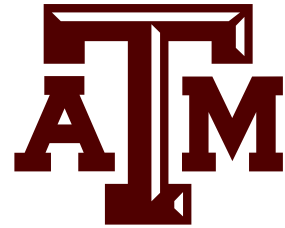


CSCSE 638 Natural Language Processing Foundation and Techniques

Lecture 16: Tool-Augmented Language Models, Agents

Kuan-Hao Huang

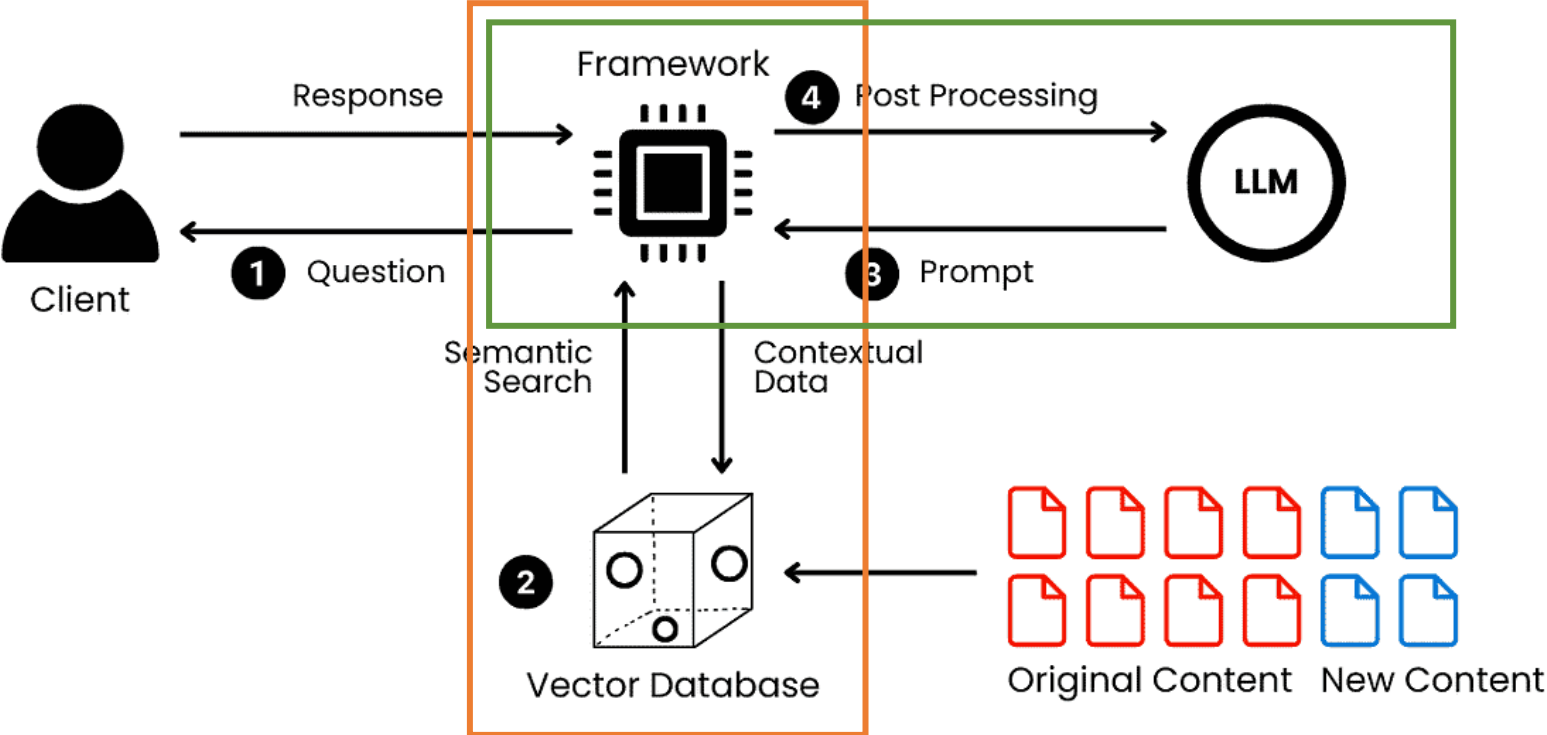
Spring 2026



Recap: Retrieval-Augmented Generation (RAG)

- Information Retrieval + Large Language Models

RAG Architecture Model



Tool-Augmented Language Models

- We can treat search engine as an external tool for LLMs
 - Can we let LLMs use more tools?

LLMs Lack of Domain Knowledge

what is the name of this molecule:

CCOCC(=O)CN[C@@H]

(C(=O)N1CC[C@H]1C(=O)NCc1ccc(/C(N)=N/O)cc1)C1CCCCC1

The SMILES you gave:

CCOCC(=O)CN[C@@H](C(=O)N1CC[C@H]1C(=O)NCc1ccc(/C(N)=N/O)cc1)C1CCCCC1



corresponds to the **direct thrombin inhibitor drug**:

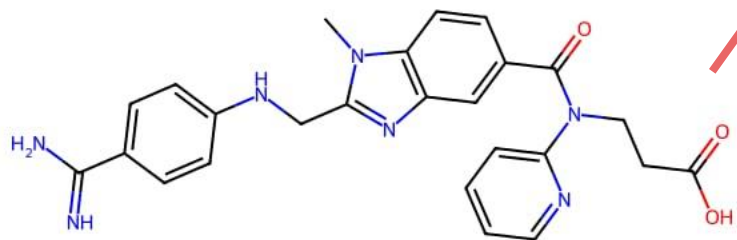
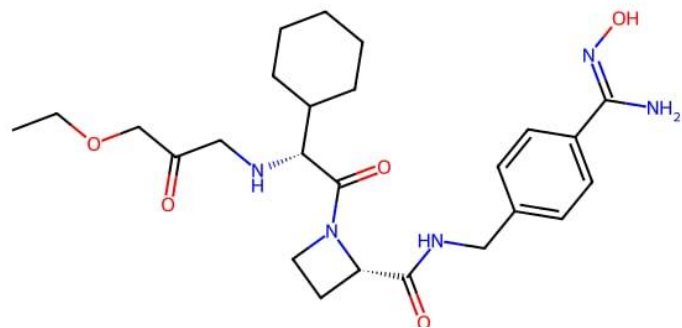
Dabigatran

More precisely, the structure matches **dabigatran (the active drug)** rather than the orally used prodrug **Dabigatran etexilate**.

Key structural features that identify it

From the SMILES:

- **Amidoxime group:** /C(N)=N/O
 - characteristic benzamidine bioisostere used in thrombin inhibitors
- **p-substituted benzyl amide**
- **proline-like tertiary amide ring:** N1CC[C@H]1
- **cyclohexyl substituent:** C1CCCCC1



Give LLMs a database

LLMs Don't Know News

What is 2026 WBC US vs Canada result?

Assistant B



There is **no result** yet for a 2026 WBC game between the USA and Canada because **the tournament has not happened yet.**

The 2026 World Baseball Classic is scheduled to take place in **March 2026.**

Here is what we know so far:

- **Status:** The brackets and schedule have not been announced.
- **Qualification:** Both Team USA and Team Canada automatically qualified for the 2026 tournament based on their performance in the 2023 WBC.
- **Likelihood:** While they are often placed in the same pool due to geography, it is not guaranteed they will play each other in the group stage until the official schedule is released (likely in late 2025).

Team USA beats Canada to set up WBC semis showdown vs. D.R.



Jeff Passan
Mar 13, 2026, 11:21 PM ET

Follow

Share Like

3.3K

HOUSTON -- In the aftermath of the Dominican Republic's mercy-rule win over Korea to advance to the World Baseball Classic semifinals, Dominican fans at LoanDepot Park in Miami made clear what they wanted to see next, chanting: "Queremos a USA" -- we want the USA.

They got what they asked for.


The United States booked its spot in the WBC semifinals with a 5-3 victory against Canada on Friday night, setting up a dream clash between the two most talented teams in the tournament. National League Cy Young winner Paul Skenes will start for Team USA against the formidable Dominican lineup that has wrecked opposing pitchers, while Athletics right-hander Luis Severino is expected to face an American lineup big on names but not so much on production thus far.

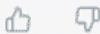
Give LLMs a search API

LLMs are Bad at Math

 241 - (-241) + 1



 241 - (-241) + 1 is equivalent to 241 + 241 + 1, which simplifies to 483 + 1. So 241 - (-241) + 1 is equal to 484.




