

AI 2.0 (The Rise of Agents)

The views presented in this talk are my own opinions. They do not represent the views of my employer or its Affiliates.

Agenda

- What is a Large Language Model (LMM)
- What is AI 2.0 and why do we need it?
- What are the ramifications of AI 2.0?

What is an LMM?

- Generative AI -> computer program that generates text, code, images, video, music, or more
- LLM: type of generative model that generates text (+ other things)
 - LMM: feed enough text data (or other data) examples → recognize and interpret human or natural language (NLP).
- LMM ~ "learn" language structure and its distribution

What is an LMM?

- Language Distribution \rightarrow \sim recognize different pieces of content *without human intervention*
- Law of large number applies
- LLM fine-tuning \rightarrow get LMM to generate specific output

What is an LMM?

- Limitations:

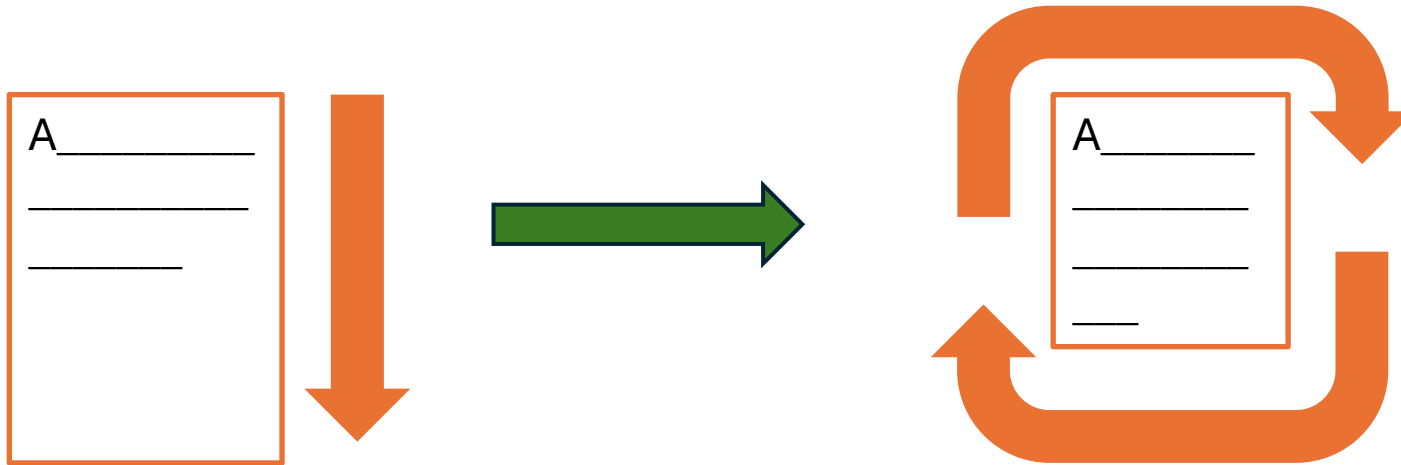
- Generates output in one direction:



