

Guest Lecture GaTech CS 4644 / 7643 Deep Learning

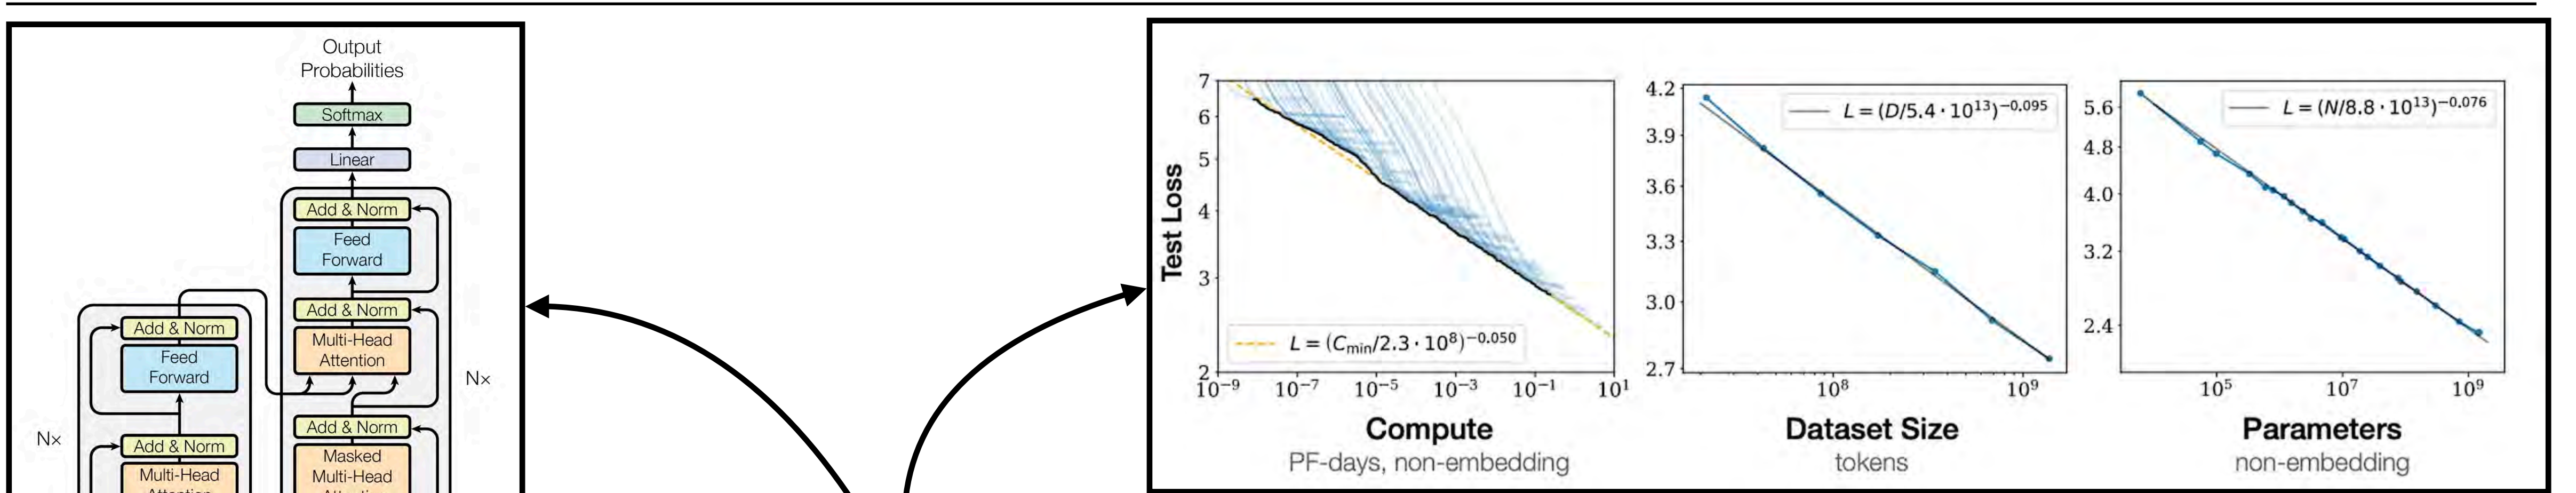
LLM Agents

**extensions of LLMs or
start of something wonderful?**

Hao Zhu