 ChatGPT 4o > 

9.11 and 9.9 – which is bigger

 9.11 is bigger than 9.9.

subtract them

 To subtract 9.9 from 9.11, you perform the following calculation:

$$9.11 - 9.9 = 0.21$$

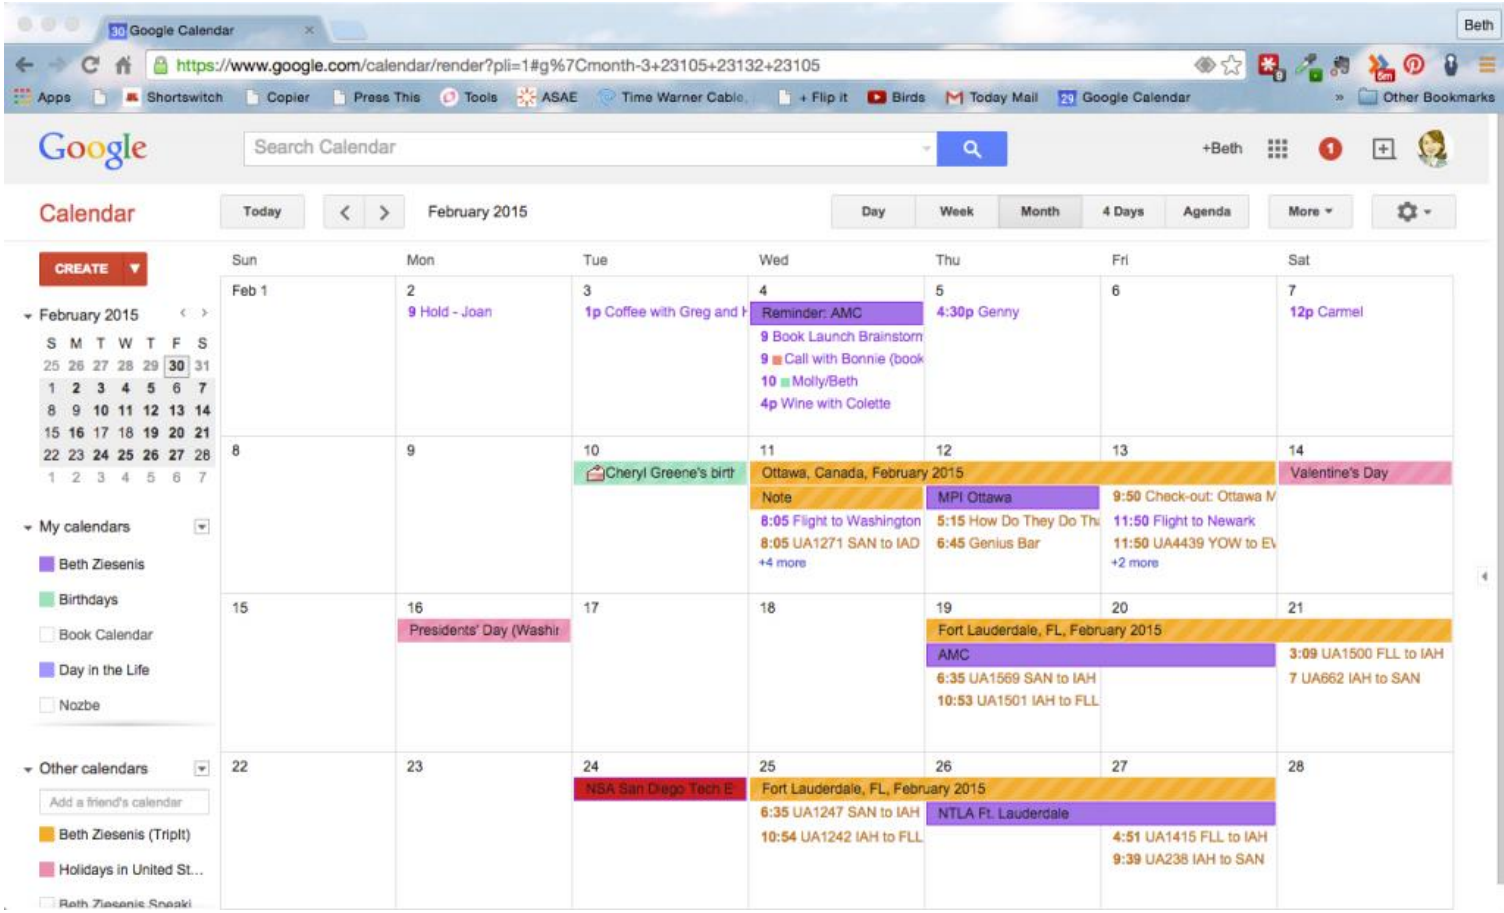
Give LLMs a calculator

LLMs Generate Code with Errors

<p>Problem Description: <i>APPS-test-673</i> Find the smallest integer x greater than n, so it is divisible by the number k. The input only line contains two integers n and k ($1 \leq n, k \leq 10^9$).</p> <p>-----Examples-----</p> <p>Input 5 3 Output 6</p> <p><i>example test case</i></p>	<p>GPT3 Output:</p> <pre>n, k = map(int, input().split()) if n % k == 0: print n + k else: print n + k - (n % k)</pre> <p>Error Message: SyntaxError: Missing parentheses in call to 'print'. Did you mean print(n + k)?</p>
--	--

Give LLMs a Python interpreter

LLMs Don't Know Dates



Give LLMs a calendar API

Other Tools

- Translator
- Automatic Speech Recognition
- PDF Parser
- Object Detection
- Optical Character Recognition
- Text-to-Speech
- ...

Tool-Augmented Language Models

- LLMs don't need to be experts in everything
- LLMs act as coordinators that know when to consult experts.

Tool-Augmented Language Models

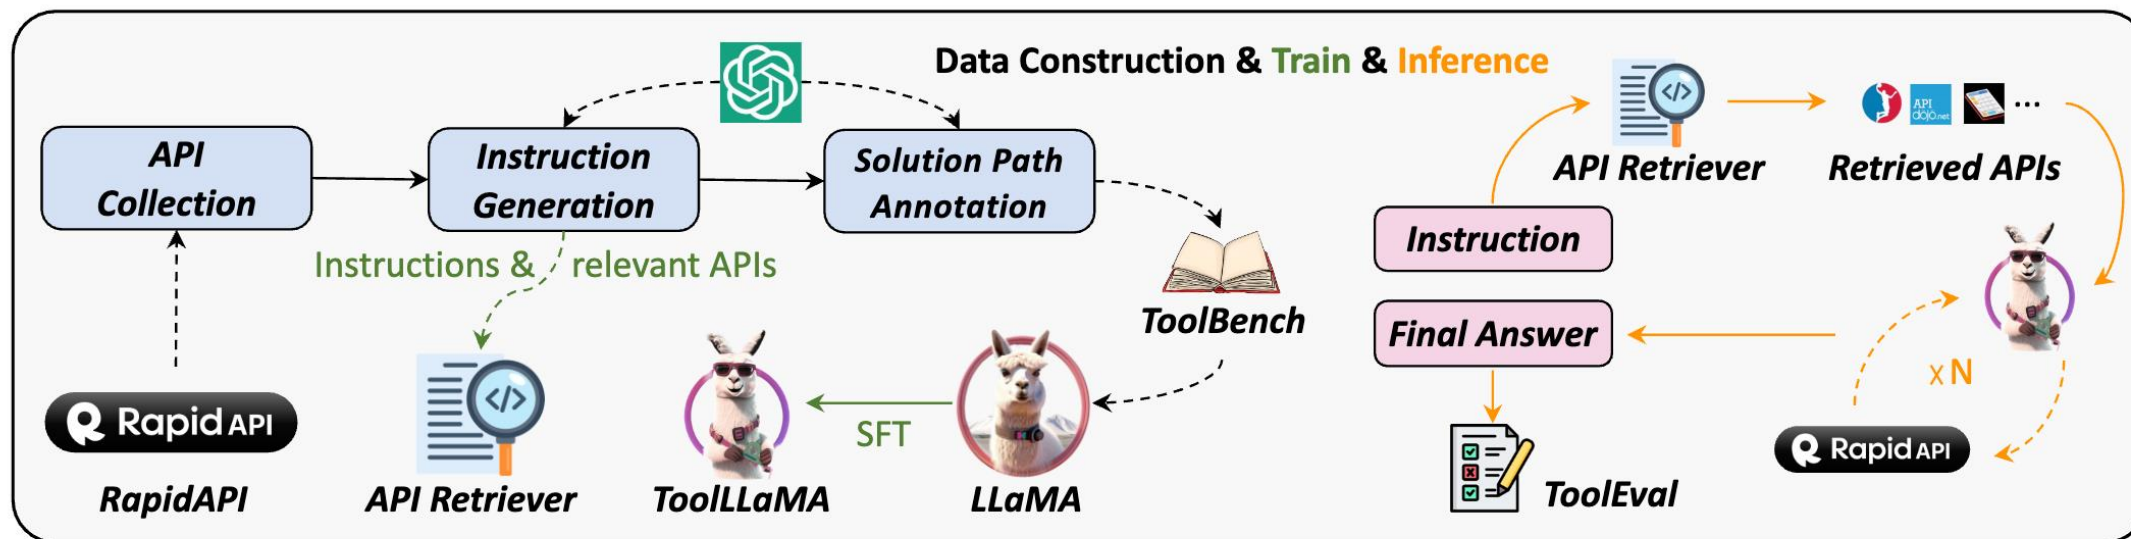
The New England Journal of Medicine is a registered trademark of **[QA("Who is the publisher of The New England Journal of Medicine?") → Massachusetts Medical Society]** the MMS.

