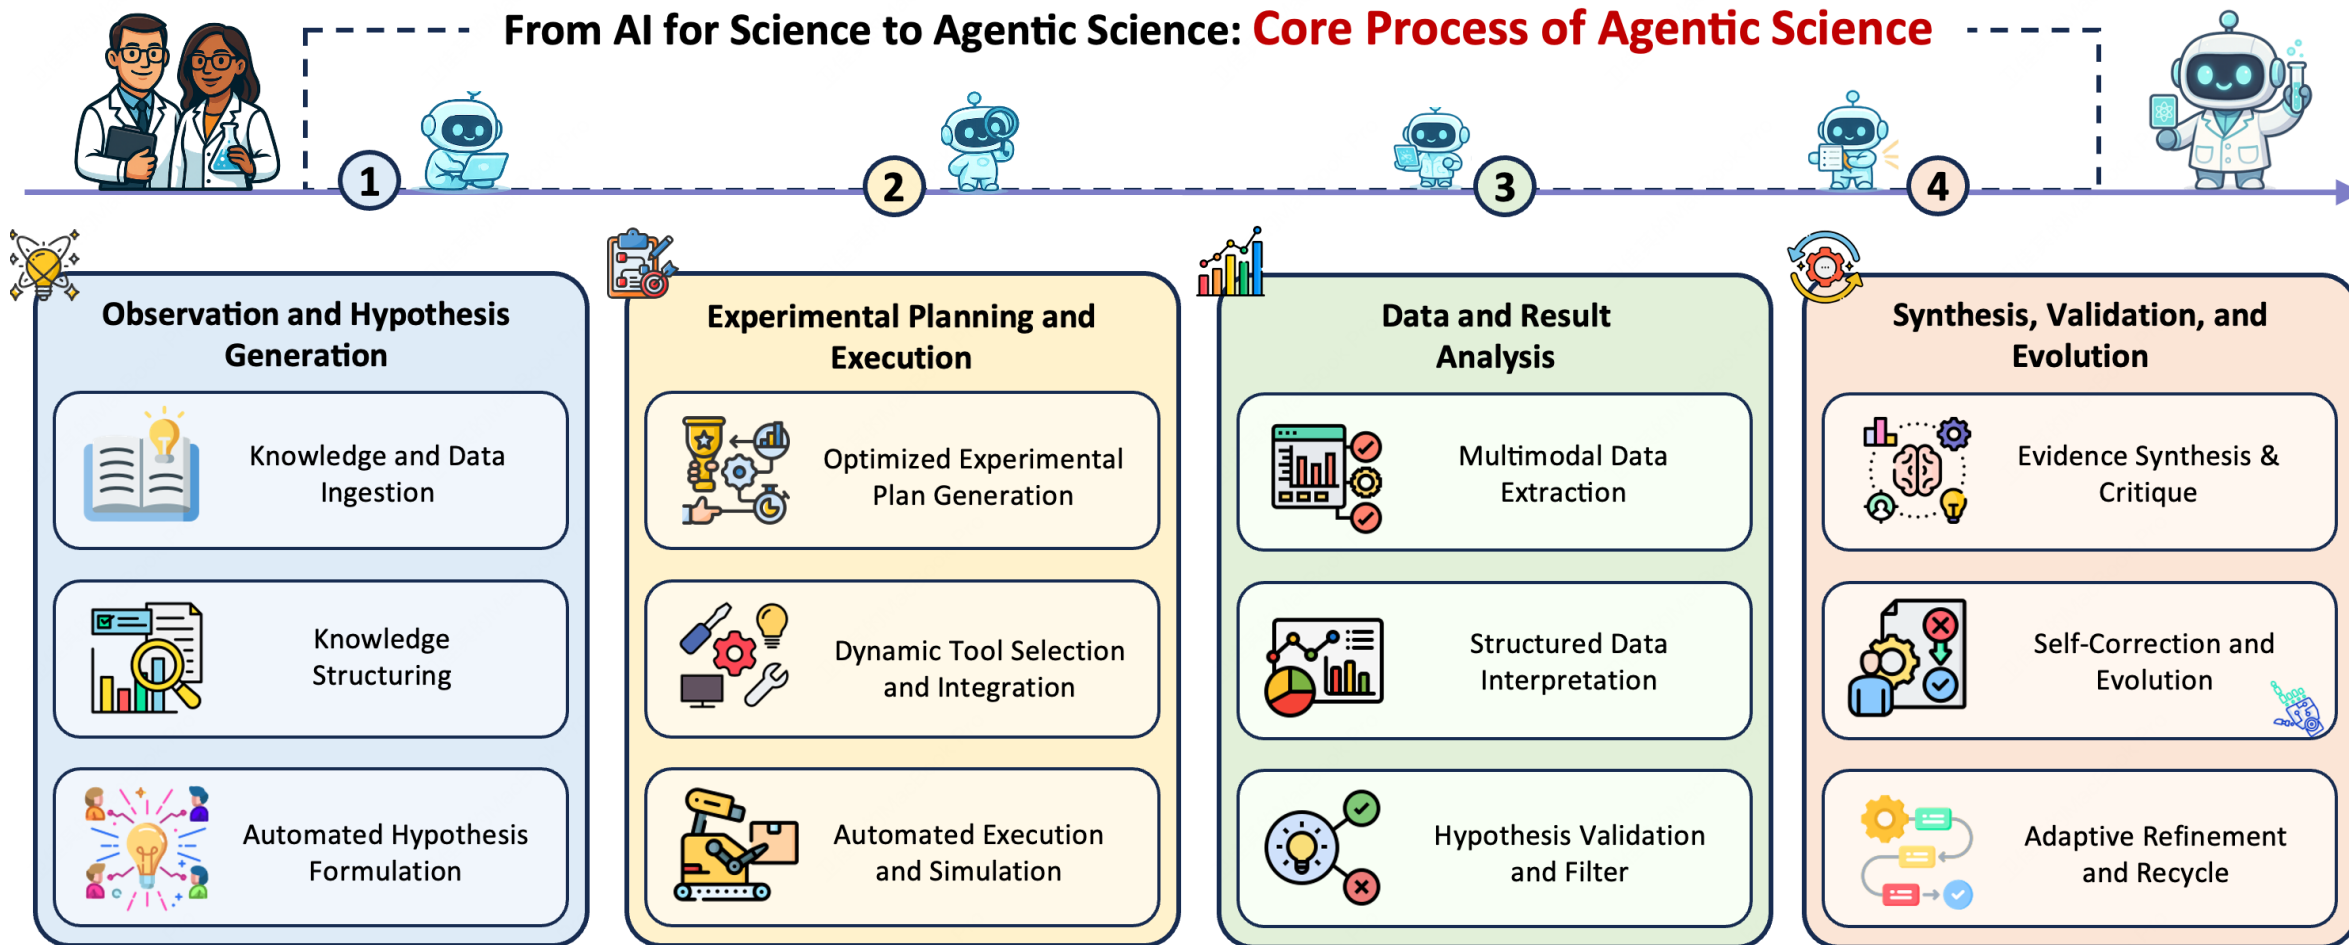
The positive integer values of x that satisfy the inequality $x^2 + 4x + 4 < 20$ are 1 and 2. The answer is 2.