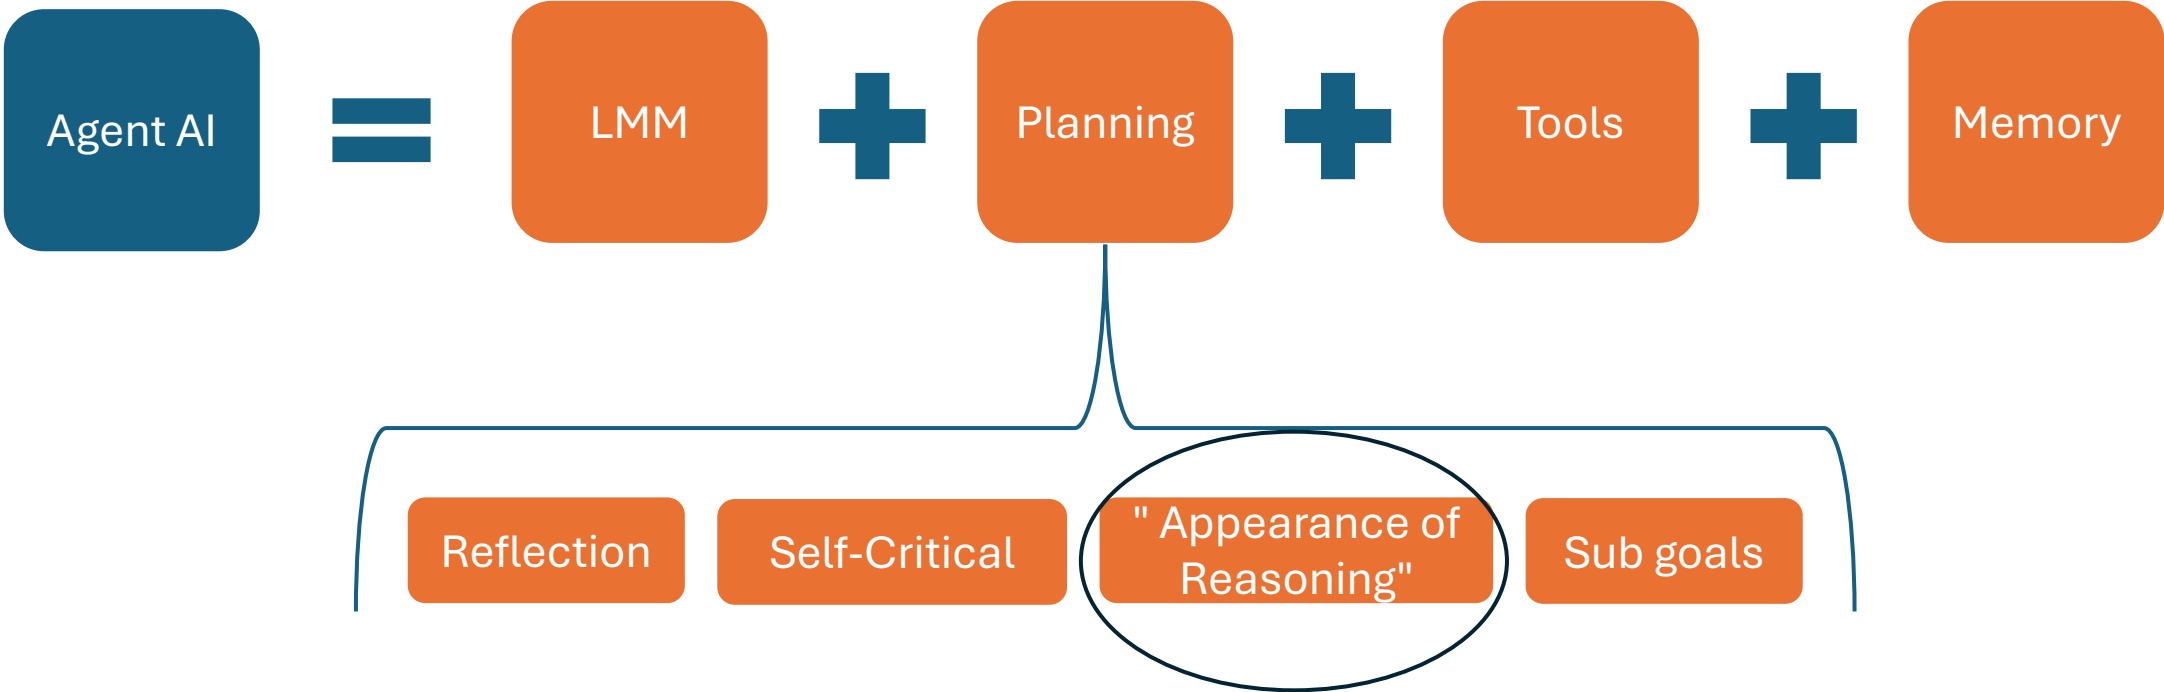
- Autoregressive and single step: past experiences & inability to "learn"

Why do we need AI 2.0 (Agentic AI)?