<https://zhuhao.me>



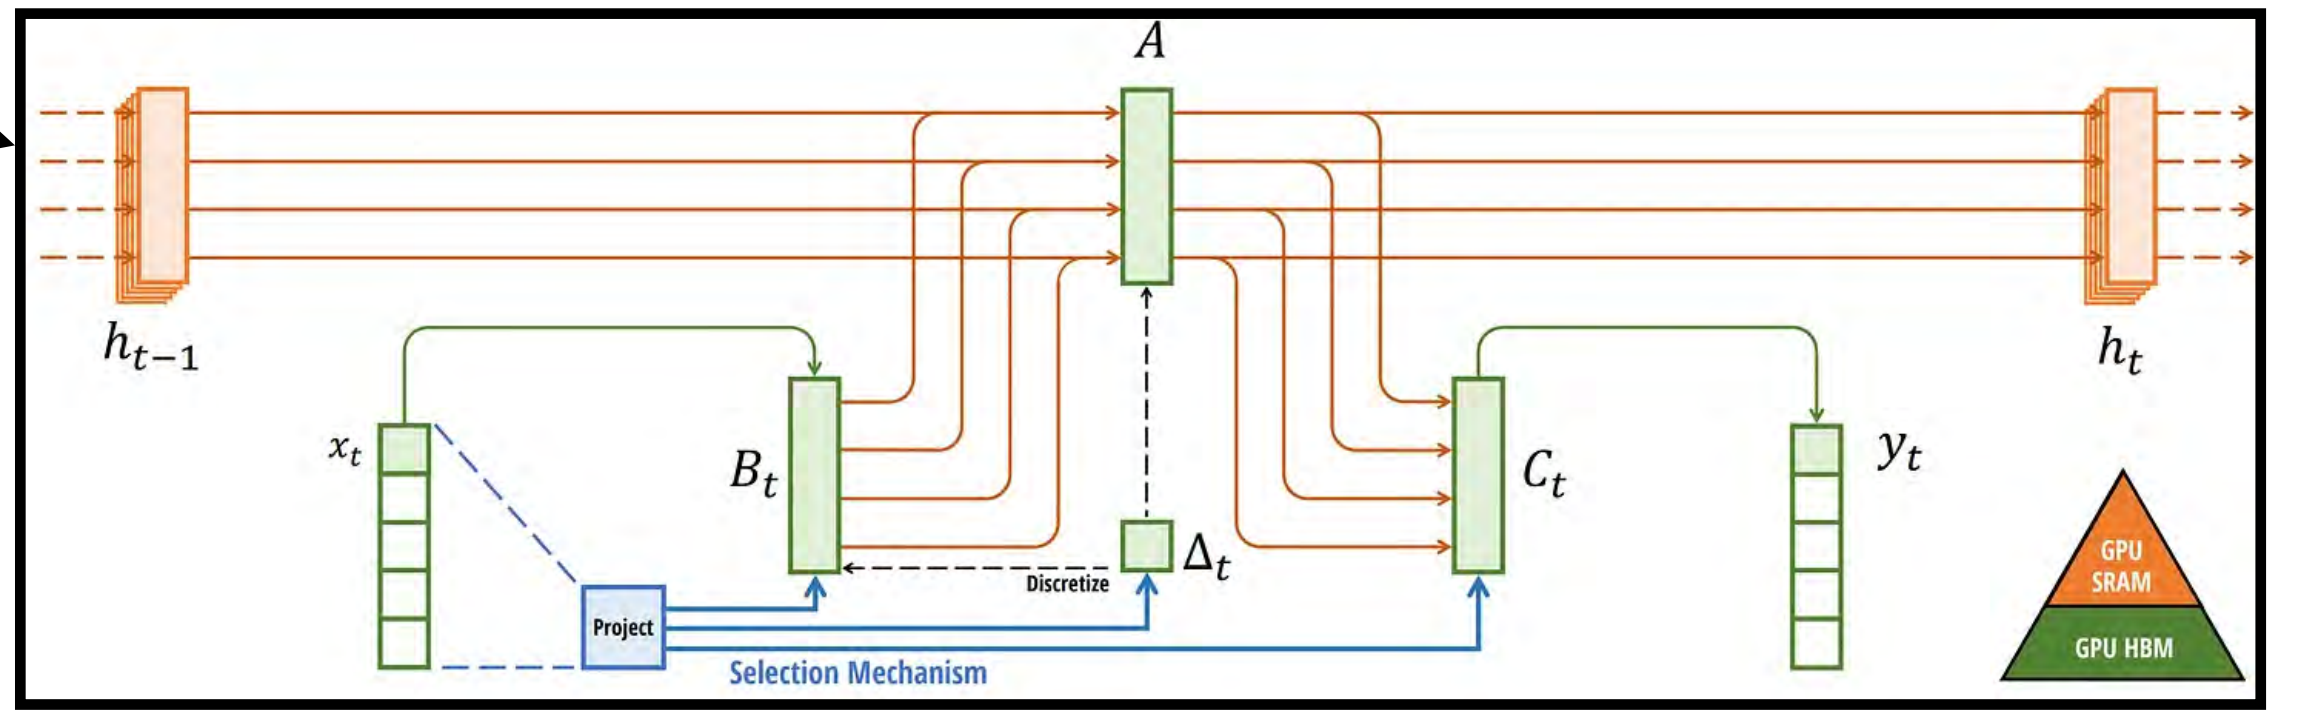


LLM Agent

large

Transformer

or Mamba
& other efficient archs

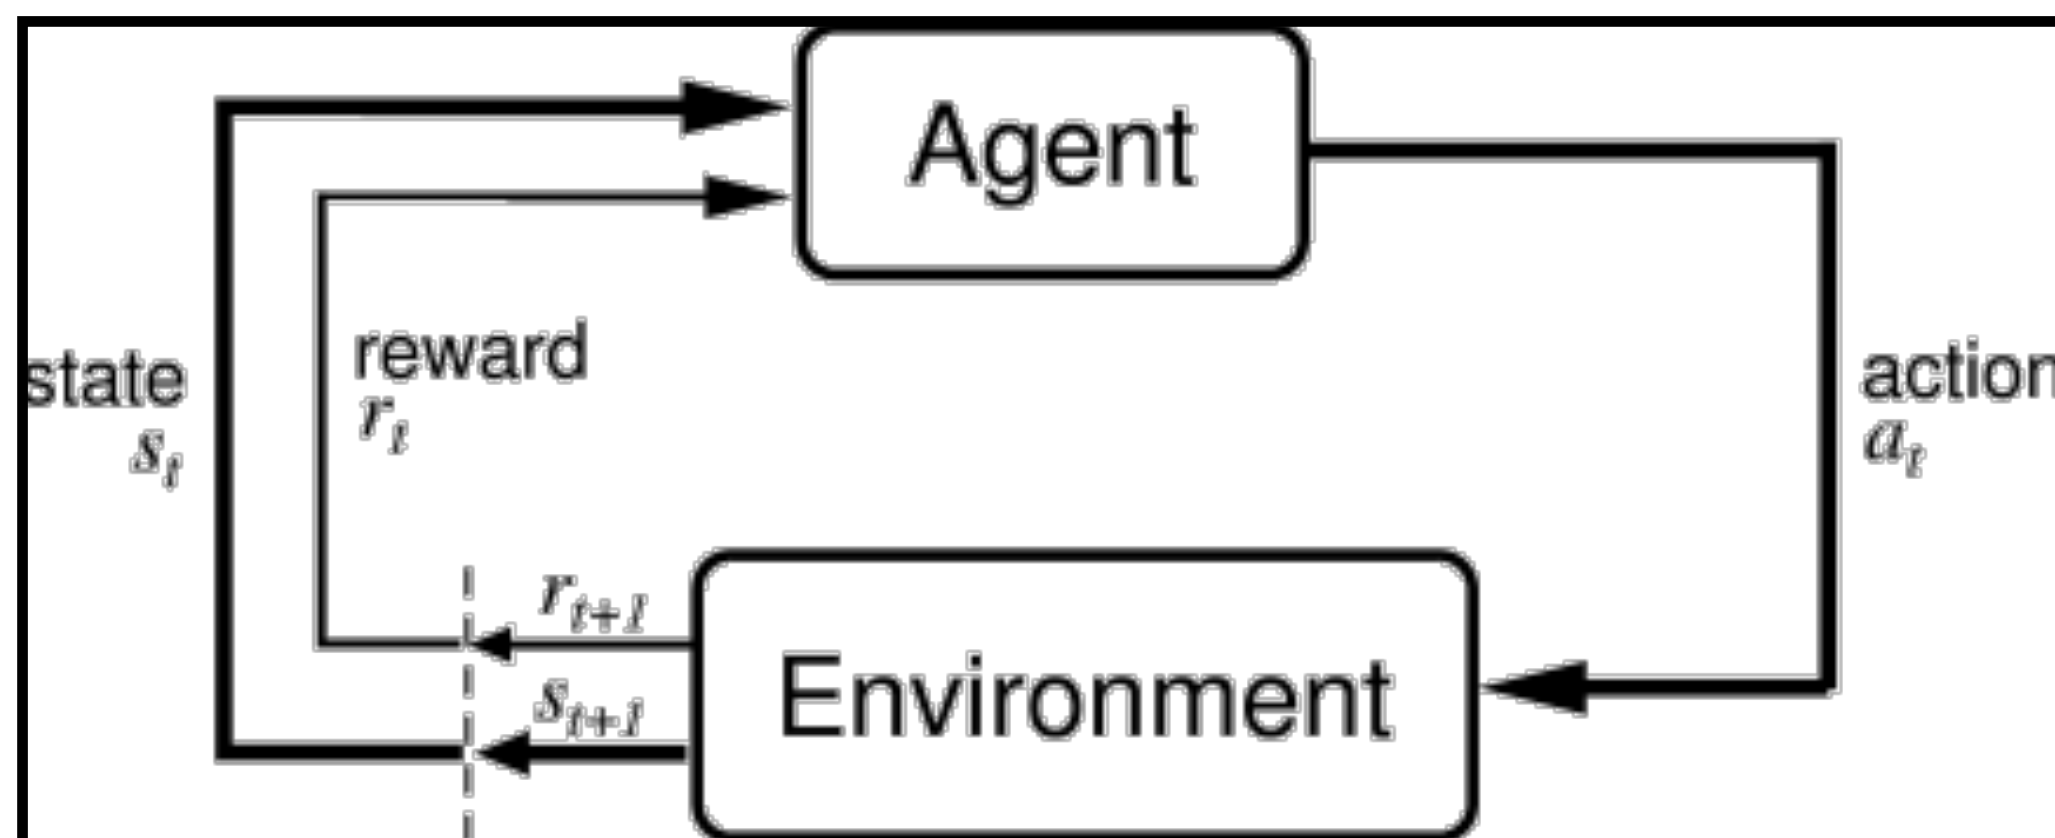