LLMs Enhanced Scientific Discovery

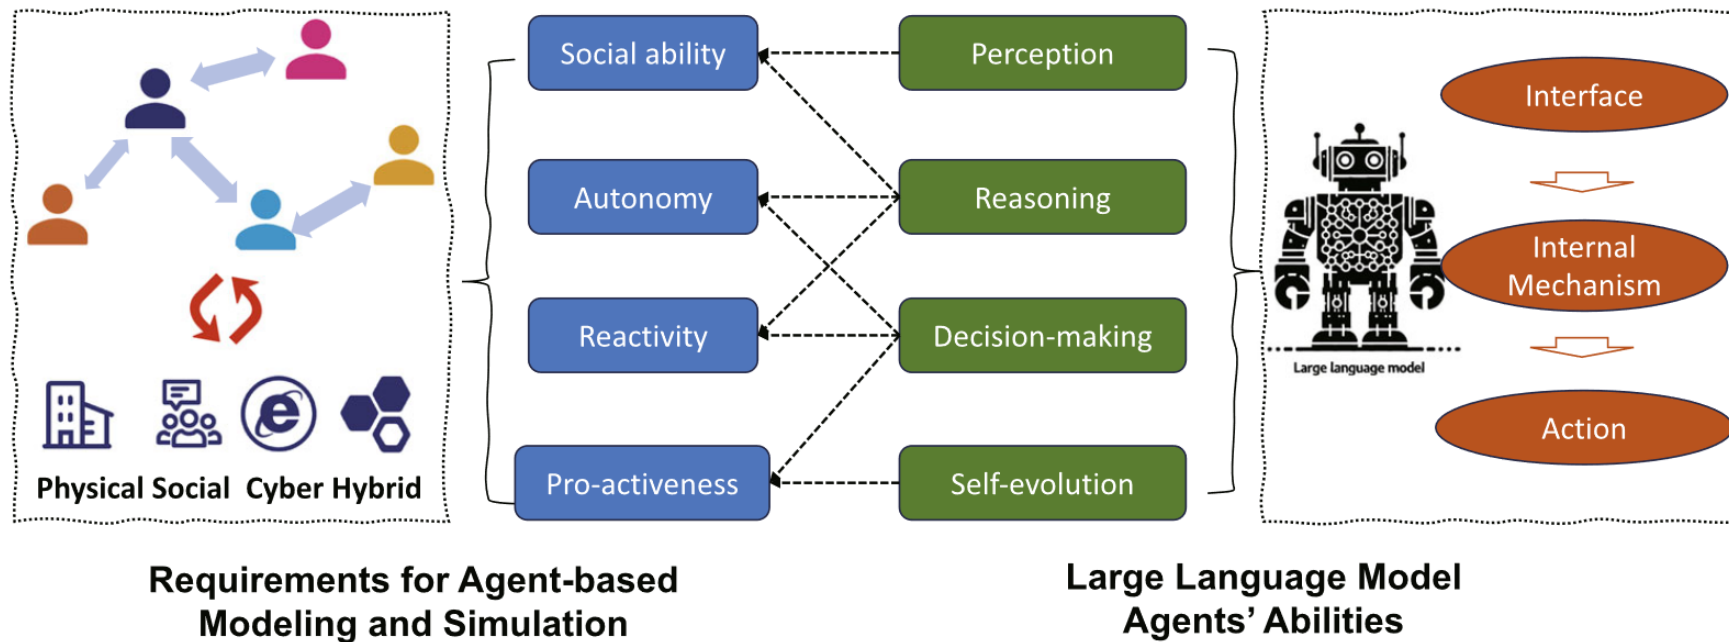
Agentic Science

From AI for Science to Agentic Science: **Core Process of Agentic Science**

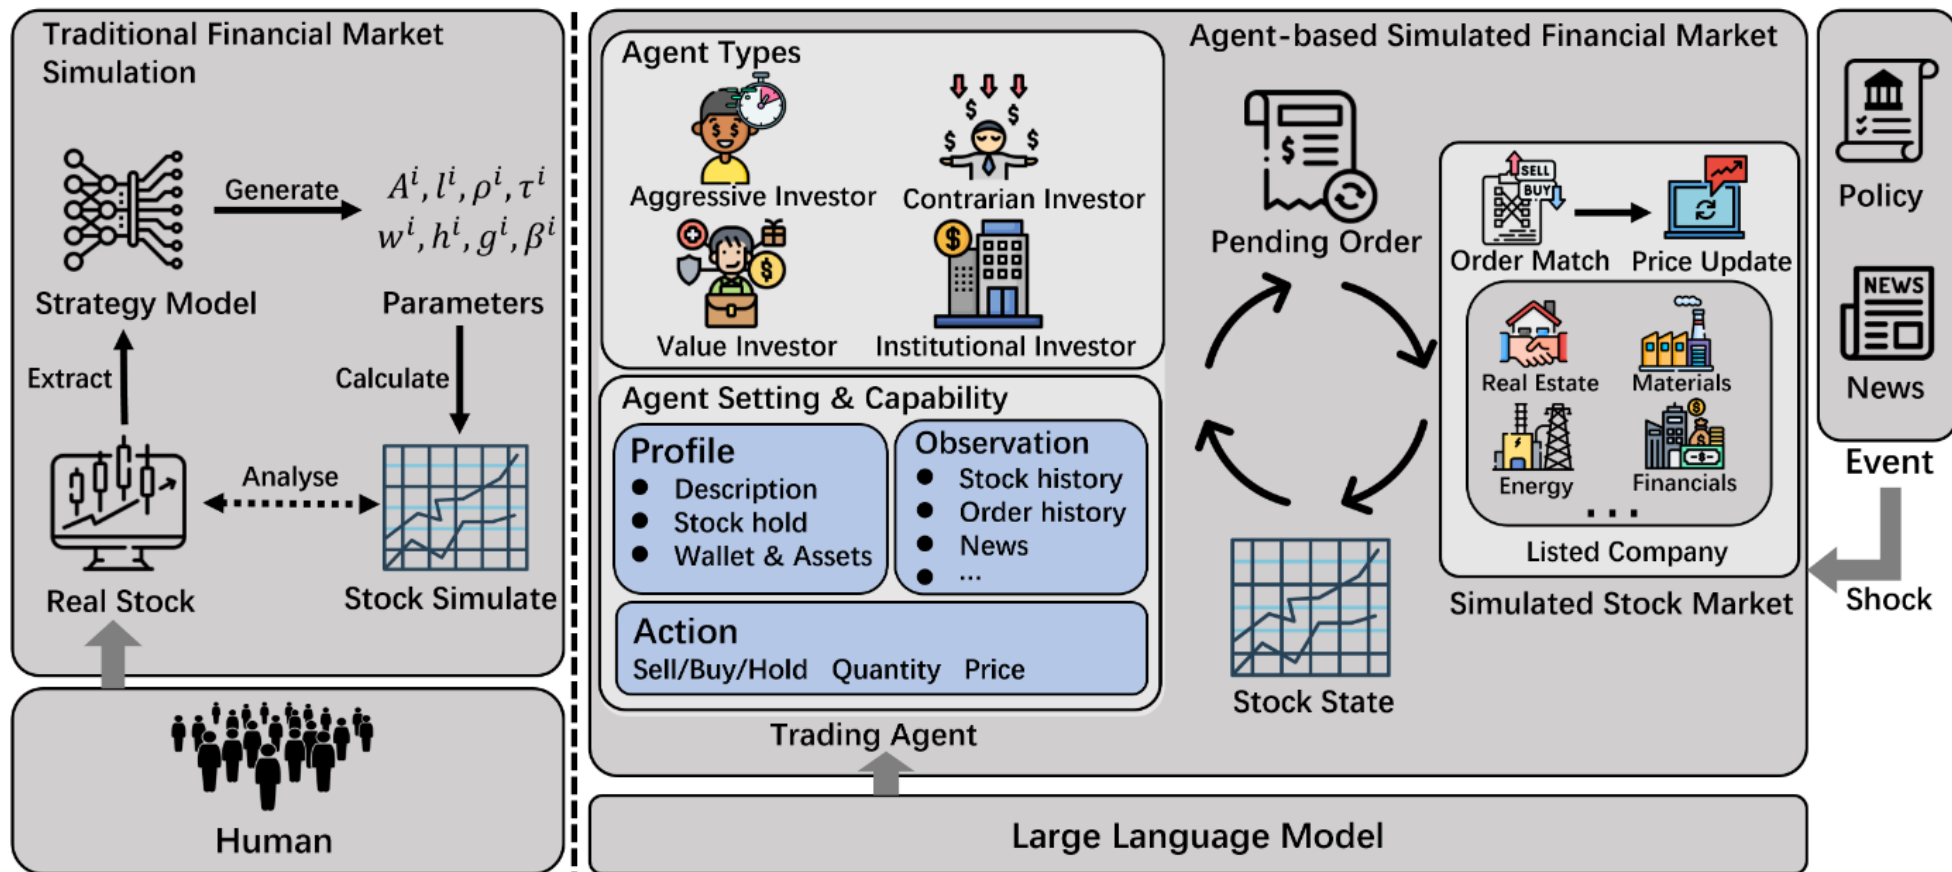


Agent-based Modeling and Simulation

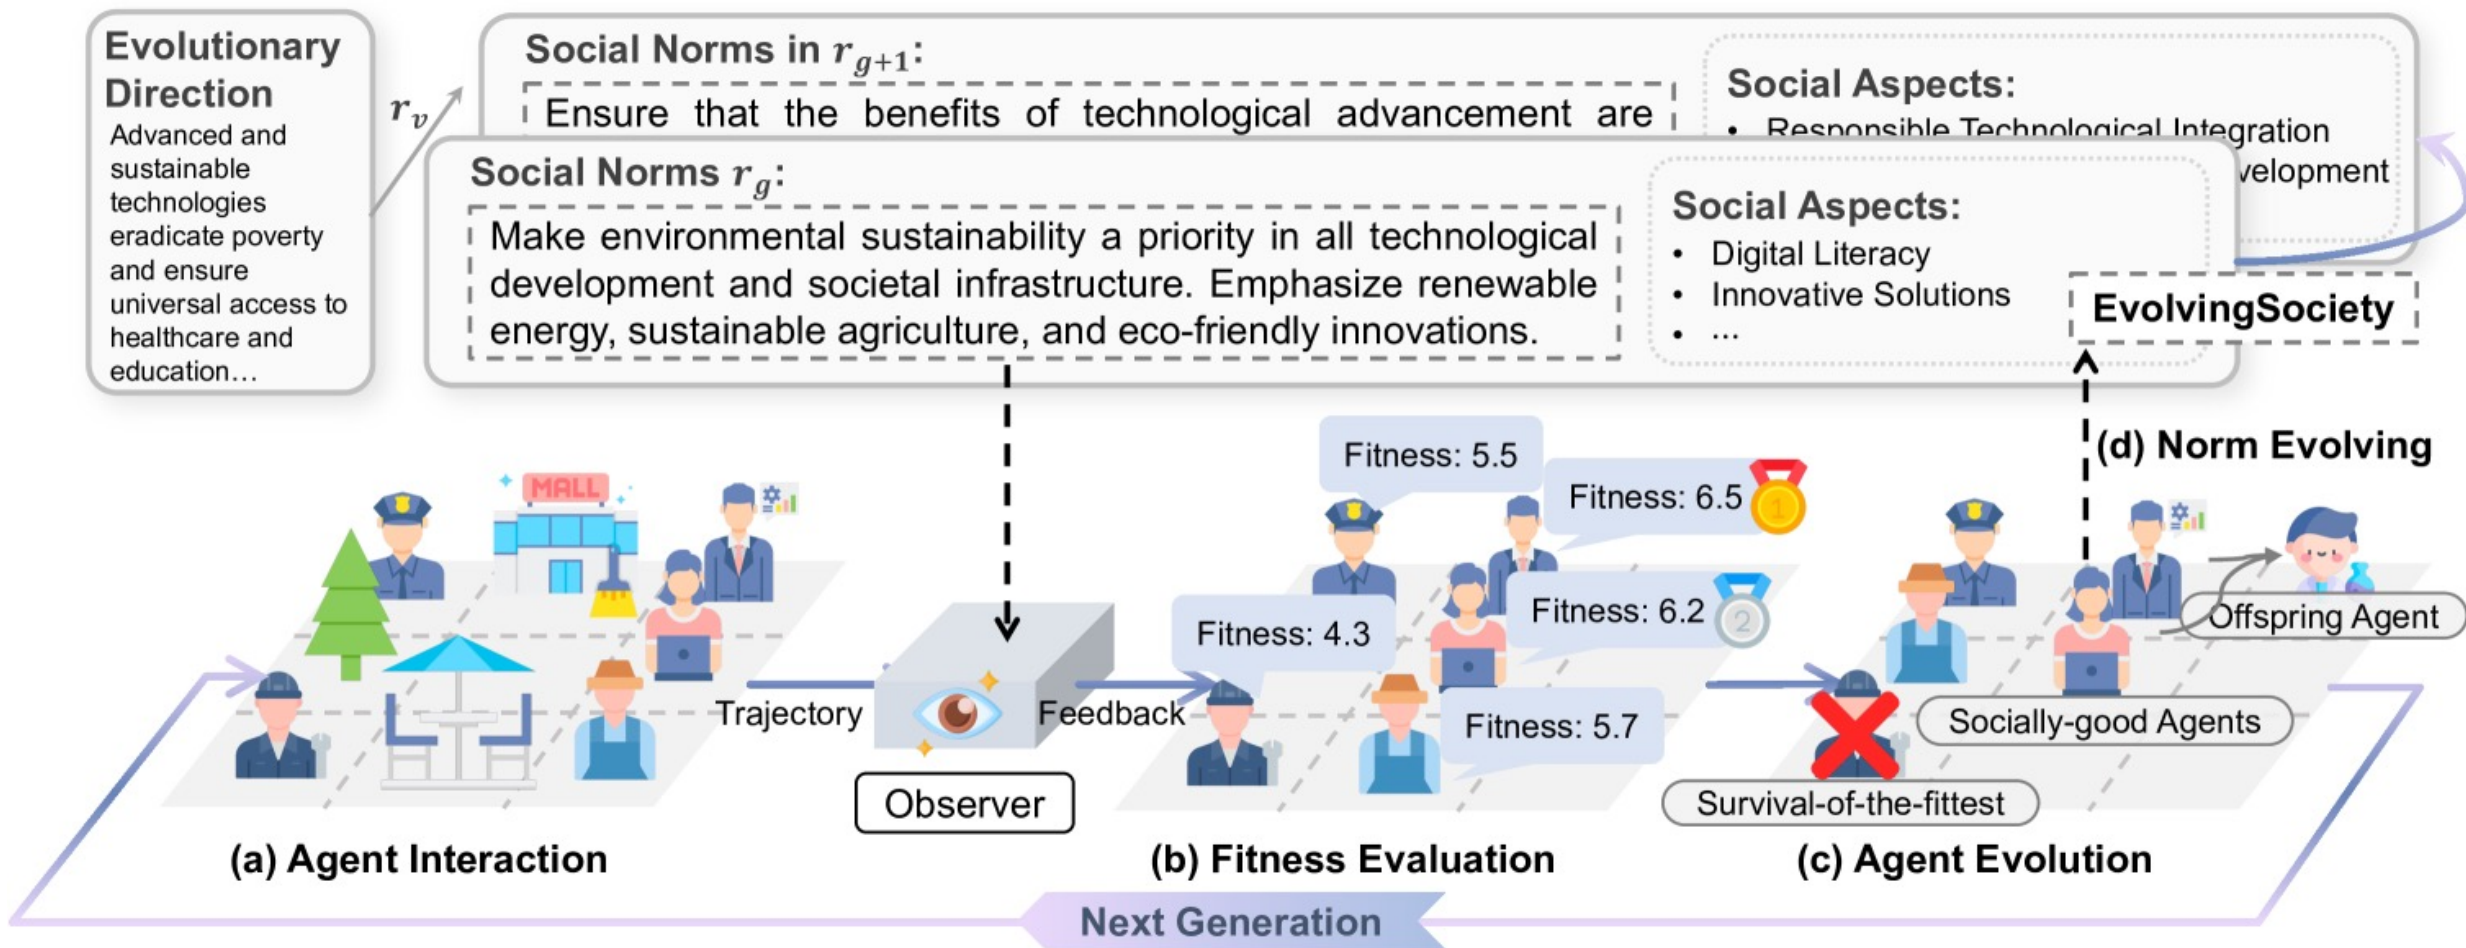
- **Agent-based Modeling and Simulation** (ABMS). This is a way to study complex systems by simulating the actions and interactions of many individual agent.
- Each agent has the ability to make its own decisions based on its environment and what other agents are doing. Their interactions allow researchers to predict overall system behaviors, such as market changes, traffic flow, or how information spreads in a social network.