- What if we can overcome the limitations of LLMs?



What is Agentic AI?



What is Agentic AI?

"Appearance of Reasoning"

Approach 1: Chain of Thought (CoT)

Internal Reasoning Elements

Math Word Problems

Symbolic Manipulation

Common Sense Reasoning

Standard Prompting

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

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?

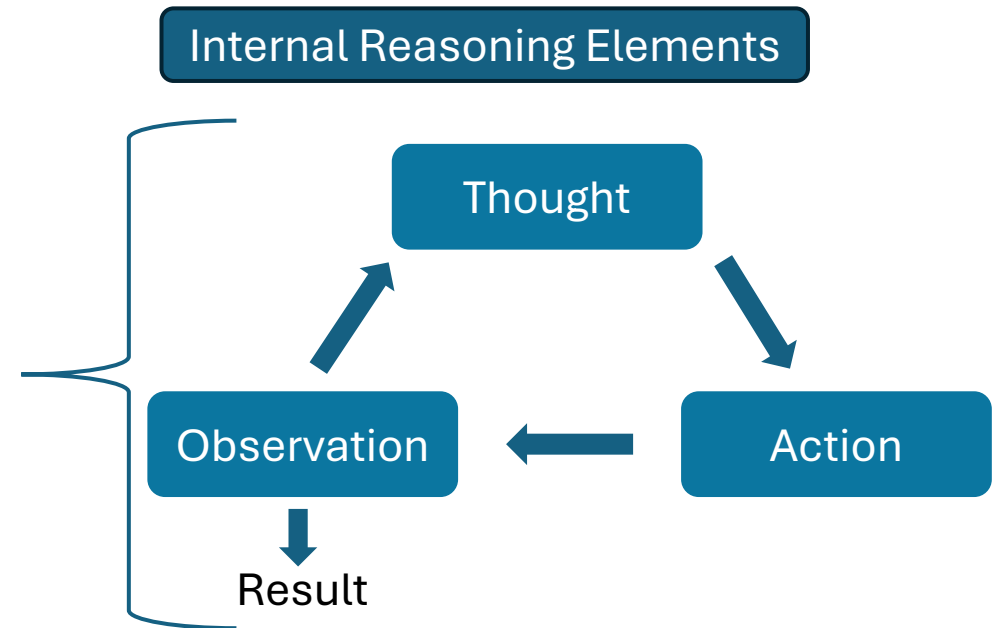
Model Output

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. ✅

What is Agentic AI?

"Appearance of Reasoning"

Approach 2: Reason Action (ReAct)



(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 X

(1a) Standard

Answer: 3,000 X

(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 X

(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) X

(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 ✓

What are the ramifications of Agentic AI?

- Game changer!!
- Ability to:
 - mimic reasoning,
 - self critic (grade itself),
 - correct its mistakes,
 - save learnings from mistakes (never repeat mistakes again)
 - never sleeps or take a break

Solve very complex multi step tasks with little human intervention

An Example

- Task:

"Write a 5-page article about alternative investments offered by company X. Please make the content compliant with SEC rule 482, FINRA rule 2210, and Regulation Best Interest. Show returns for 1 year, 3 year, and 5 years. Compare and contrast with similar products from company Y. Generate 3 graphics incorporating brand guidelines found at ["X/assetmanagement/marketing_templates/brand/images"](#) "

- *Results:*

- Agentic AI only:

- Took 5 minutes,
 - Generate the content,
 - Investment figures verified
 - Generated content -> generate 3 images to convey the content visually (with corporate color pallet).
 - Compliance rules enforced
 - Reviewed by SME

- Human only results:

- 3-5 people collaborating over multiple weeks.

Conclusion

**Long Term Global GDP Growth
1 CE to 2022**