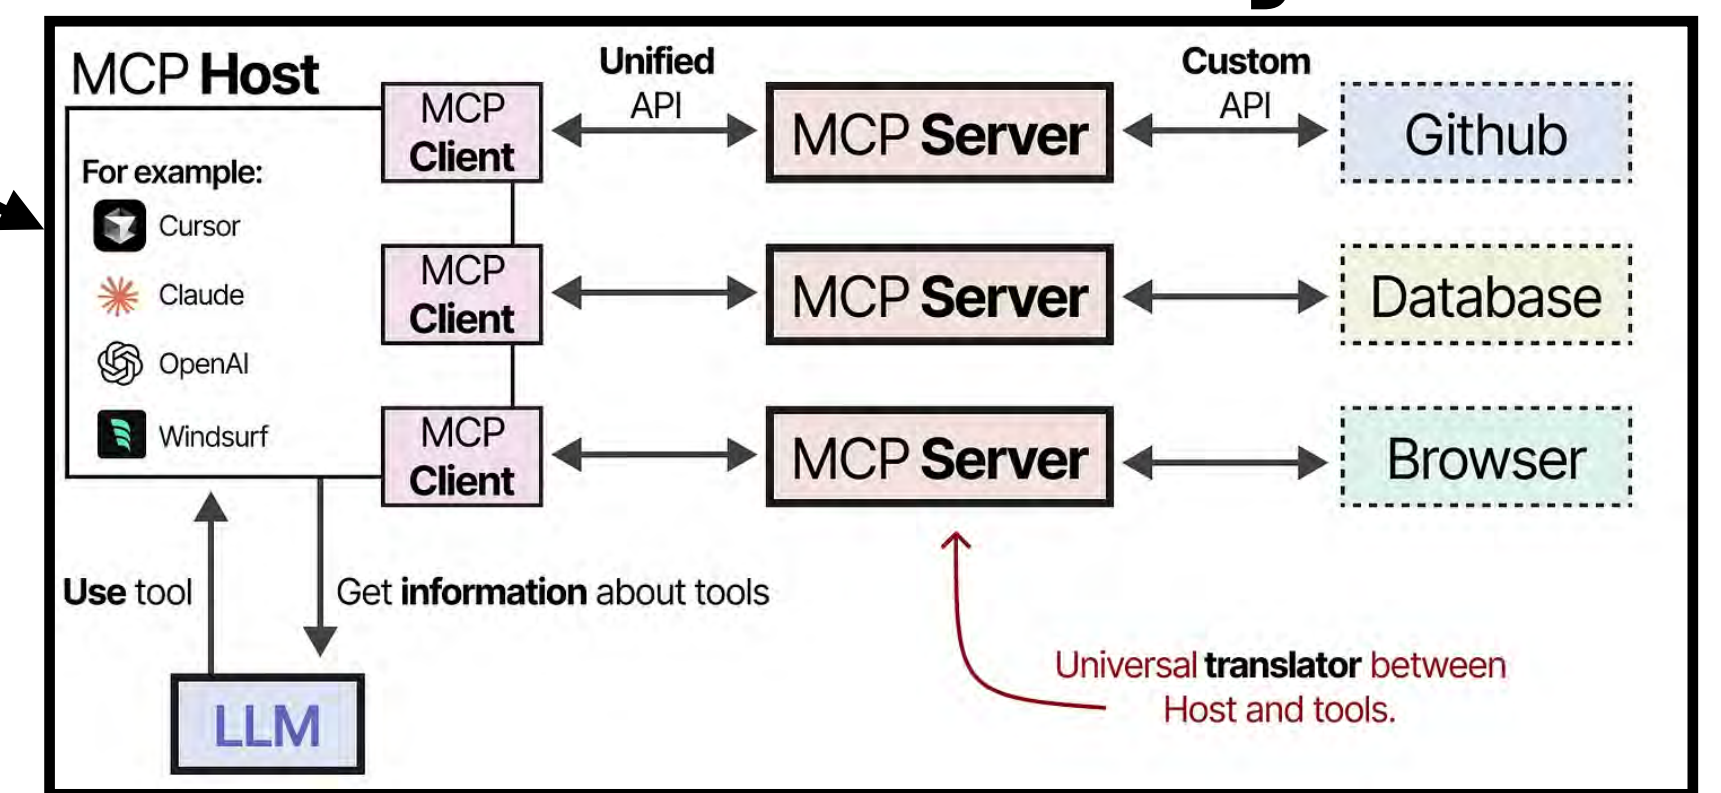
assistants?

models
connected to
tools/
memory etc.

LLM Agent



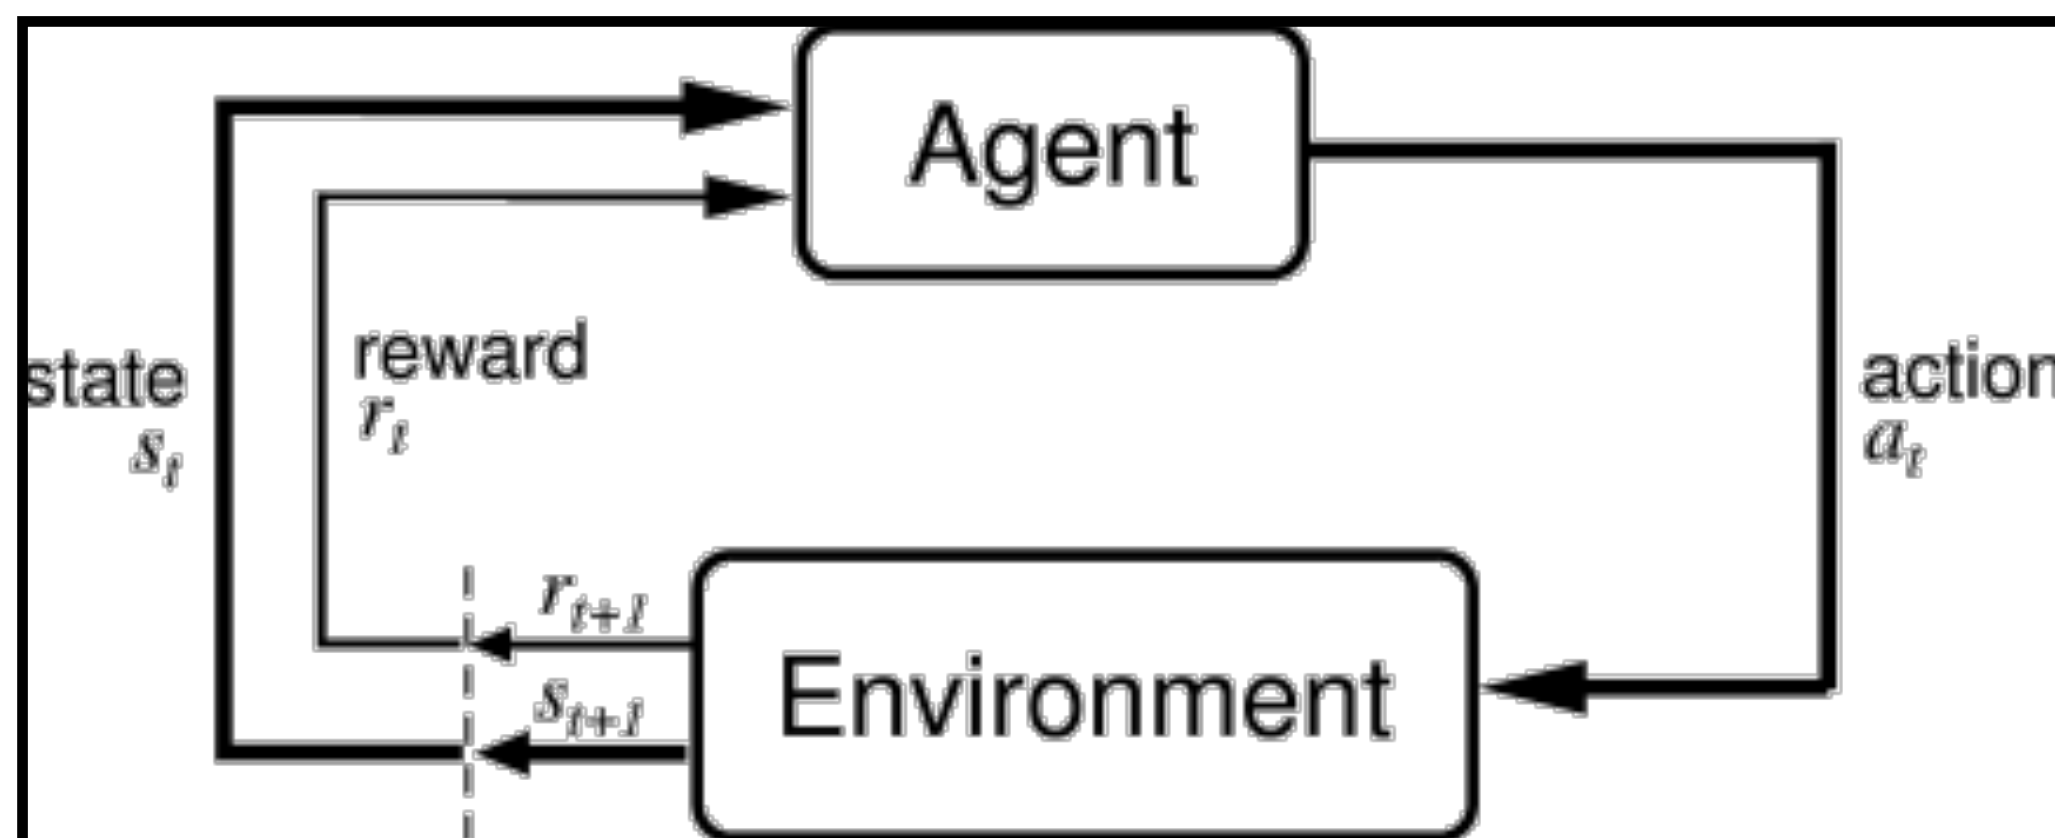
decision
makers



LLM Agent = ?