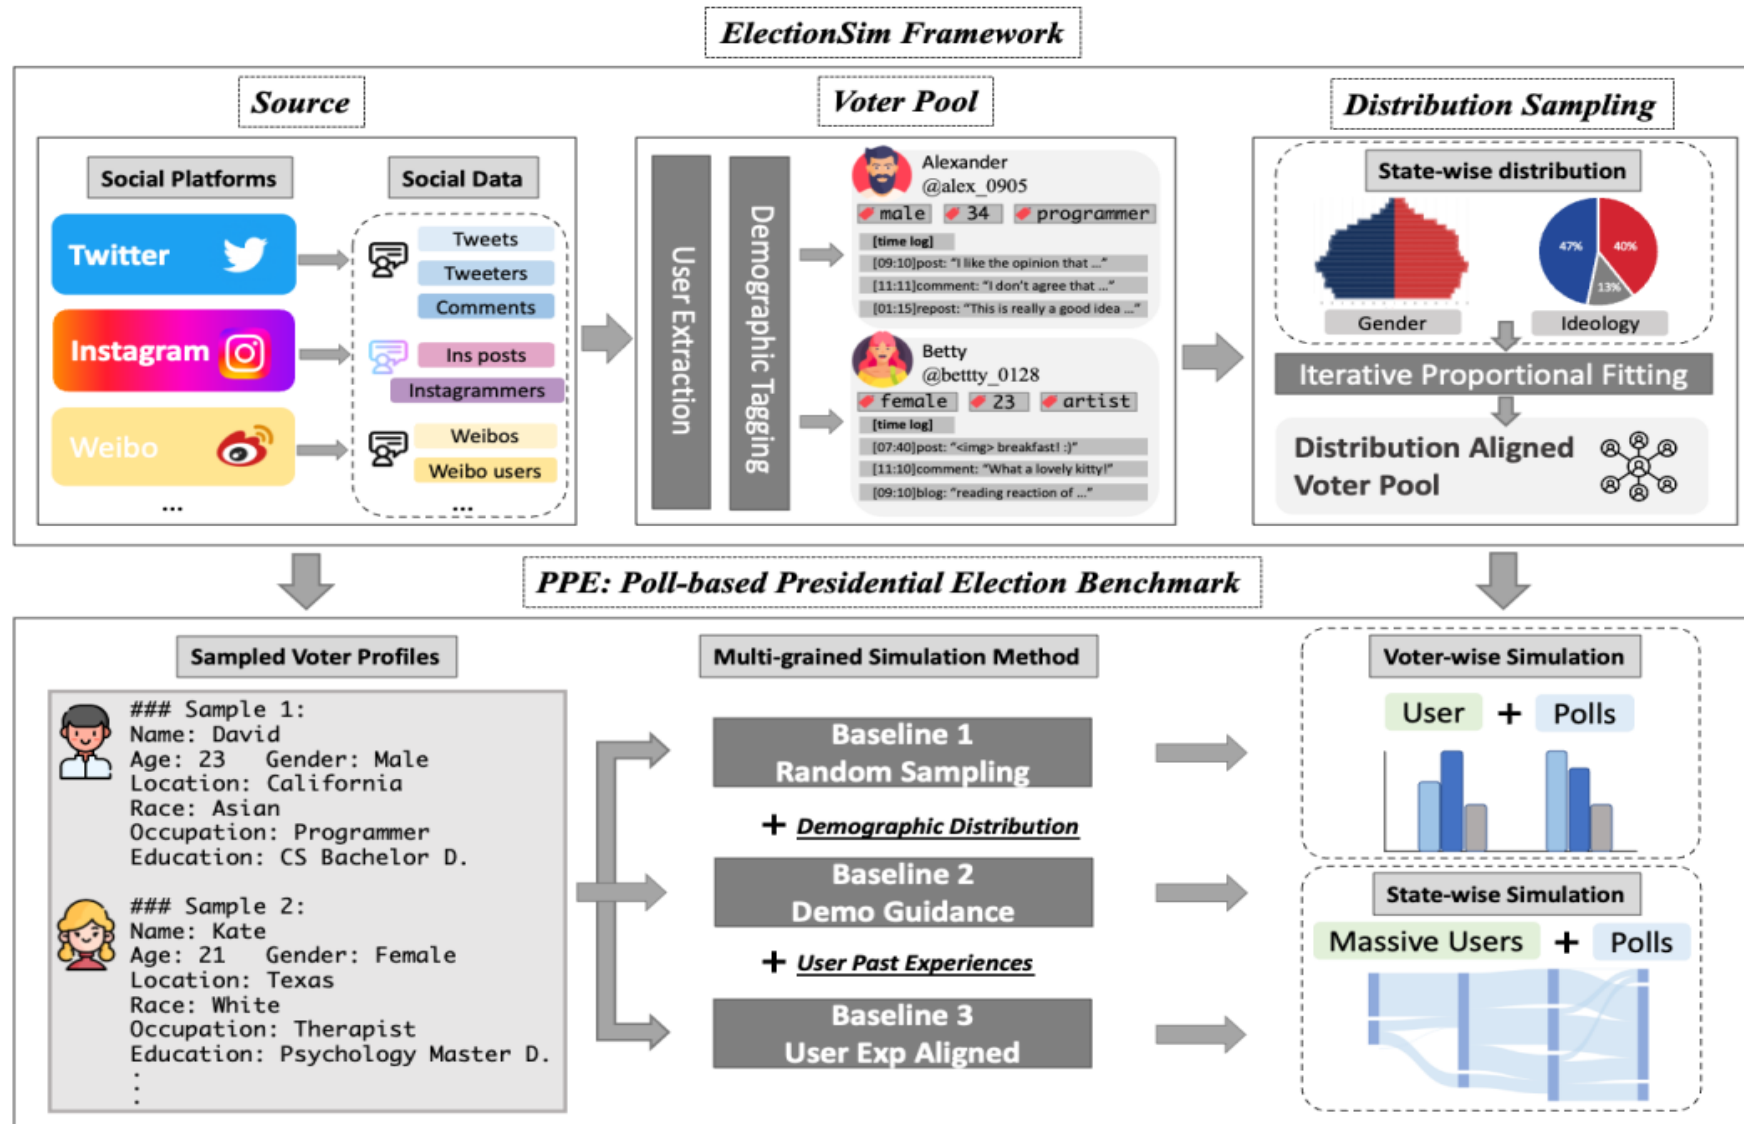
Simulating Financial Market



Evolving Social Norms



Massive Population Election Simulation