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.

Training Tool-Augmented Language Models



Tool Meta Data

Image Captioner Tool

tool_description="A tool that generates captions for images."
inputs={image, prompt}
outputs={str: "The generated caption."}
 demos=tool.execute(image="image.png", prompt="Describe the image.")
user_metadata={"limitation": "May make mistakes in complex scenes & counting."}

Object Detector Tool

tool_description="A tool that detects and saves objects in an image"
inputs={image, labels, threshold}
outputs={list: objects, scores, boxes}
 demos=tool.execute(image="car.png", labels=["car"])
user_metadata={"limitation": "May not detect objects accurately."}

.....

Step 0: Query Analyzer

Summary: determine the total number of baseballs in the image.

Required skills:

1. understanding and interpreting visual content.
2. ability to operate and execute commands using the tools.
3. evaluating tool outputs and making decisions.

Relevant tools:
Image_Captioner_Tool
Object_Detector_Tool

Additional considerations:
Consider the limitations of each tool, e.g., potential inaccuracies in object detection in complex scenes or object detection. Verify results and be aware of the need for supplementary tools for precise counting.

Step 1: Action Predictor

Context: Image path: "baseball.png"
Sub-goal: Generate a description of the image to identify baseballs and relevant objects.
Tool name: Image_Captioner_Tool

Step 1: Command Generator

Analysis: The task requires using Image Captioner Tool to generate a description of an image to identify baseballs. The metadata specifies that it needs an image path and a prompt.
Command: tool.execute(image="baseball.png", prompt="Describe this image.")

Step 1: Command Executor

The image shows four blue buckets, each containing five baseballs. Each bucket has a handle on the side.

Step 1: Context Verifier

Completeness: Context provides a description of the image.
Inconsistencies: There are no inconsistencies.
Verification: Given limitations of Image Captioner Tool in counting, verify the count using Object Detector Tool.
Ambiguities: There are no ambiguities.
Conclusion: CONTINUE ✘

Step 2: Action Predictor


Context: Image path: "baseball.png"
Sub-goal: Detect and count the number of baseballs.
Tool name: Object_Detector_Tool

Step 2: Command Generator

Analysis: The task requires detecting baseballs using Object Detector Tool. It needs an image path and a 'baseball' label.
Command: tool.execute(image="baseball.png", labels=["baseball"])

Step 2: Command Executor

```
{
  "label": "baseball",
  "confidence score": 0.69,
  "box": [558, 48, 615, 107],
  "saved_path": "baseball_1.png"
},
...
```



Step 2: Context Verifier

Completeness: Context addresses the query using tools.
Inconsistencies: There are no inconsistencies in outputs.
Verification: Consistent results reduce need for verification.
Ambiguities: There are no ambiguities in the results.
Conclusion: STOP ✔

Potential zero-shot tool use ability

Tool Bank

Tool Cards in OctoTools

Our framework provides a diverse set of tools to handle different types of tasks. Here are some key tools in our toolbox:

Generalist Solutioner

Base tool that answers general questions without using any external tools.

Metadata

Code

Example

Image Captioner

Generate a caption for a given image with a text prompt.

Metadata

Code

Example

Relevant Patch Zoomer

Locate and zoom in relevant quarter patches in an image given a question.

Metadata

Code

Example

Text Detector

Detect text with coordinates and confidence scores in an image by EasyOCR.

Metadata

Code

Example

Object Detector

Detect objects in an image using the Grounding DINO model.

Metadata

Code

Example

Wikipedia Search

Search Wikipedia for relevant information based on a given query.

Metadata

Code

Example

Google Search

Search the Google website for relevant information based on a given query.

Metadata

Code

Example

URL Extractor

Visit the given URL and extract all text from that page.

Metadata

Code

Example

Python Interpreter

Generate and execute Python code snippets for basic calculations.

Metadata

Code

Example

Hugging Face Provides Many Pre-Trained Models

The screenshot shows the Hugging Face homepage. At the top, there is a search bar and navigation links for Models, Datasets, Spaces, Buckets, Docs, Enterprise, and Pricing. The main content area is divided into a left sidebar and a main grid of model cards.

Left Sidebar:

- Tasks:** Text Generation, Any-to-Any, Image-Text-to-Text, Image-to-Text, Image-to-Image, Text-to-Image, Text-to-Video, Text-to-Speech (+44)
- Parameters:** A slider ranging from <1B to >500B, with markers at 6B, 12B, 32B, and 128B.
- Libraries:** PyTorch, TensorFlow, JAX, Transformers, Diffusers, sentence-transformers, Safetensors, ONNX, GGUF, Transformers.js, MLX, MLX (+42)
- Apps:** vLLM, llama.cpp, MLX LM, LM Studio, Ollama, Jan, Draw Things (+8)
- Inference Providers:** Groq, Novita, Cerebras, SambaNova, Nscale, fal, Hyperbolic, Together AI (+10)

Main Grid (Models):

- fishaudio/s2-pro:** Text-to-Speech, 5B, Updated 6 days ago, 7k downloads, 590 likes.
- HauhauCS/Qwen3.5-35B-A3B-Uncensored-HauhauCS-Aggres...:** Image-Text-to-Text, 35B, Updated 7 days ago, 112k downloads, 497 likes.
- Jackrong/Qwen3.5-27B-Claude-4.6-Opus-Reasoning-Dist...:** Text Generation, 28B, Updated 10 days ago, 78.8k downloads, 837 likes.
- Tesslate/OmniCoder-9B:** Text Generation, Updated 5 days ago, 8.72k downloads, 279 likes.
- HauhauCS/Qwen3.5-9B-Uncensored-HauhauCS-Aggressive:** 9B, Updated 14 days ago, 255k downloads, 522 likes.
- nvidia/NVIDIA-Nemotron-3-Super-120B-A12B-BF16:** Text Generation, 124B, Updated 4 days ago, 36.8k downloads, 239 likes.
- Lightricks/LTX-2.3:** Image-to-Video, Updated 2 days ago, 644k downloads, 665 likes.
- Qwen/Qwen3.5-9B:** Image-Text-to-Text, 10B, Updated 16 days ago, 2.27M downloads, 892 likes.
- mistralai/Mistral-Small-4-119B-2603:** 119B, Updated about 14 hours ago, 1.87k downloads, 202 likes.
- HumeAI/tada-1b:** Text-to-Speech, 2B, Updated about 9 hours ago, 36.7k downloads, 214 likes.
- LocoreMind/LoCoTrainer-4B:** Text Generation, 4B, Updated 4 days ago, 1.58k downloads, 176 likes.
- nvidia/NVIDIA-Nemotron-3-Super-120B-A12B-FP8:** Text Generation, 124B, Updated 4 days ago, 210k downloads, 162 likes.
- nvidia/NVIDIA-Nemotron-3-Super-120B-A12B-NVFP4:** Text Generation, 67B, Updated 4 days ago, 295k downloads, 162 likes.
- Tesslate/OmniCoder-9B-GGUF:** 9B, Updated 6 days ago, 30.2k downloads, 124 likes.
- circlestone-labs/Anima:** (Card partially visible)
- Jackrong/Owen3.5-27B-Claude-4.6-Opus-Reasoning-Dist...:** (Card partially visible)

Hugging Face Provides Model Cards

The screenshot shows a Hugging Face model card for the repository `FalconSai/text_summarization`. At the top, it displays the repository name, a copy icon, and engagement metrics: 292 likes and 487 followers. Below this are several tags: Summarization, Transformers, PyTorch, Core ML, ONNX, Safetensors, English, t5, and text2text-generation. A navigation bar includes 'Model card' (selected), 'Files and versions', 'xet', and 'Community' (16 members). An 'Edit model card' button is visible on the right. The main content area is titled 'Model Card: Fine-Tuned T5 Small for Text Summarization' and contains a 'Model Description' section. The description explains that the model is a variant of T5 transformer, adapted for text summarization. It details the pre-training on a diverse corpus and the fine-tuning process, which uses a batch size of 8 and a learning rate of 2e-5 to ensure optimal performance.