agent thinking

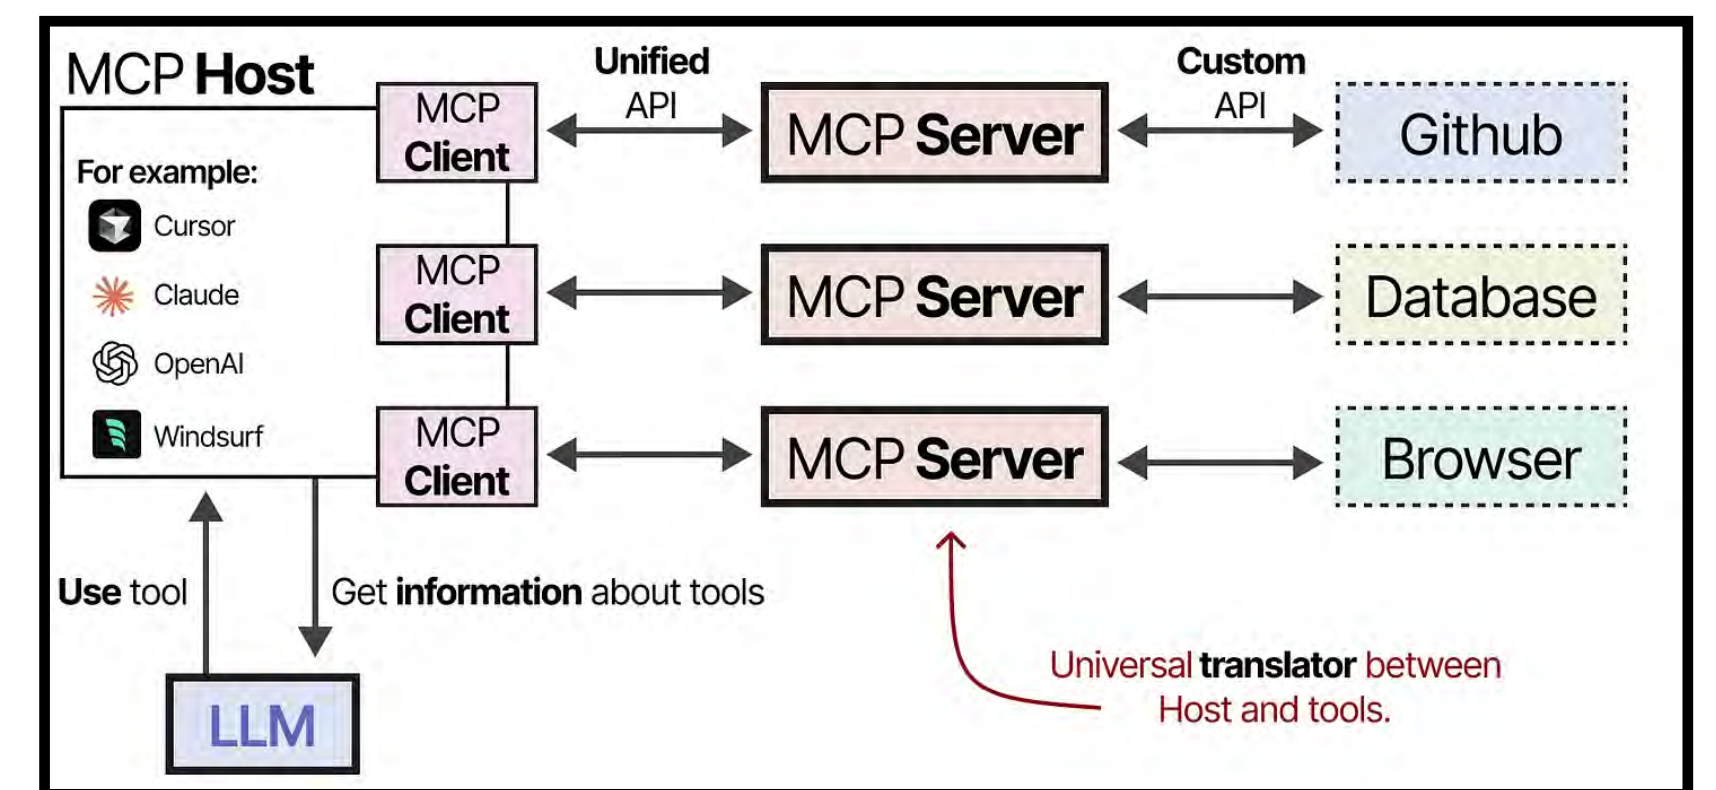
decision makers
powered by LLMs



or

LLM thinking

LLMs connected to
tools/memory etc.

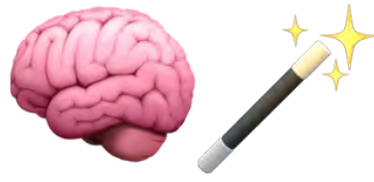


two useful thinking tools



agent thinking

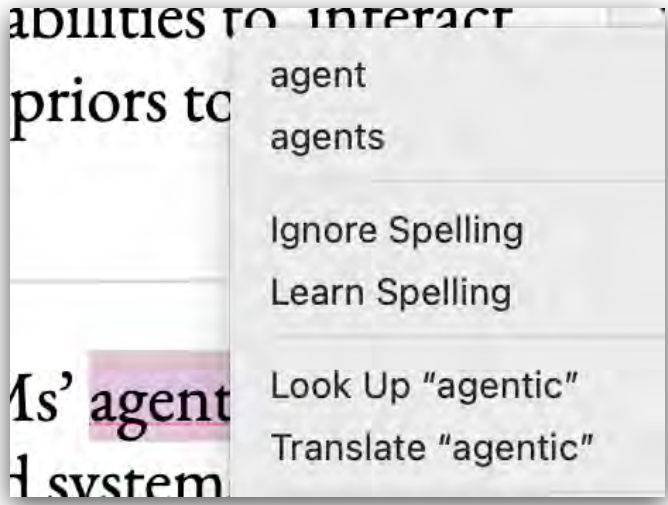
How to improve the models' capabilities to interact with the world? Are LLMs good priors to start with?



LLM Thinking

How to make full use of the LLMs' agentic capabilities? What algorithms and systems we should on top of them?

Yeah, I know. "Agentic" is weird.