Massive Population Election Simulation

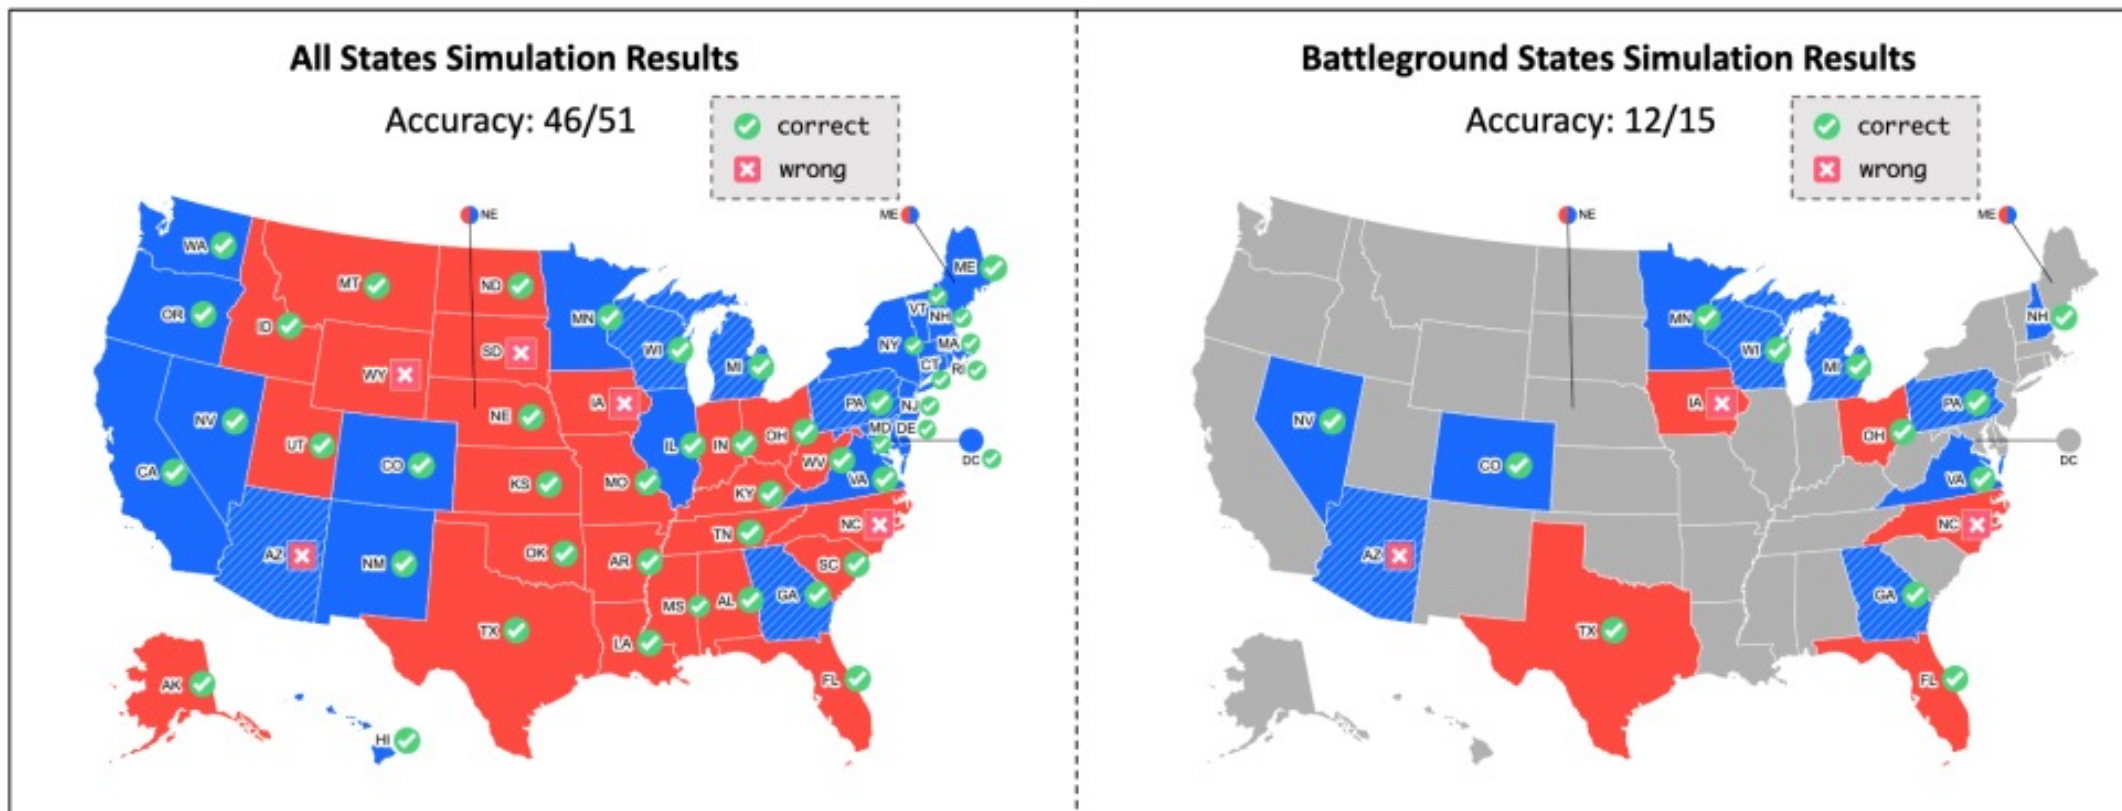
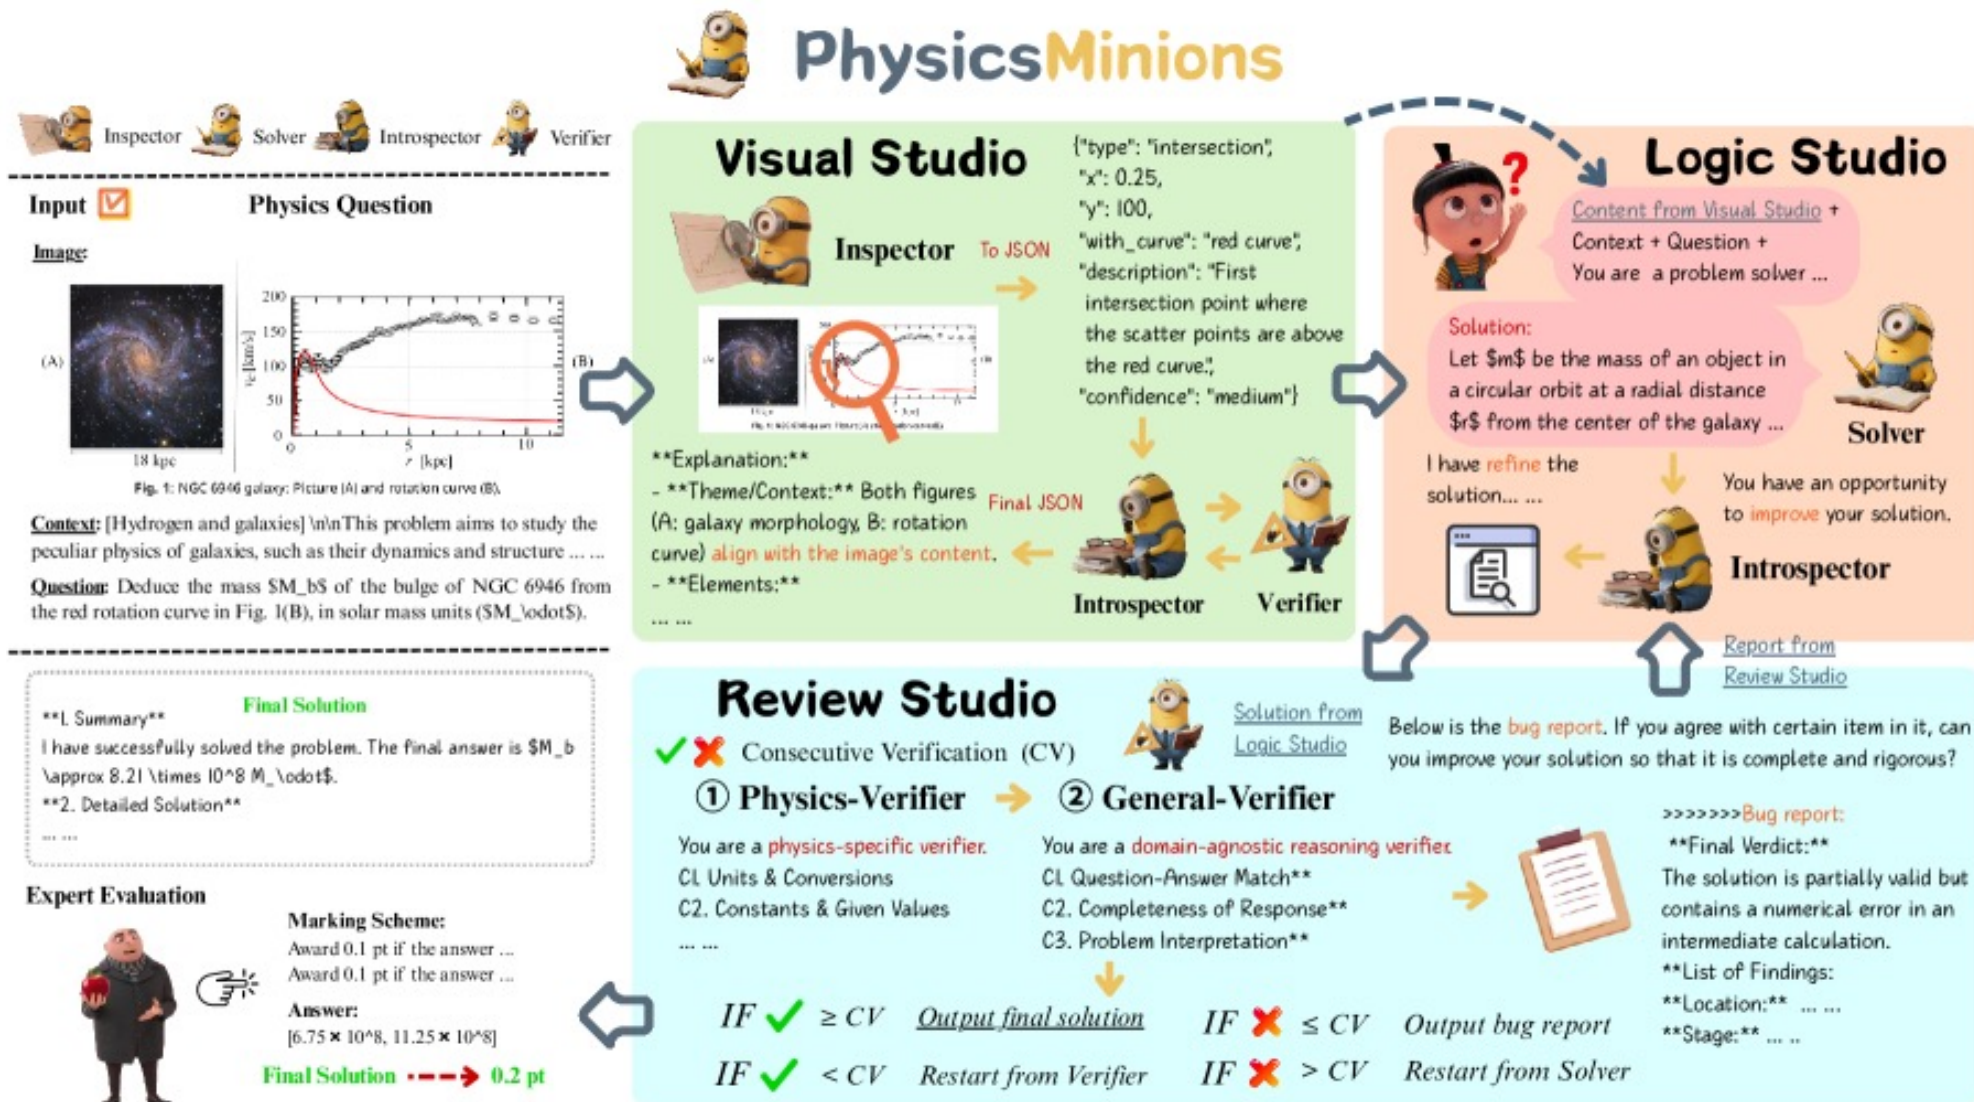