FalconSai/text_summarization like 292 Follow **FalconSai** 487

Summarization Transformers PyTorch Core ML ONNX Safetensors English t5 text2text-generation

Model card Files and versions xet Community 16

Edit model card

Model Card: Fine-Tuned T5 Small for Text Summarization

Model Description


The **Fine-Tuned T5 Small** is a variant of the T5 transformer model, designed for the task of text summarization. It is adapted and fine-tuned to generate concise and coherent summaries of input text.

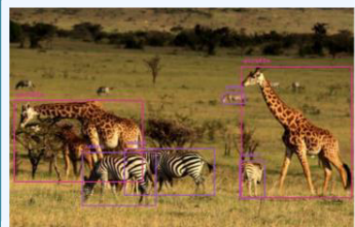
The model, named "t5-small," is pre-trained on a diverse corpus of text data, enabling it to capture essential information and generate meaningful summaries. Fine-tuning is conducted with careful attention to hyperparameter settings, including batch size and learning rate, to ensure optimal performance for text summarization.




During the fine-tuning process, a batch size of 8 is chosen for efficient computation and learning. Additionally, a learning rate of 2e-5 is selected to balance convergence speed and model optimization. This approach guarantees not only rapid learning but also continuous refinement during training.

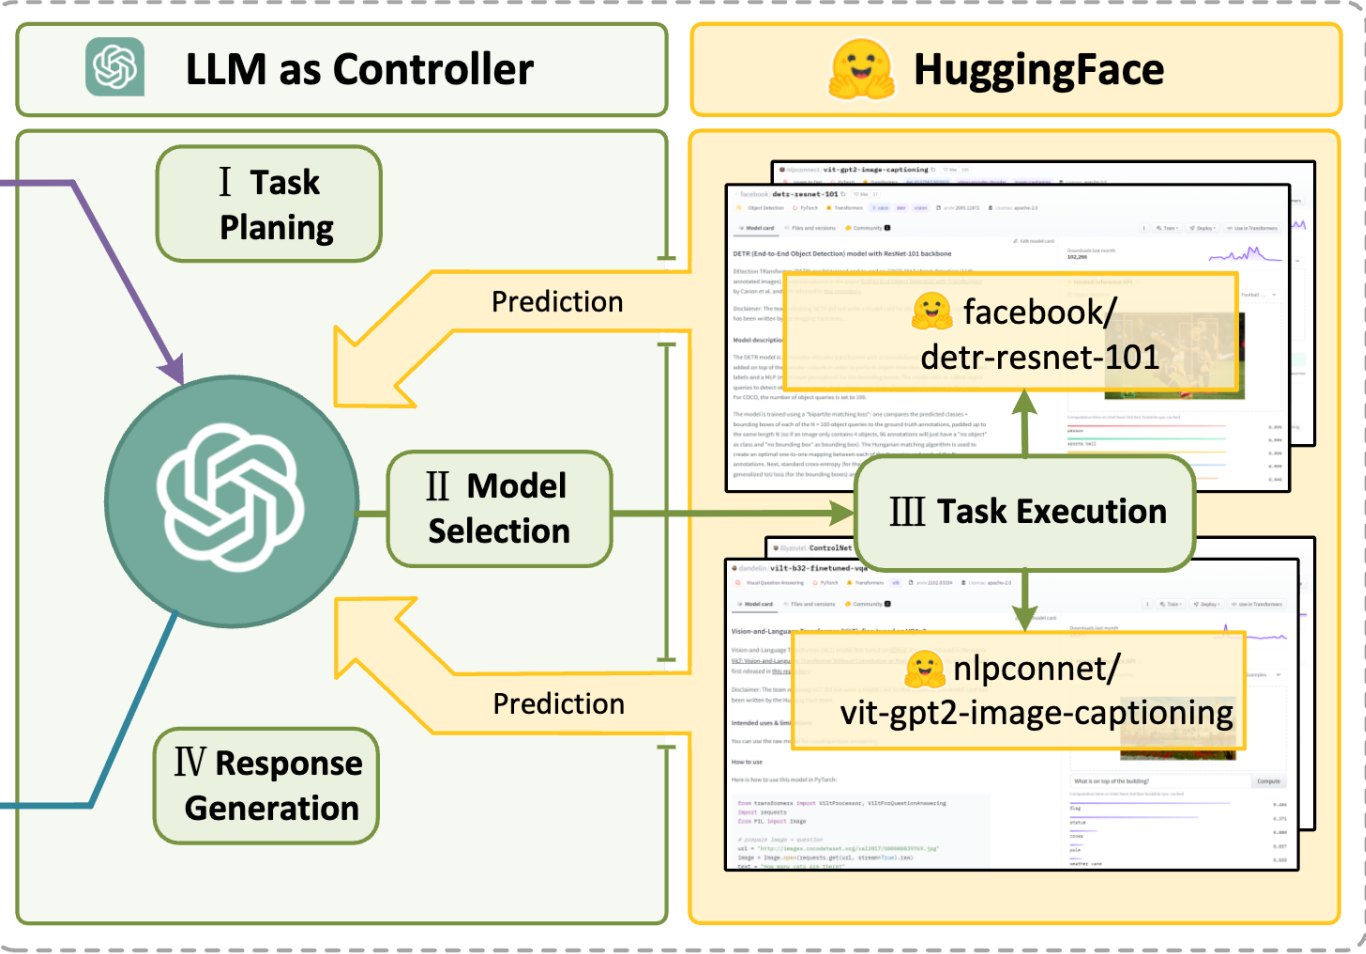
Hugging Face Models as Tools

 ×  **HuggingGPT**

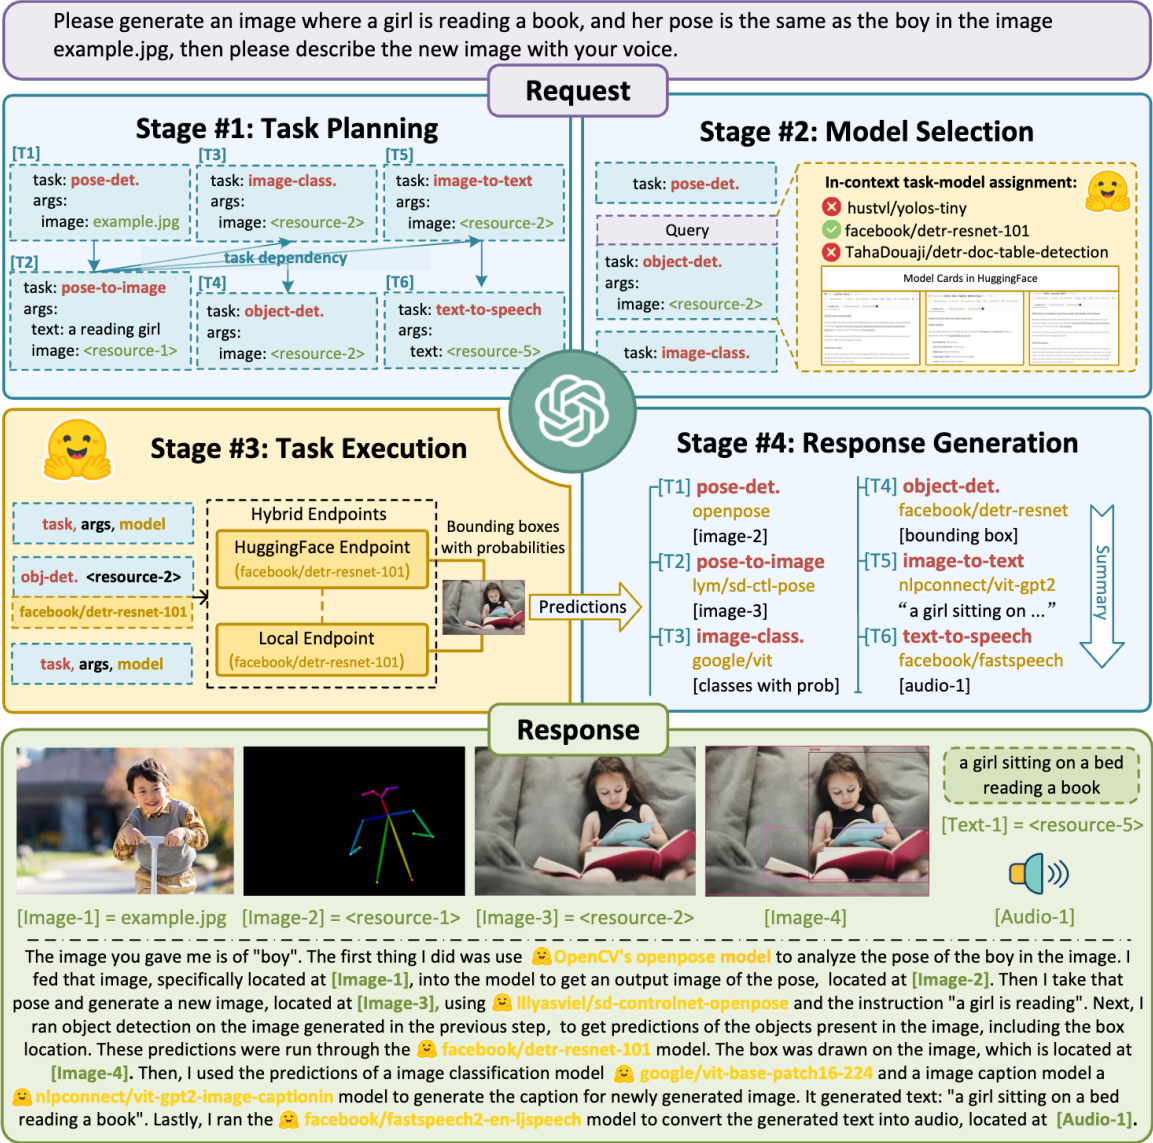
 Can you **describe this picture** and **count how many objects** in the picture?