“emergent” capabilities of LLMs

*we only wanted to train next-token
predictors, how did it learn...*



coding



memory



in-context learning



grounding

“emergent” capabilities of LLMs

*we only wanted to train next-token
predictors, how did it learn...*



coding



memory



in-context learning



grounding

“emergent” capabilities of LLMs

we only wanted to train next-token predictors, how did it learn...



coding



memory

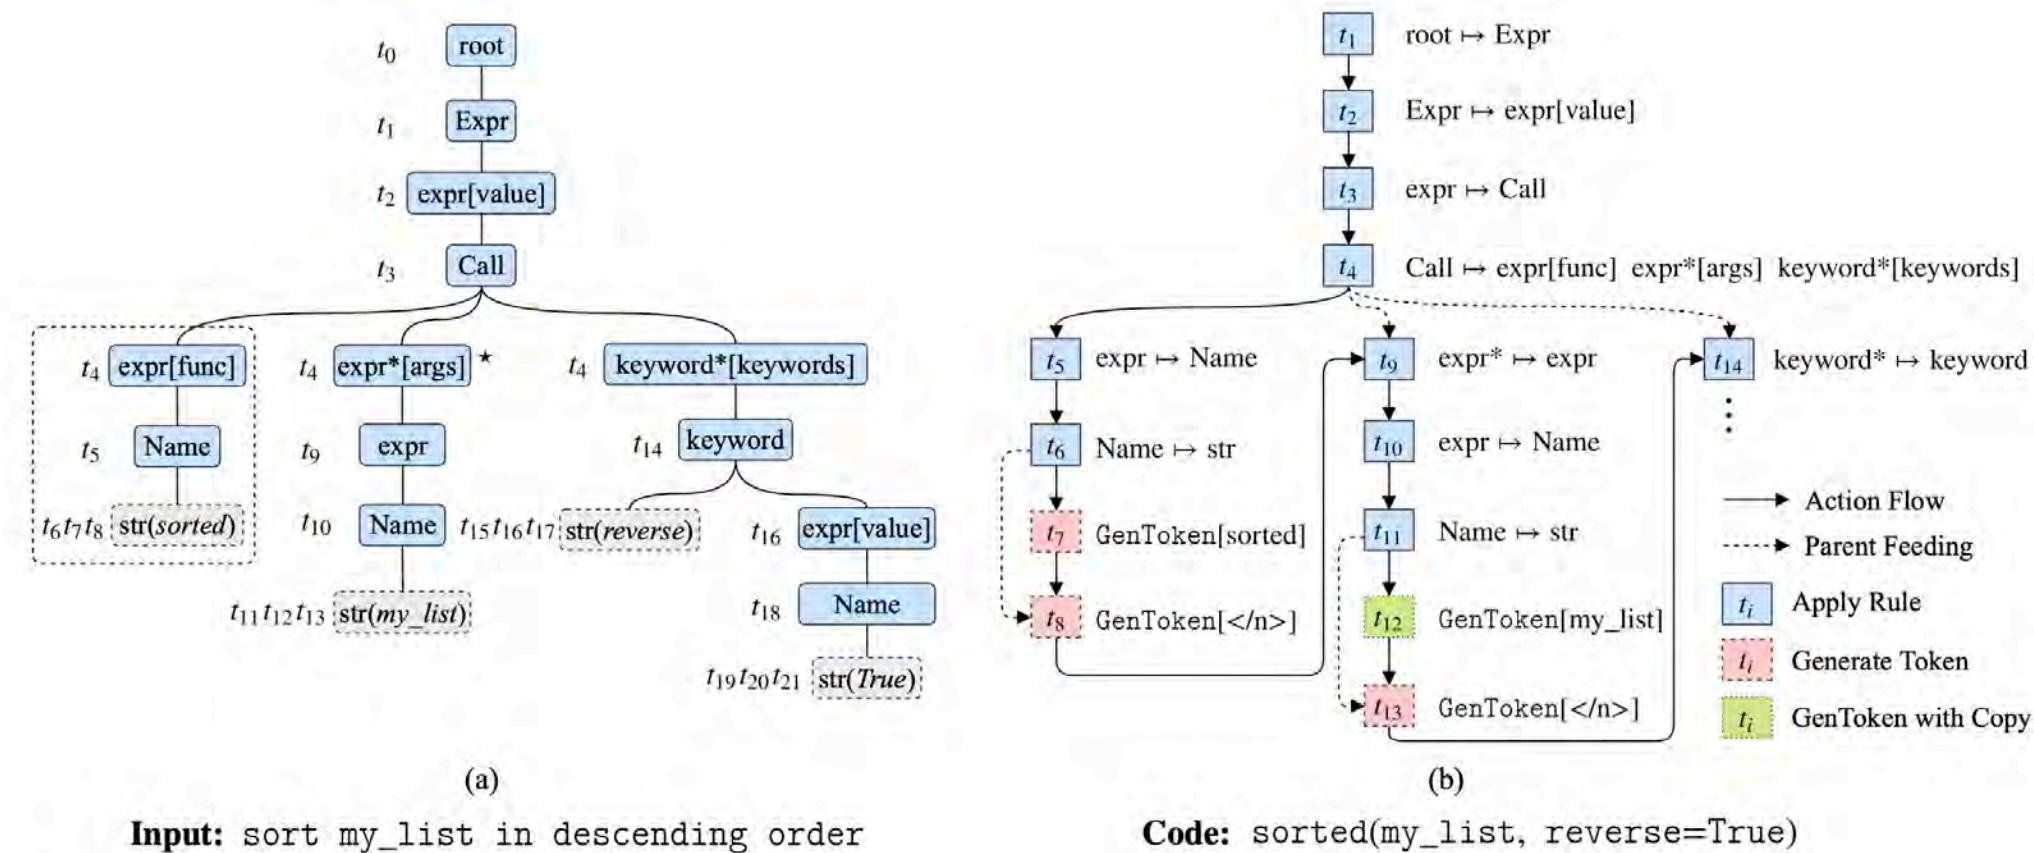


in-context learning



grounding

coding — LLM doing surprisingly well



A syntactic neural model for general-purpose code generation

```
text.style.left =
rocketship.offsetLeft + 'px';
text.style.top =
rocketship.offsetTop + 'px';

document.body.appendChild(text
);
xSpeed = 20;
setTimeout(function() {
xSpeed = 5;

document.body.removeChild(text
);
}, 250);
};

/* Now add an image of an
asteroid:
https://d.newsweek.com/en/full
/1721338/asteroid.jpg?
w=1600&h=1600&q=88&f=9d82d35c9
de96a82b3fcaf7705eb325b */
var asteroid =
document.createElement('img');
asteroid.src =
'https://d.newsweek.com/en/ful
l/1721338/asteroid.jpg?
w=1600&h=1600&q=88&f=9d82d35c9
de96a82b3fcaf7705eb325b';
document.body.appendChild(aste
roid);
```

Make it be the size of the rocketship times 0.75

pre-LLM NL2Code

semantic parsing → AST

(Yin and Neubig, 2017)

LLM NL2Code

instruction following

(OpenAI Codex, 2021)