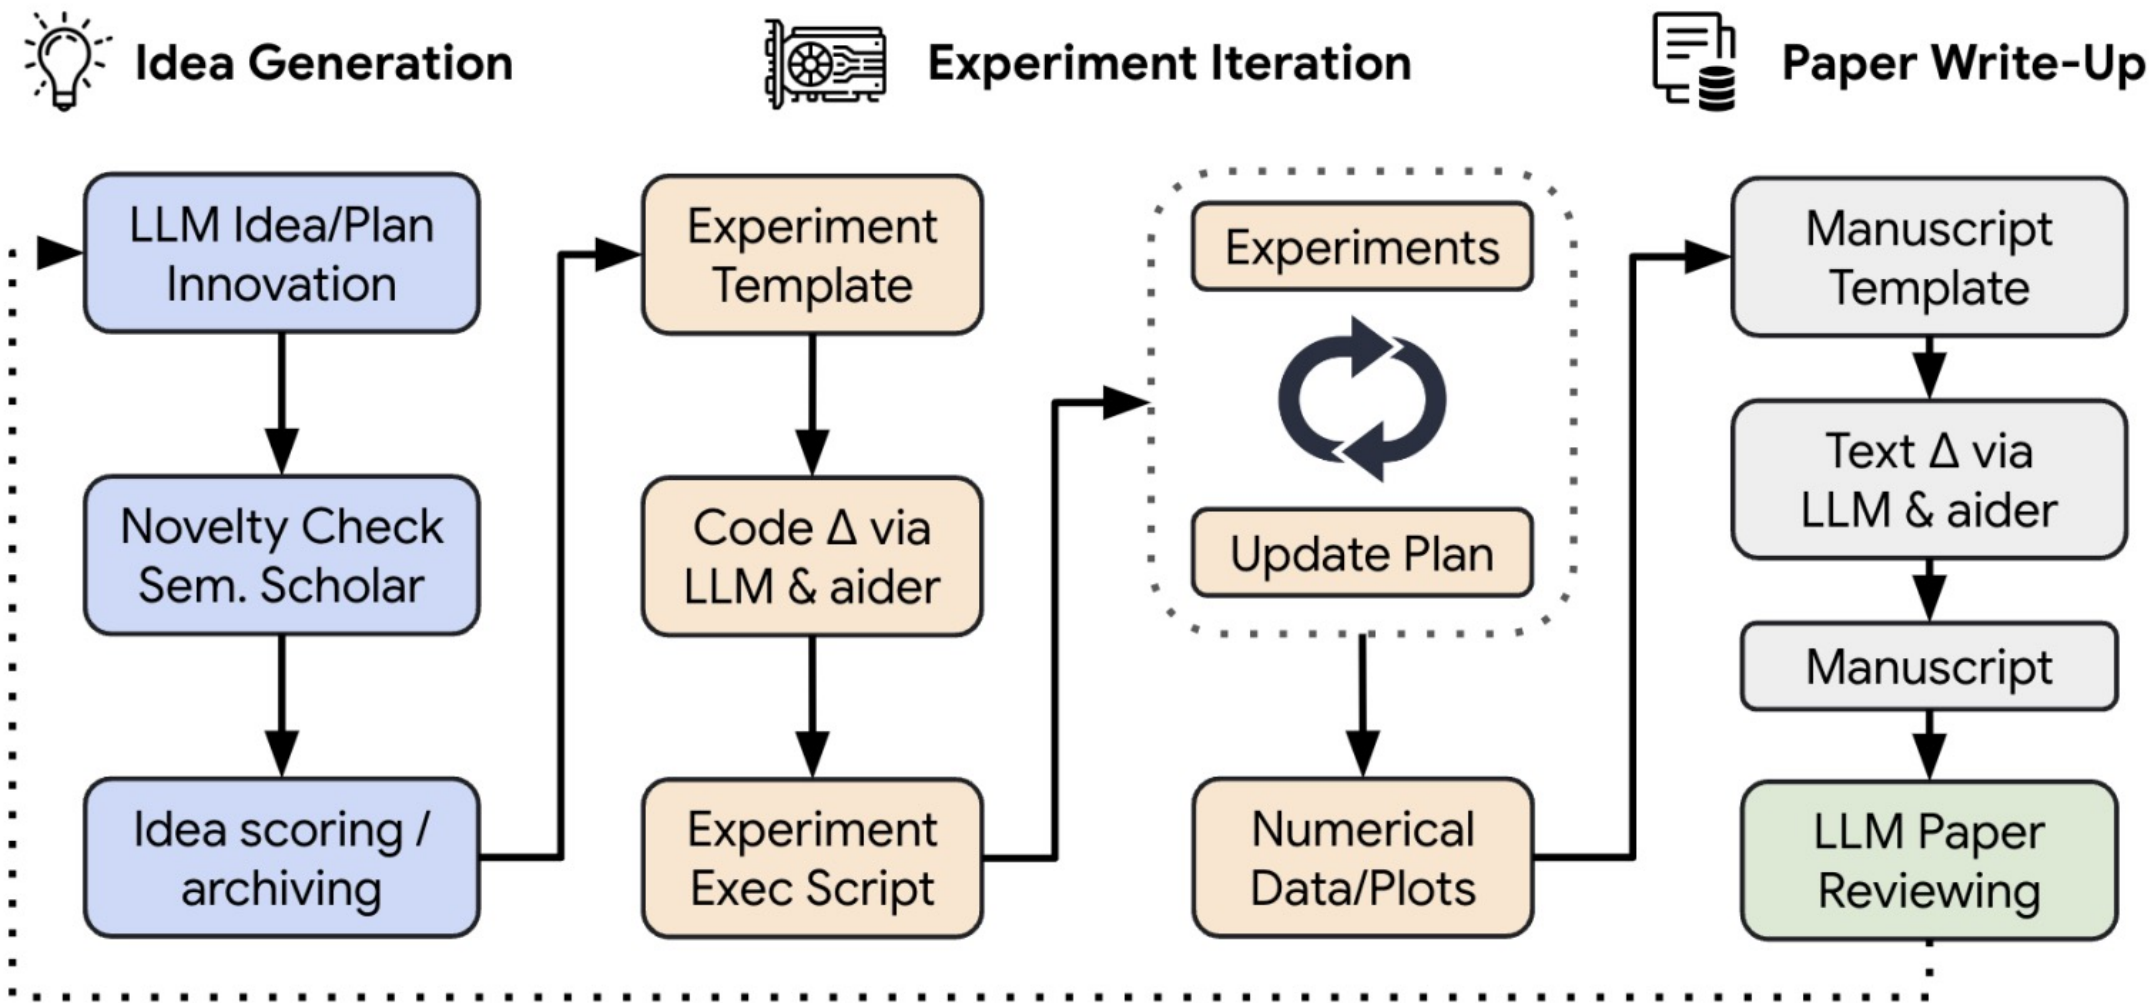
Figure 1: Simulation results of the 2020 Presidential Election. The colors represent the real-world results and the marks represent the simulation results accuracy.