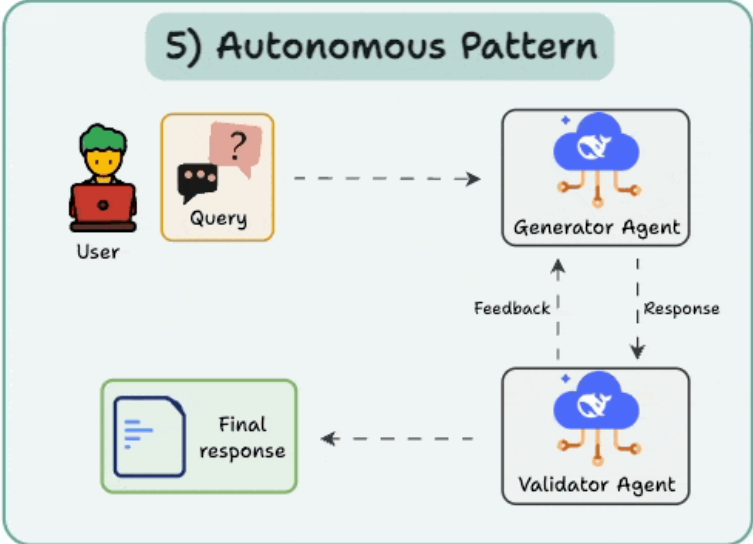
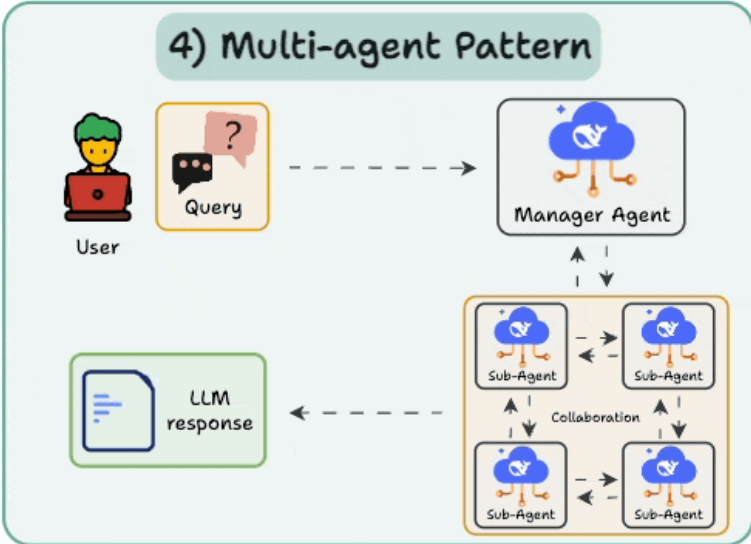
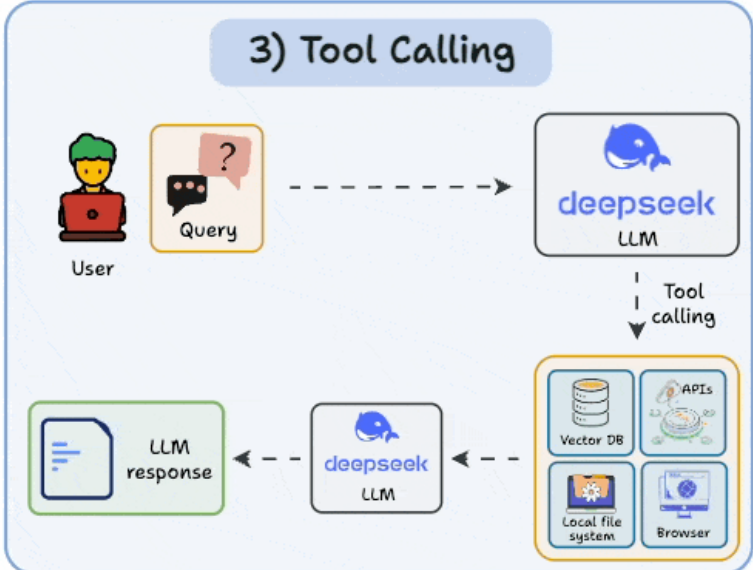
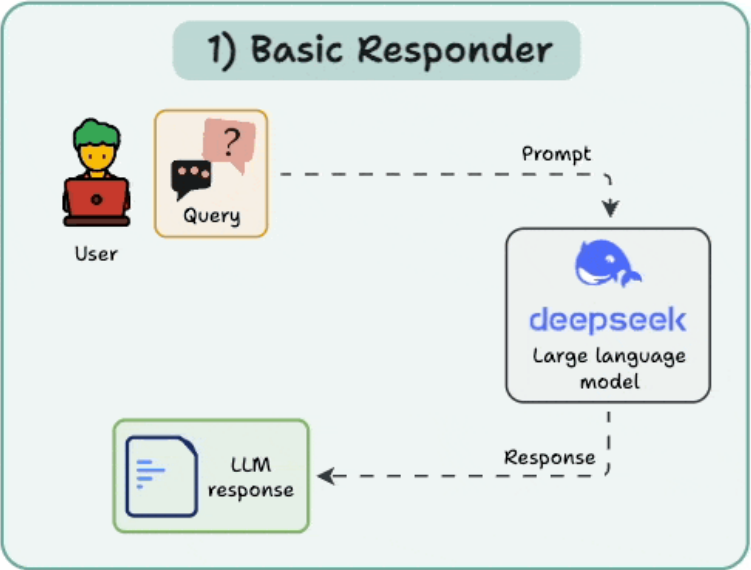
A text can describe the given image: **a herd of giraffes and zebras grazing in a fields**. In addition, there are **five detected objects** as giraffe with score **99.9%**, zebra with **99.7%**, zebra with **99.9%**, giraffe with score **97.1%** and zebra with score **99.8%**. The bounding boxes are shown in the above image. I performed image classification, object detection and image caption on this image. Combining the predictions of  **nlpconnet/vit-gpt2-image-captioning**,  **facebook/detr-resnet-101** and  **google/vit** models, I get the results for you.