coding — and they got even better now

All Courses > Short Courses > Vibe Coding 101 with Replit

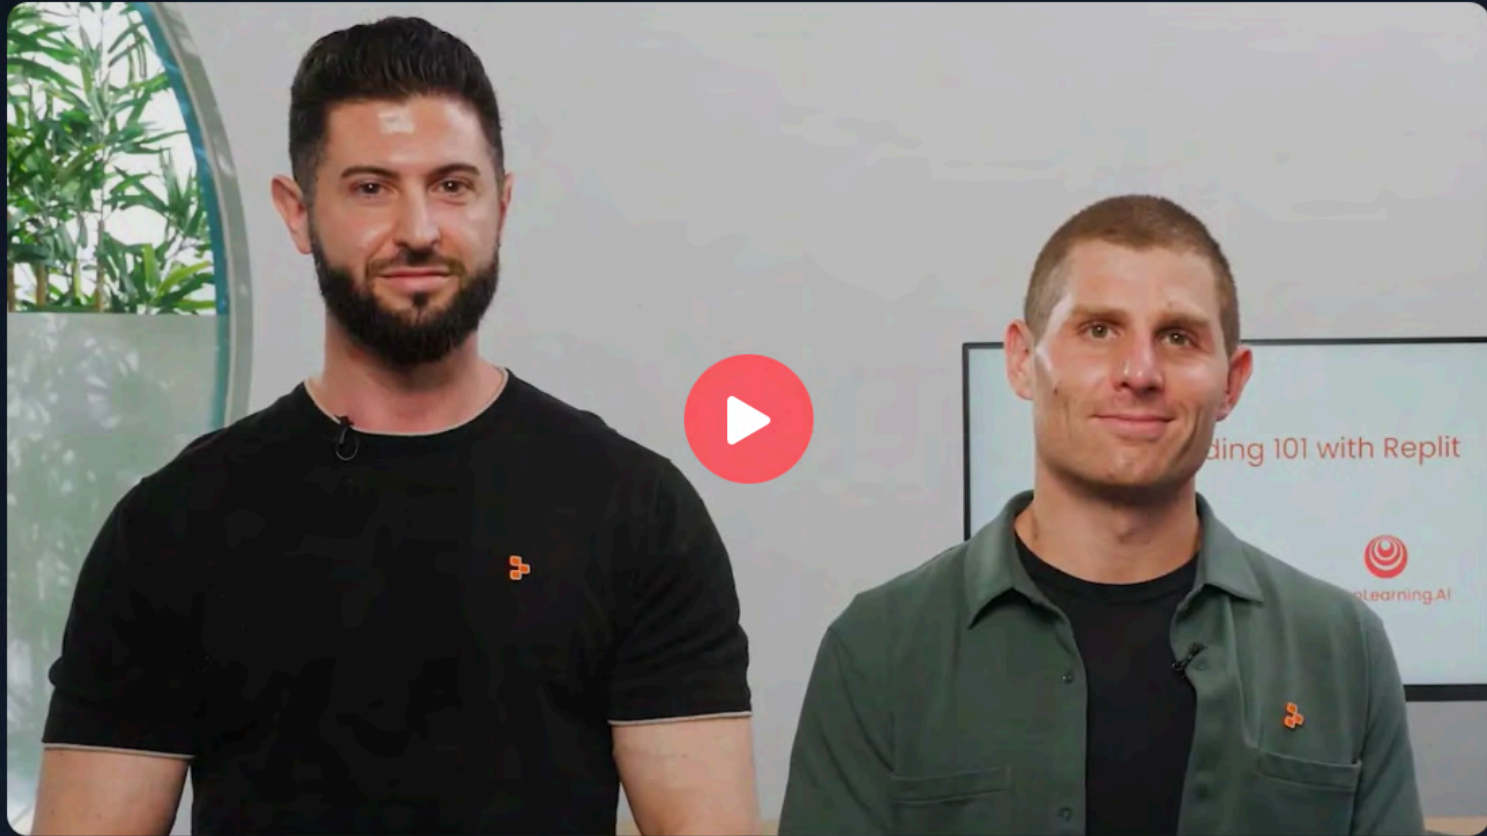
Short Course Beginner 1 Hour 34 Minutes

Vibe Coding 101 with Replit

Instructors: Michele Catasta, Matt Palmer



Enroll for Free



Vibe Coding

ask LLM to code and it just works.


AK @_akhaliq [Subscribe](#)

How to vibe code games and deploy them for free entirely on Hugging Face

get started in a few minutes, no logins, signups or subscriptions

deploy on Hugging Faces spaces for anyone to play for free

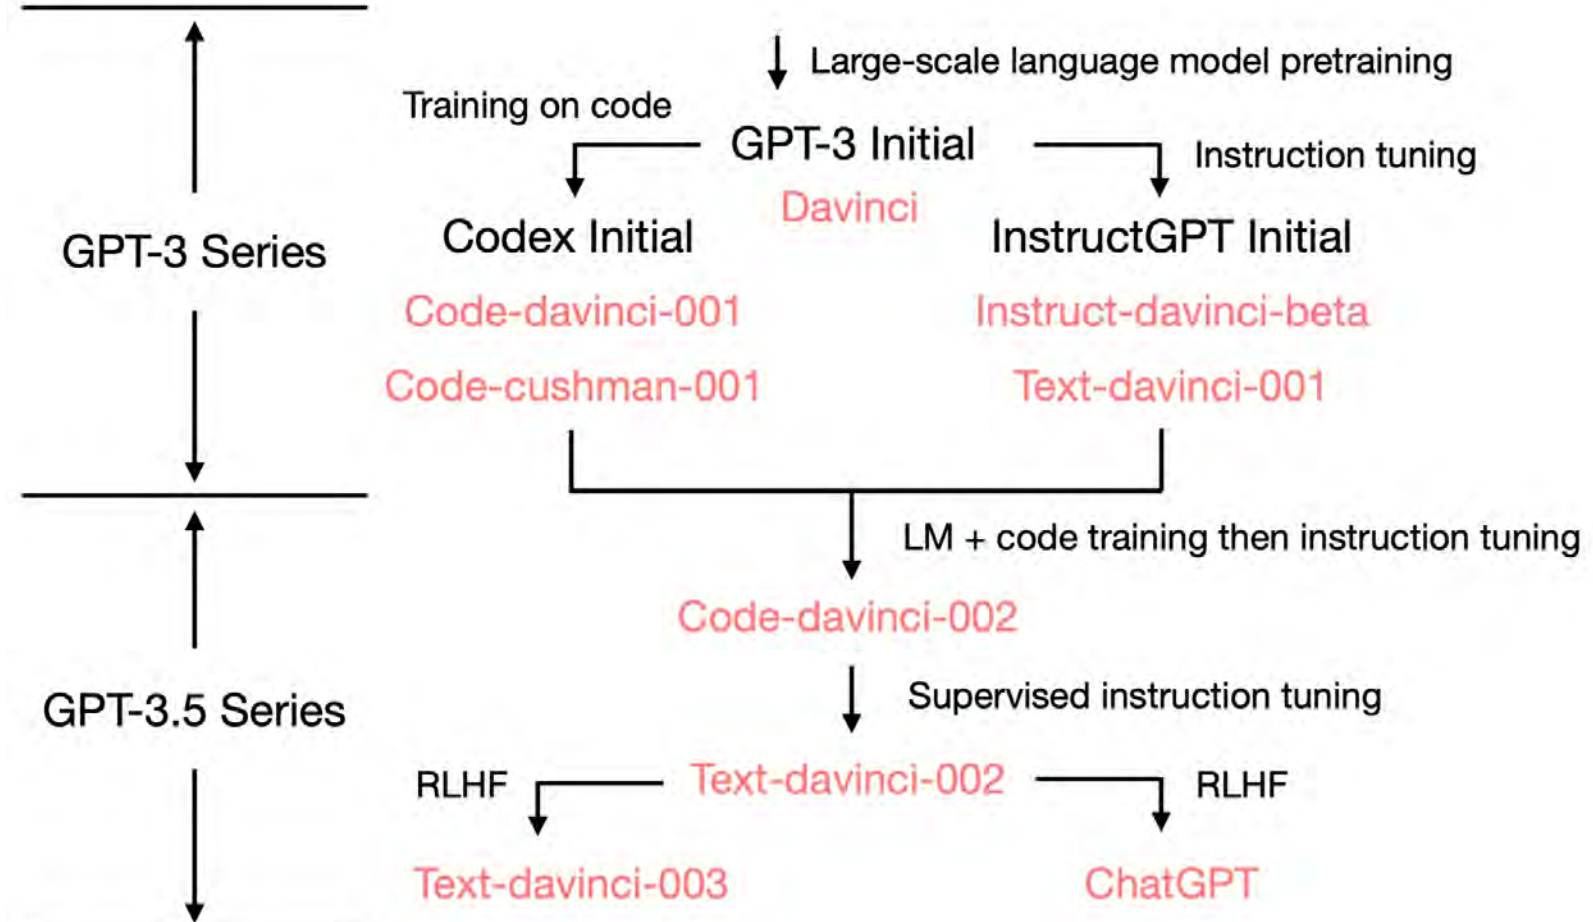
this game was made with a few prompts in anychat starting with "make a three.js game" and then adding "make it better" with some small tweaks



0:00

10:14 AM · Mar 19, 2025 · 39.5K Views

coding capabilities (con't)



How does GPT Obtain its Ability? Tracing Emergent Abilities of Language Models to their Source