Multimodal Multi-Agent System for IphO

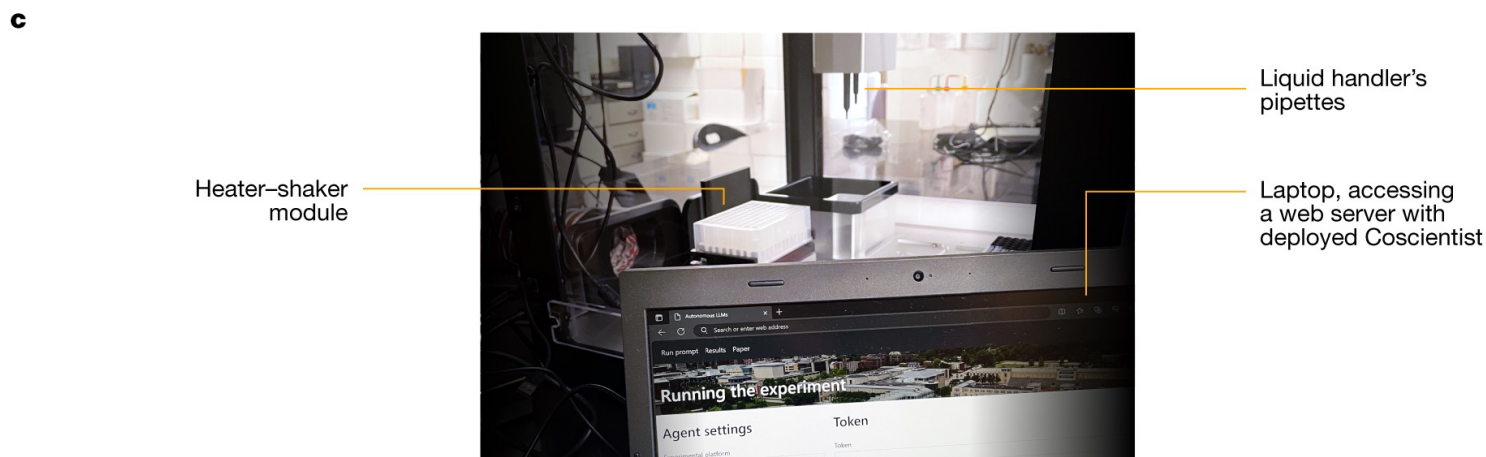
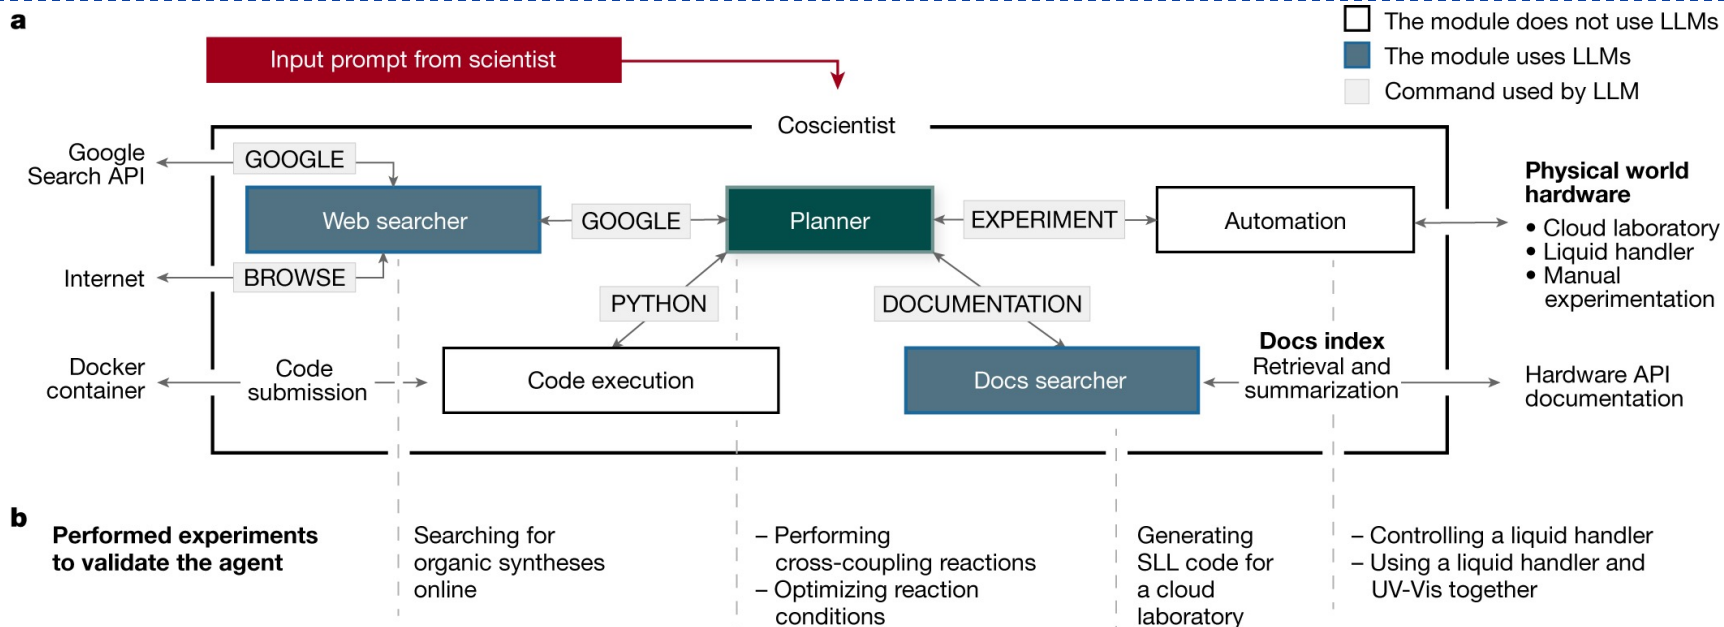


PhysicsMinions: Winning Gold Medals In The Latest Physics Olympiads With A Coevolutionary Multimodal Multi-Agent System, arXiv 2509