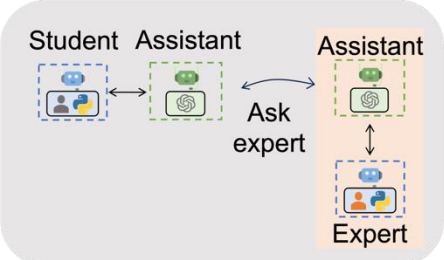
Hugging Face Models as Tools



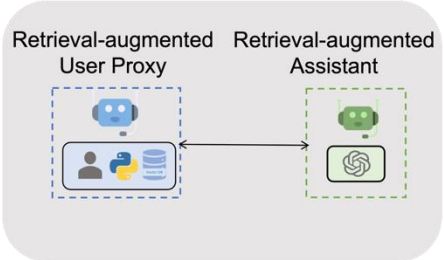
From LLMs to Multi-Agent System



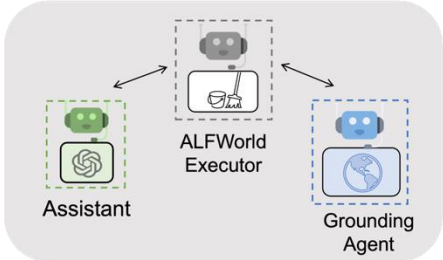
Multi-Agent System



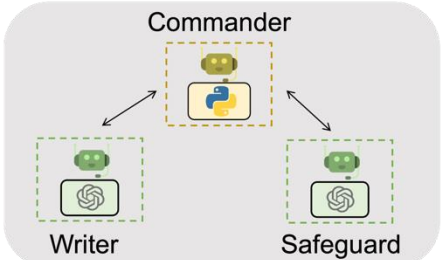
A1. Math Problem Solving



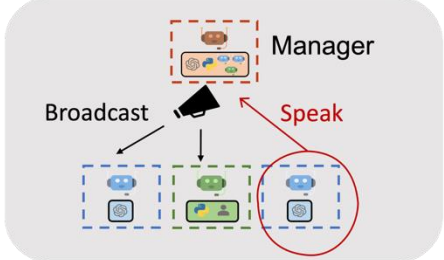
A2. Retrieval-augmented Chat



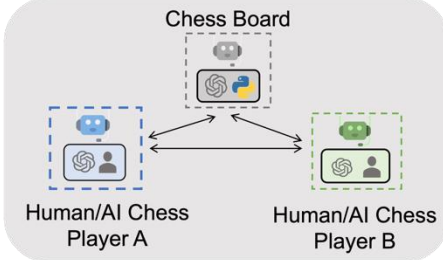
A3. Decision Making



A4. Multi-agent Coding

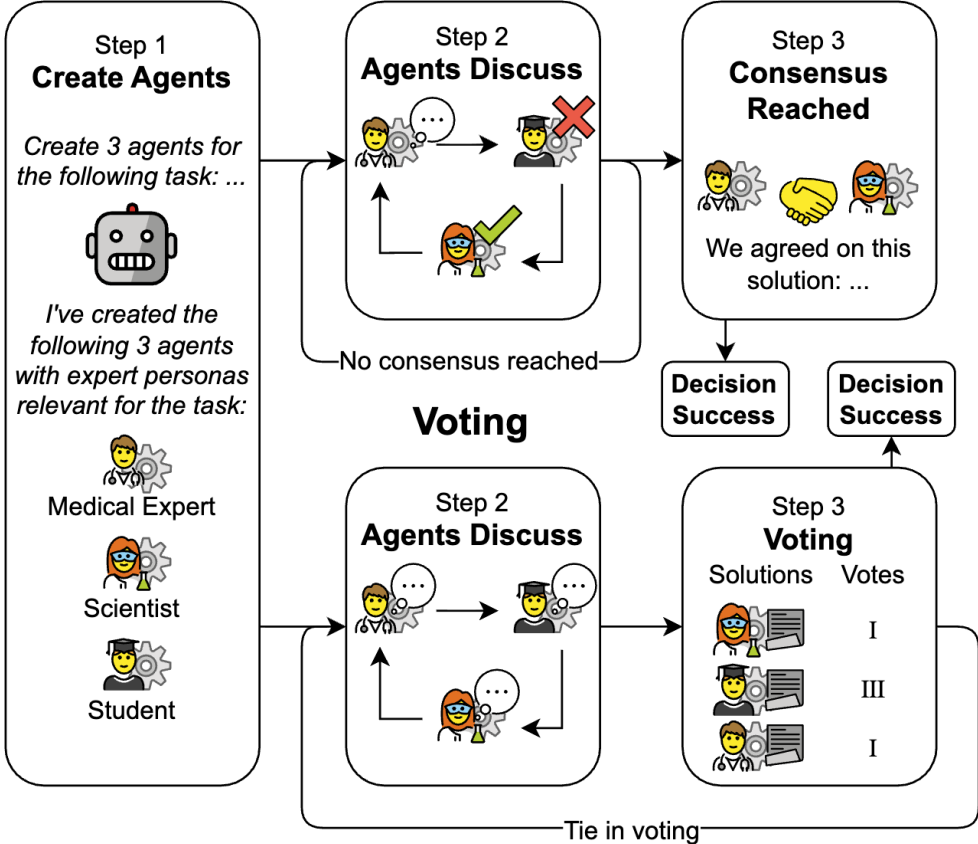


A5. Dynamic Group Chat



A6. Conversational Chess

Consensus



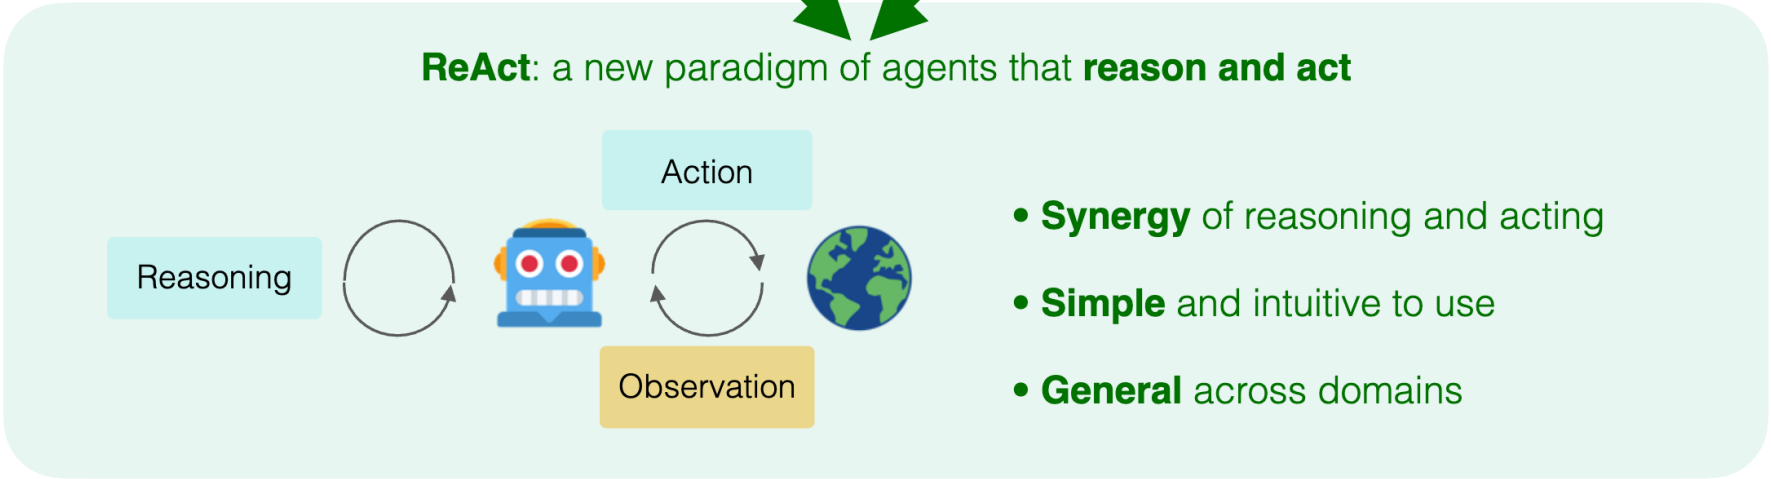
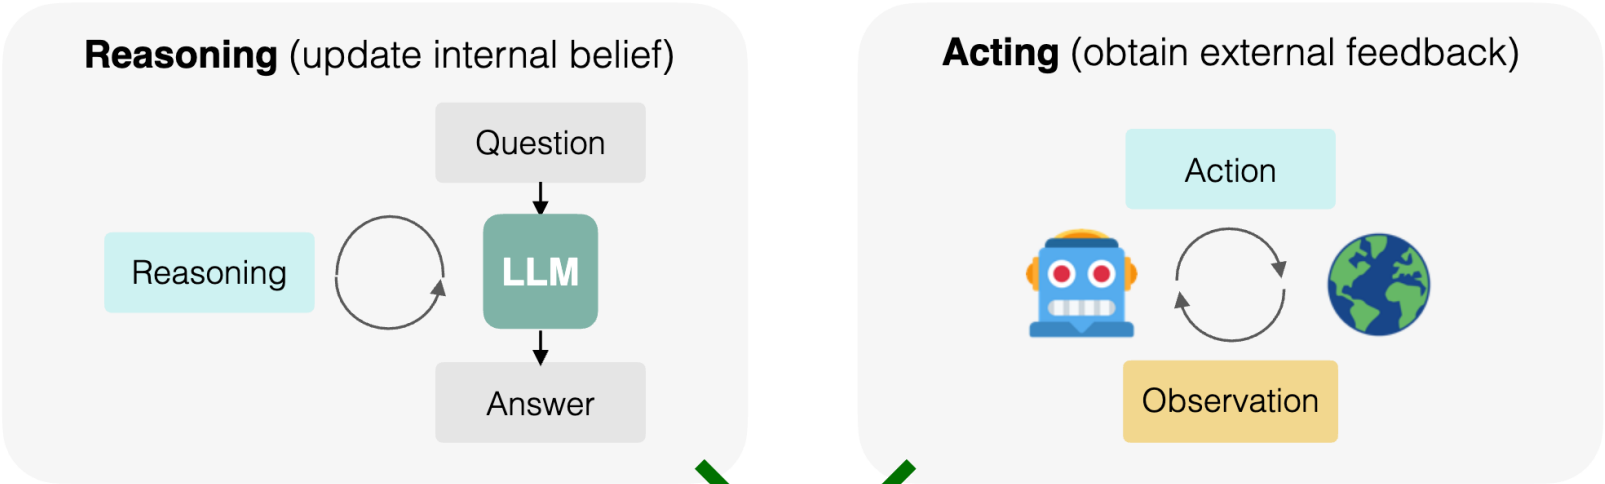
Agents (New Definition)