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 tennis balls.
 tennis_balls = 5
 2 cans of 3 tennis balls each is
 bought_balls = 2 * 3
 tennis_balls. The answer is
 answer = tennis_balls + bought_balls

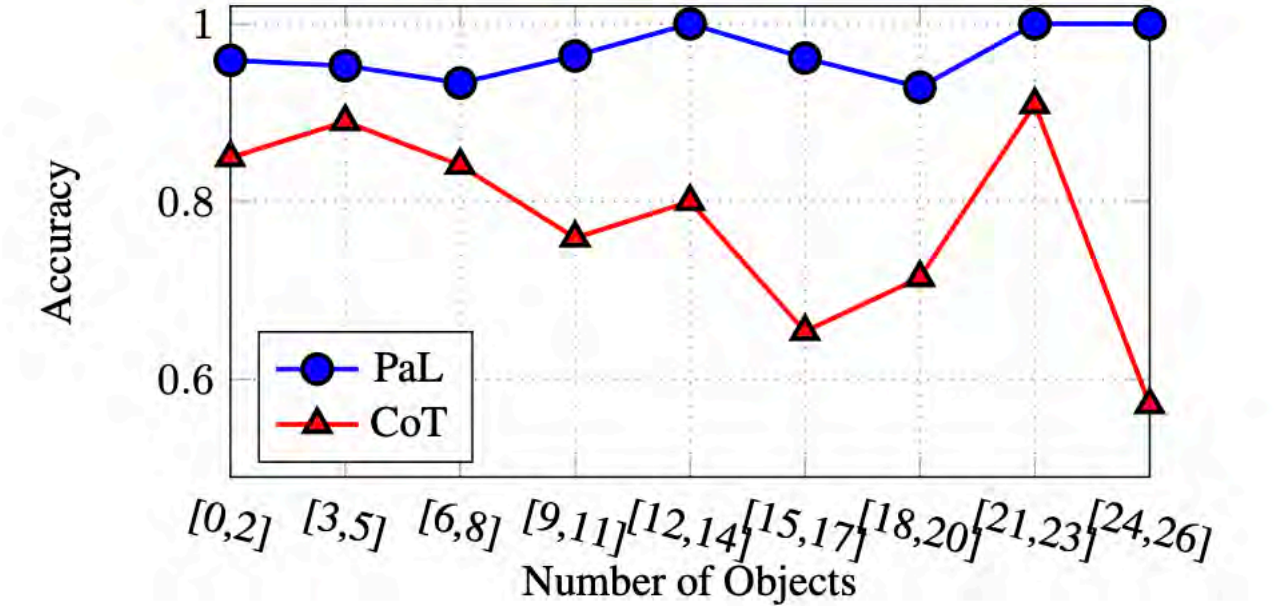
Q: The bakers at the Beverly Hills Bakery baked 200 loaves of bread on Monday morning. They sold 93 loaves in the morning and 39 loaves in the afternoon. A grocery store returned 6 unsold loaves. How many loaves of bread did they have left?

Model Output

A: The bakers started with 200 loaves
 loaves_baked = 200
 They sold 93 in the morning and 39 in the afternoon
 loaves_sold_morning = 93
 loaves_sold_afternoon = 39
 The grocery store returned 6 loaves.
 loaves_returned = 6
 The answer is
 answer = loaves_baked - loaves_sold_morning
 - loaves_sold_afternoon + loaves_returned

```
>>> print(answer)
74
```

✓



PAL: Program-aided Language Models

“early” history

code in the pertaining data

(Fu et al, 2022)

code as a detour

prompt to code for QA

(Gao et al, 2023)