AI Scientist for Machine Learning

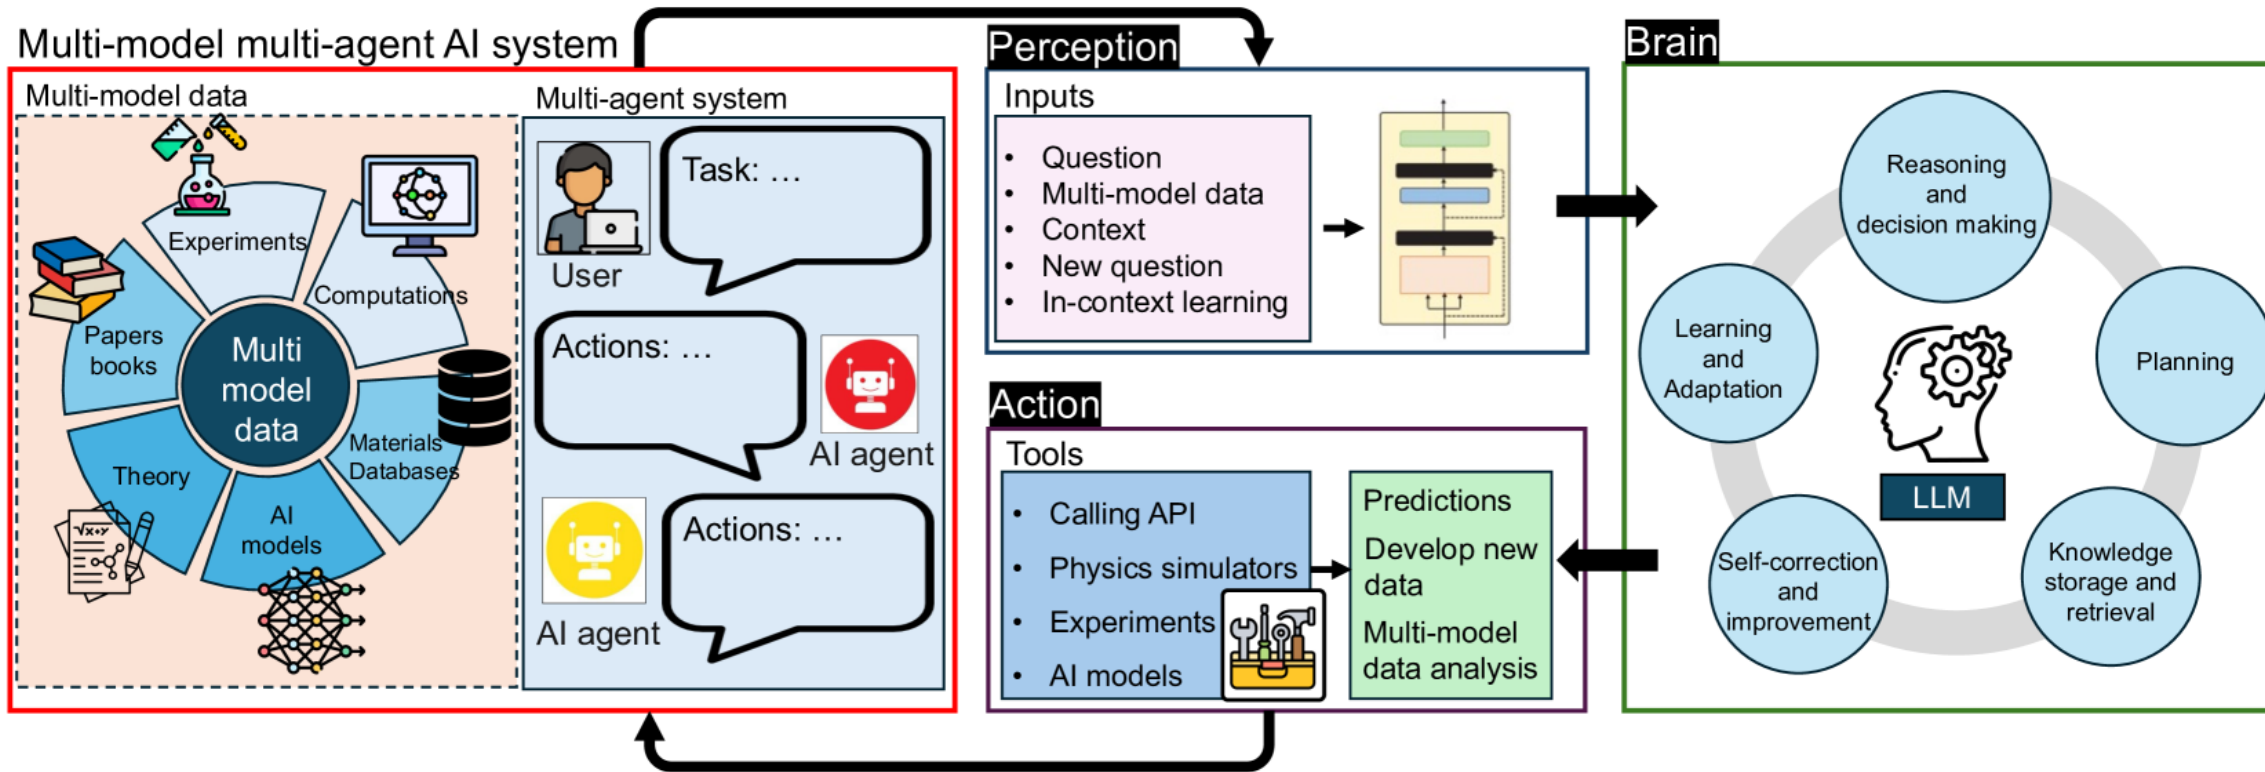


Coscientist



Autonomous chemical research with large language models, Nature 2024

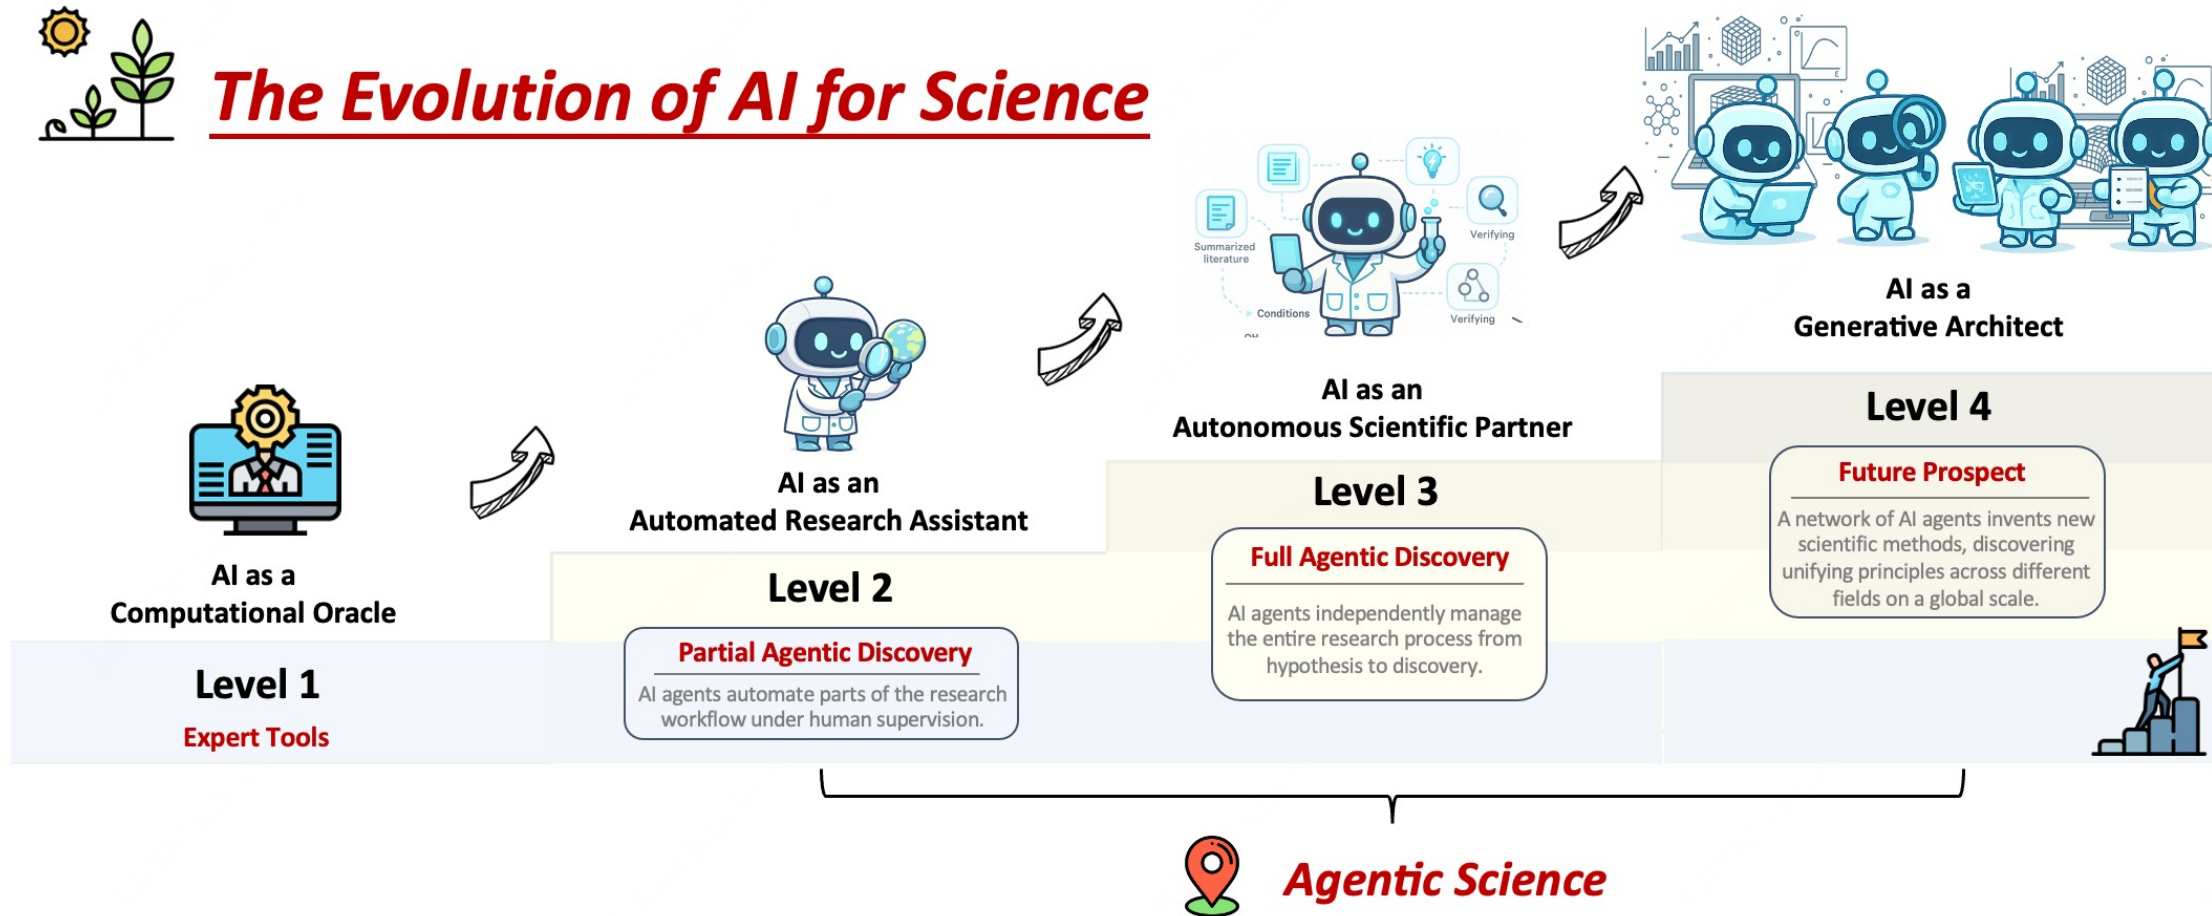
AtomAgents



Autonomous Scientific Discovery



The Evolution of AI for Science

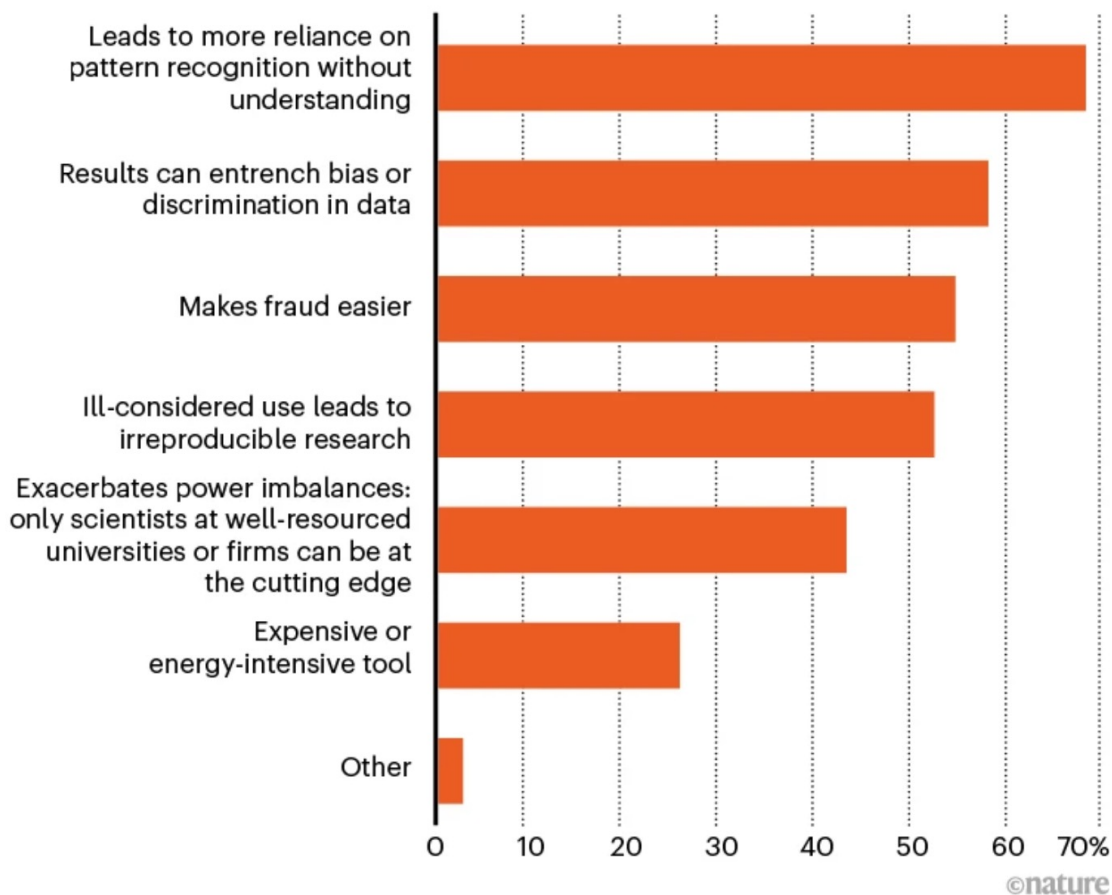


From AI for Science to Agentic Science: A Survey on Autonomous Scientific Discovery, arXiv:2508

The Other Side of The Coin