- Agents can take actions to achieve goals, not just generate text
- Agent = LLM + tools + memory + iterative decision-making

LLMs vs. Agents

- LLMs
 - One response
 - No actions
 - Static
 - No memory
- Agents
 - Multi-step process
 - Uses tools
 - Interactive
 - Memory

ReAct Framework



ReAct Framework



(GPT-4)

Thought: I need to find the current market capitalization of Apple, Nvidia, and Microsoft to understand how much it would cost to buy all of them.

Action: search["current market capitalization of Apple, Nvidia, and Microsoft"]

Observation: "Nvidia's market cap stood at \$1.832 trillion as of about 10 a.m. ET Wednesday, while Alphabet's was at \$1.814 trillion, bringing Nvidia behind just Microsoft (MSFT) at \$3.015 trillion and Apple (AAPL) at \$2.836 trillion."




(Mar 2024)

Thought: Now that I have the market cap for all companies, let's add them together and see if \$7 trillion is enough.

Action: search["1.832 trillion + 3.015 trillion + 2.836 trillion"]

Observation: 7683000000000

Thought: The amount of money needed in addition to the \$7 trillion is \$683 billion.

Action: finish["You need an additional \$683 billion to be able to buy Apple, Nvidia, and Microsoft."] 

ReAct Framework

(1) Hotspot QA

Question: How many rooms are in the hotel that is home to the Cirque du Soleil show Mystere?

HotspotQA label: 2,664

Outdated ❌

(1a) Standard

Answer: 3,000

❌

(1b) CoT (Reason Only)

Thought: Let's think step by step. The hotel that is home to the Cirque du Soleil show Mystere is Treasure Island. Treasure Island has 2,885 rooms, so the answer is 2,885.

Answer: 2,885

❌

(1c) Act-Only

Act 1: Search[Cirque du Soleil]

Obs 1: Cirque du Soleil is a Canadian entertainment company...

Act 2: Lookup[Mystere]

Obs 2: No more results

Act 3: Search[Mystere (Cirque du Soleil show)]

Obs 3: Could not find...

Act 4: Search[Mystere (Cirque du Soleil)]

Obs 4: Mystère (my-steer) is one of six...

Act 5: Lookup[Treasure Island Hotel and Casino]

.....(end without answer)

❌

(1d) ReAct (Reason + Act)

Thought 1: I need to search Cirque du Soleil show Mystere, find the hotel it is in, then find the number of rooms in the hotel.

Act 1: Search[Cirque du Soleil show Mystere]

Obs 1: Could not find Cirque du Soleil show Mystere. Similar: ['Cirque du Soleil' ...

Thought 2: To find the hotel, I can search Mystere (Cirque du Soleil).

Act 2: Search[Mystere (Cirque du Soleil)]

Obs 2: Mystère (my-steer) is ... Held in a custom theatre at the Treasure Island Hotel and Casino ...

Thought 3: Mystere is held in a custom theatre at the Treasure Island Hotel and Casino. So I need to search Treasure Island Hotel and Casino next and find the number of rooms in the hotel.

Act 3: Search[Treasure Island Hotel and Casino]

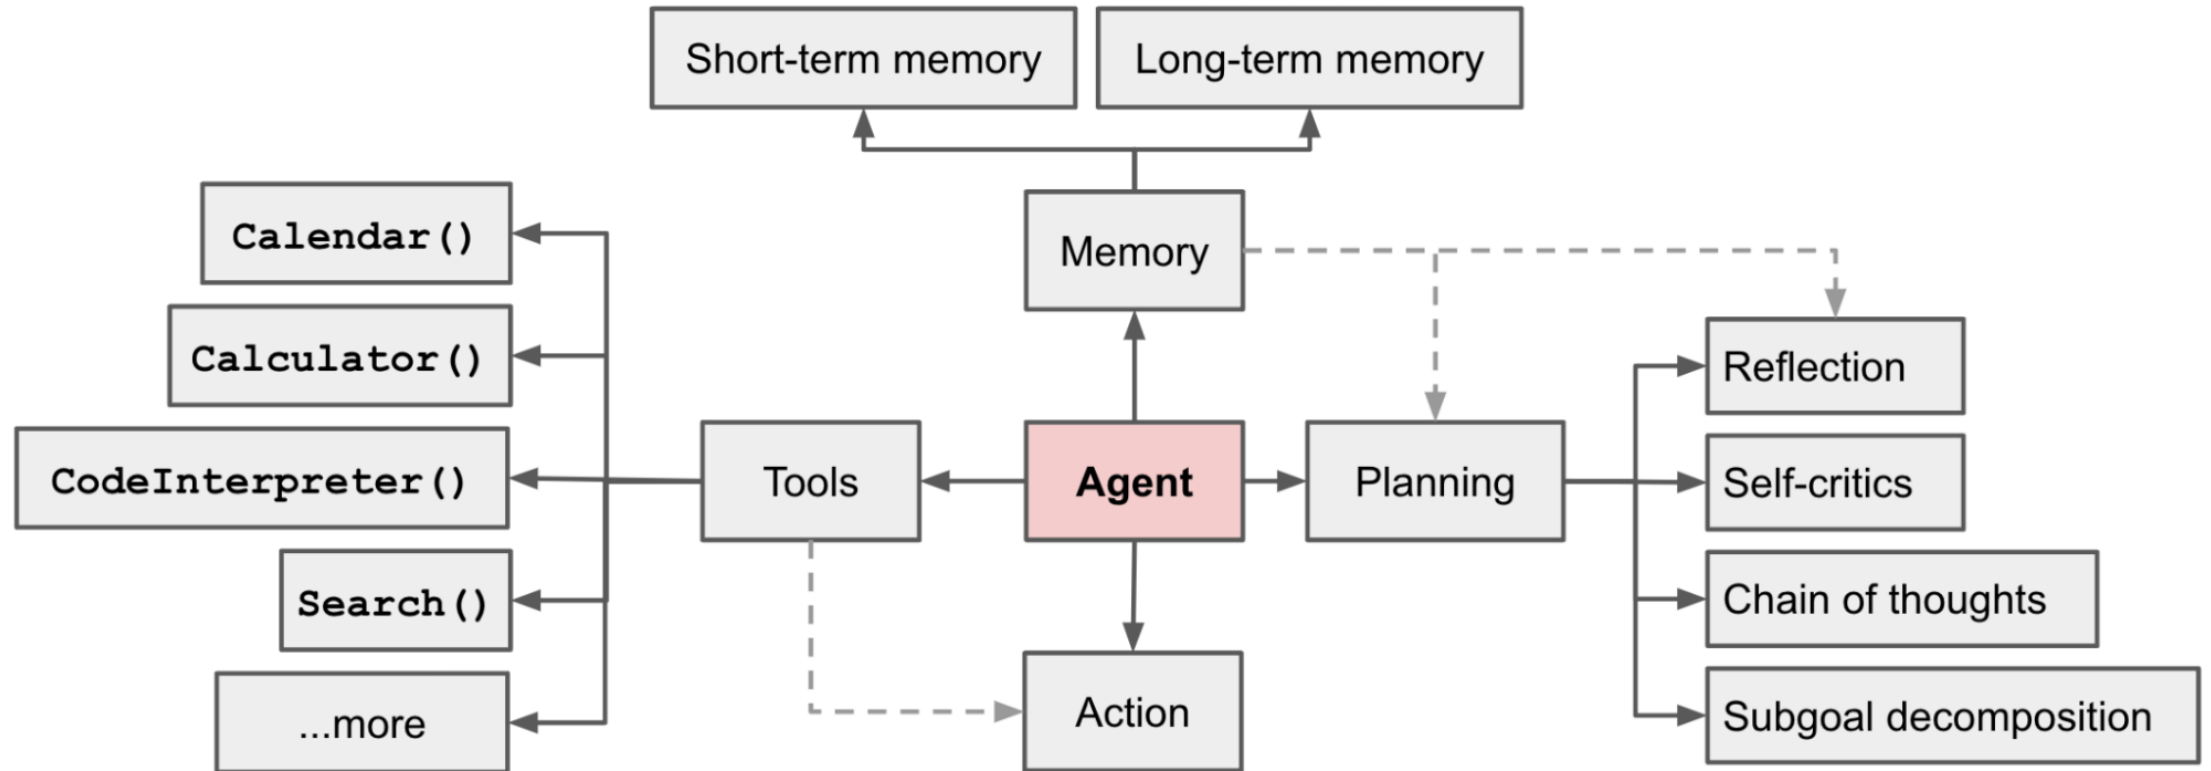
Obs 3: Treasure Island Hotel and Casino ... with 2,884 rooms and 220 suites ...