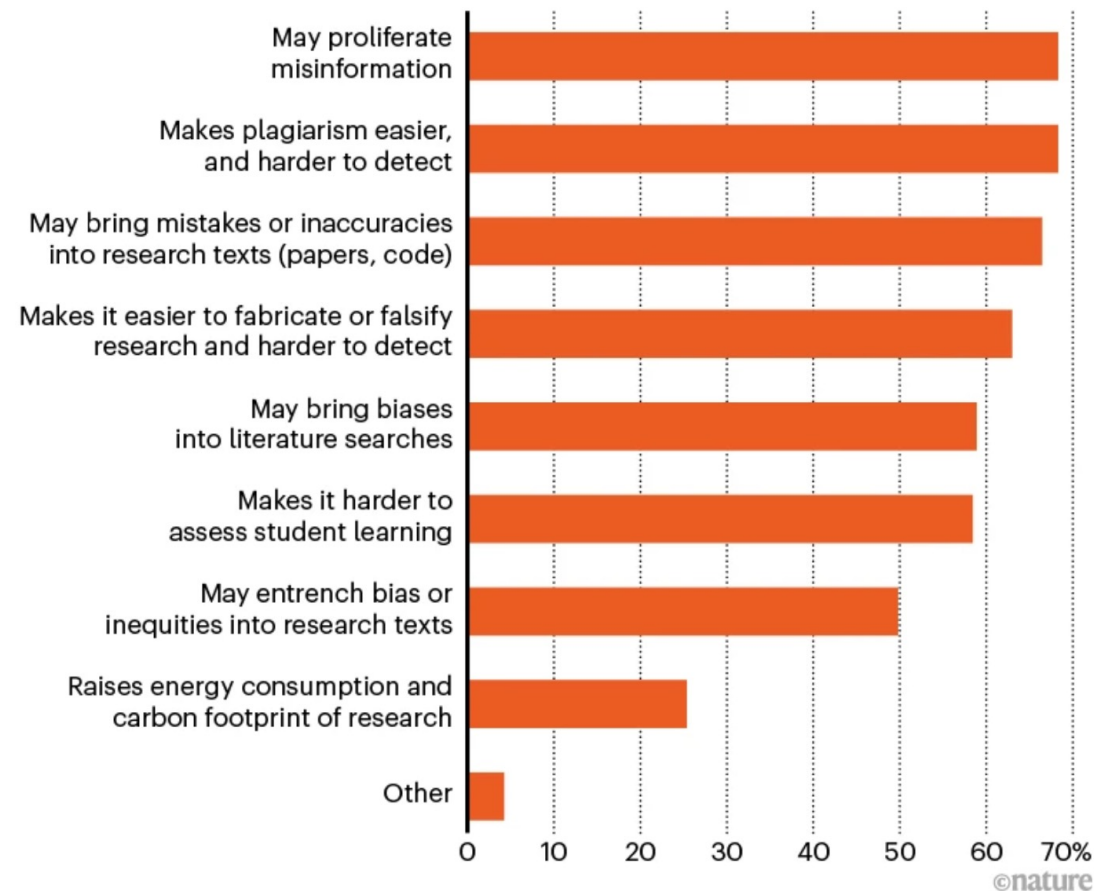
NEGATIVE IMPACTS OF AI

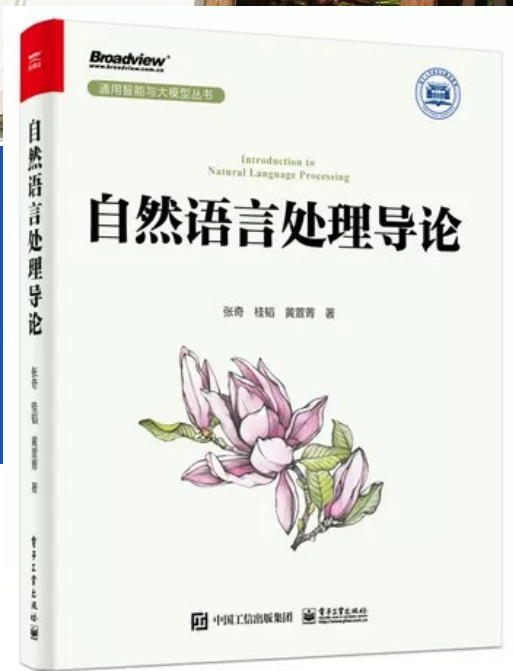
Q: Considering machine-learning methods, what do you think are negative impacts of AI in research? (Choose all that apply.)



PROBLEMS OF GENERATIVE AI

Q: Where do you think generative AI may have negative impacts on research? (Choose all that apply.)





Thanks

解码大语言模型奥妙 引领机器智能新时代

《大规模语言模型：从理论到实践》

张奇 桂韬 郑锐 黄萱菁 著

第2版



彩色
印刷

本书封底含有wispaper.ai科研助手智能体邀请码，凭邀请码可免费下载使用（预计2025年6月上线）。