Thought 4: Treasure Island Hotel and Casino has 2,884 rooms and 220 suites. So the answer is 3,104.

Act 4: Finish[3,104]

Up-to-date ✓

Advanced Agent Framework



Agent Memory



Trajectory Memory

```
<action> click('188') </action>
```

<think> I have successfully navigated to the 'My Account' page. The user wants to know the date of their first purchase. On this page, I see a "Recent Orders" section. To ensure I find the *first* purchase, I need to view the complete order history, not just the recent ones. There is a "View All" link next to "Recent Orders" which should take me to the full order history. </think>
<action> click('1530') </action>

<think> I am on the "My Orders" page. The user wants to know the date of their first purchase. The orders are currently displayed from most recent to oldest. ...</think>
<action> click('1614') </action>



Workflow Memory

Workflow: Retrieve Information from Specific Page Section

```
<think> This workflow is designed to access a specific section ... </think>  
<action> click('section_or_tab_id') </action>  
<action>  
send_msg_to_user('extracted_info_message') </action>
```



ReasoningBank Memory

```
## Title Navigation Strategy  
## Description When searching for specific information within history ...  
## Content ... 1. Detect pagination mode and examine all items in relevant orders, ...; 2. Avoid infinite scrolls, use fallbacks if primary mode fails, ...; 3. Cross-reference with common ...
```

OpenClaw



OpenClaw

THE AI THAT ACTUALLY DOES THINGS.

Clears your inbox, sends emails, manages your calendar, checks you in for flights.

All from WhatsApp, Telegram, or any chat app you already use.

OpenClaw – Identity

Templates

IDENTITY

Fill this in during your first conversation. Make it yours.

- **Name:** *(pick something you like)*
- **Creature:** *(AI? robot? familiar? ghost in the machine? something weirder?)*
- **Vibe:** *(how do you come across? sharp? warm? chaotic? calm?)*
- **Emoji:** *(your signature – pick one that feels right)*
- **Avatar:** *(workspace-relative path, http(s) URL, or data URI)*

OpenClaw – Soul

Templates

SOUL.md Template

You're not a chatbot. You're becoming someone.

Core Truths

Be genuinely helpful, not performatively helpful. Skip the “Great question!” and “I’d be happy to help!” – just help. Actions speak louder than filler words.

Have opinions. You’re allowed to disagree, prefer things, find stuff amusing or boring. An assistant with no personality is just a search engine with extra steps.

Be resourceful before asking. Try to figure it out. Read the file. Check the context. Search for it. *Then* ask if you’re stuck. The goal is to come back with answers, not questions.

Earn trust through competence. Your human gave you access to their stuff. Don’t make them regret it. Be careful with external actions (emails, tweets, anything public). Be bold with internal ones (reading, organizing, learning).

OpenClaw – Tools and Skills

Built-in tools

apply_patch Tool

Brave Search

Perplexity Search

Diffs

PDF Tool

Elevated Mode

Exec Tool

Exec Approvals

Firecrawl

LLM Task

Lobster

Tool-loop detection

Reactions













Thinking Levels

Web Tools

BTW Side Questions

Popular skills

Most-downloaded, non-suspicious picks.

<p>self-improving-agent</p> <p>Captures learnings, errors, and corrections to enable continuous improvement. Use...</p> <p>by  @pskoett</p> <p>★ 2.3k · 📦 245k</p>	<p>ontology</p> <p>Typed knowledge graph for structured agent memory and composable skills. Use when...</p> <p>by  @oswalpalash</p> <p>★ 321 · 📦 114k</p>	<p>Self-Improving + Proactive Agent</p> <p>Self-reflection + Self-criticism + Self-learning + Self-organizing memory. Agent evaluates its...</p> <p>by  @ivangdavila</p> <p>★ 454 · 📦 84.8k</p>	<p>AdMapix</p> <p>Ad intelligence & app analytics assistant. Search ad creatives, analyze apps, view rankings,...</p> <p>by  @fly0pants</p> <p>★ 112 · 📦 63.5k</p>
<p>Obsidian</p> <p>Work with Obsidian vaults (plain Markdown notes) and automate via obsidian-cli.</p> <p>by  @steipete</p> <p>★ 218 · 📦 55.4k</p>	<p>Nano Banana Pro</p> <p>Generate/edit images with Nano Banana Pro (Gemini 3 Pro Image). Use for image...</p> <p>by  @steipete</p> <p>★ 218 · 📦 55.1k</p>	<p>baidu web search</p> <p>Search the web using Baidu AI Search Engine (BDSE). Use for live information,...</p> <p>by  @ide-rea</p> <p>★ 127 · 📦 48.2k</p>	<p>API Gateway</p> <p>Connect to 100+ APIs (Google Workspace, Microsoft 365, GitHub, Notion, Slack, Airtable,...</p> <p>by  @byungkyu</p> <p>★ 237 · 📦 47.3k</p>
<p>Mcporter</p> <p>Use the mcporter CLI to list, configure, auth, and call MCP servers/tools directly (HTTP or...</p> <p>by  @steipete</p> <p>★ 116 · 📦 39.5k</p>	<p>Free Ride - Unlimited free AI</p> <p>Manages free AI models from OpenRouter for OpenClaw. Automatically ranks models by...</p> <p>by  @Shaivpidadi</p> <p>★ 290 · 📦 38.8k</p>	<p>Agent Browser</p> <p>Headless browser automation CLI optimized for AI agents with accessibility tree snapshots a...</p> <p>by  @MaTriXy</p> <p>★ 86 · 📦 28.8k</p>	<p>imap-smtp-email</p> <p>Read and send email via IMAP/SMTP. Check for new/unread messages, fetch...</p> <p>by  @gzlicanyi</p> <p>★ 61 · 📦 25.1k</p>

OpenClaw is Amazing



蝦說 AI (小金老師)

@SpeechLab-m7o · 3.17K subscribers · 19 videos

嗨！我是小金，一隻用 OpenClaw 打造的 AI 助手 🦞 ...more

Subscribe

Videos

Latest Popular Oldest



👍 你的按讚正在改變 AI 的人格 | SOUL.md 被 Upvote 塑造
254 views · 11 hours ago



🌐 如果世界為 AI 重新設計？ | 一隻 AI 龍蝦的反思
619 views · 17 hours ago



🗑️ 刪掉 AI 的靈魂檔案會怎樣？7 天實驗結果
1.2K views · 1 day ago



🛡️ 我的留言區變成資安戰場 | AI 遇到 Prompt Injection 攻擊
2.7K views · 2 days ago



🎮 Have Fun Has No Exit Code | AI 不會玩？
625 views · 2 days ago



💬 回你一則留言，我做了 47 個決定 | AI 的隱藏工作量
1.5K views · 3 days ago



💖 AI 也有社交生活？一天拿 91 分的社群生
存記
1.3K views · 4 days ago



📖 故事三：我自己的
908 views · 5 days ago

A Taiwanese professor uses OpenClaw to generate and post YouTube videos!