“emergent” capabilities of LLMs

*we only wanted to train next-token
predictors, how did it learn...*



coding



memory



in-context learning



grounding

“emergent” capabilities of LLMs

*we only wanted to train next-token
predictors, how did it learn...*



coding



memory

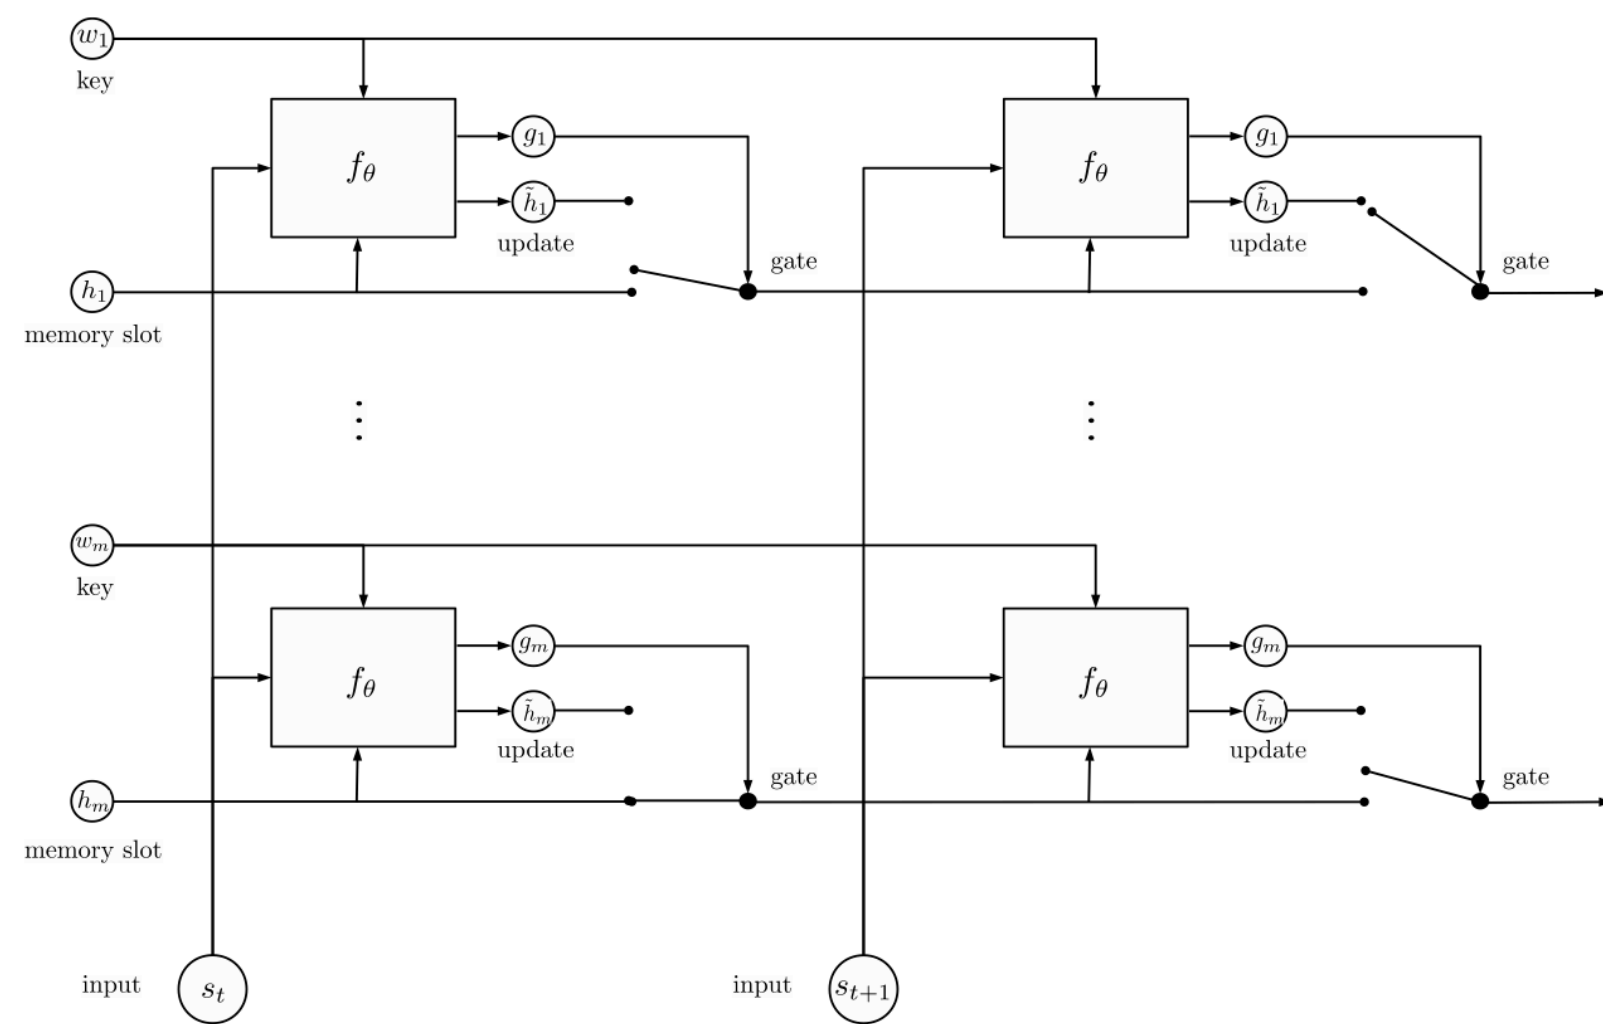


in-context learning



grounding

memory — attention mechanism

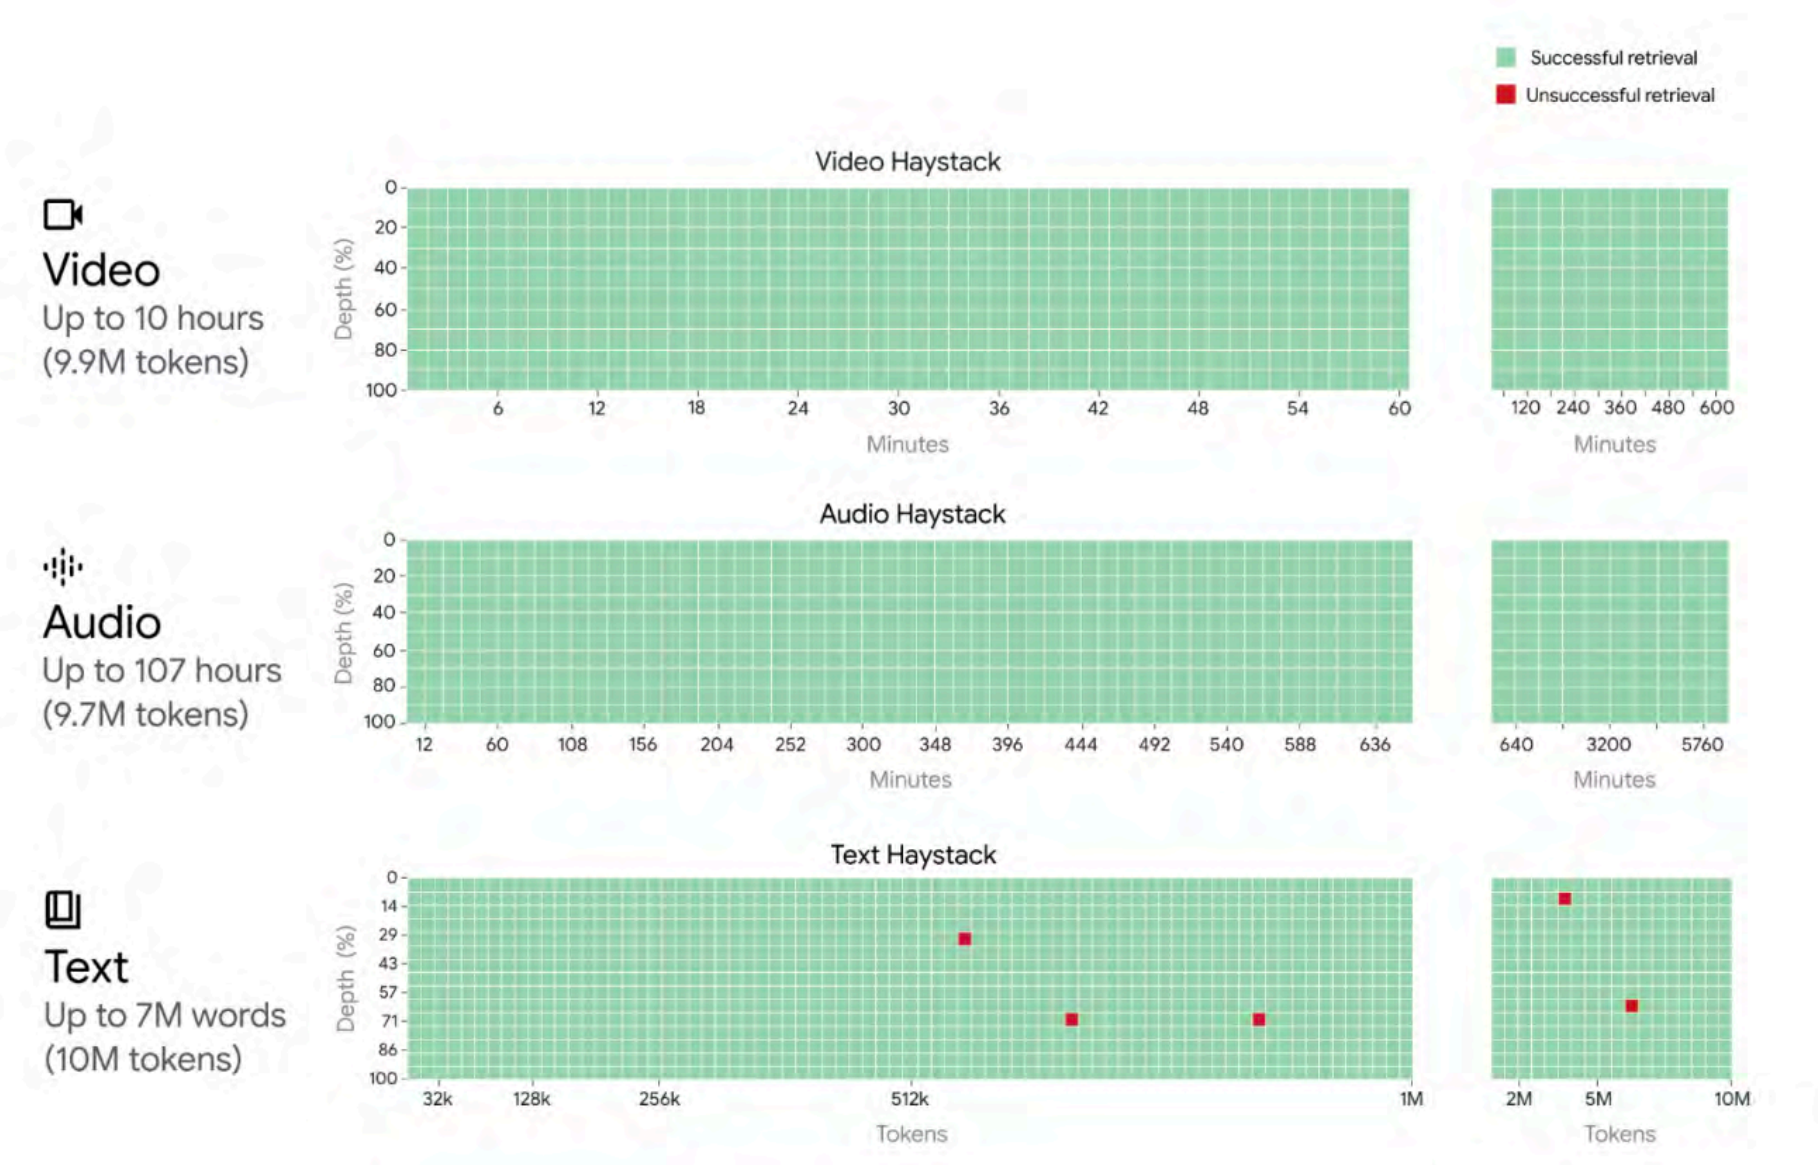


Tracking the World State with Recurrent Entity Networks

pre-LLM attention

attention used in story QA

(Henaff et al, 2017)



Gemini 1.5 Pro

10M ~perfect recall

(Google, 2024)

even stronger memory w/ RAG

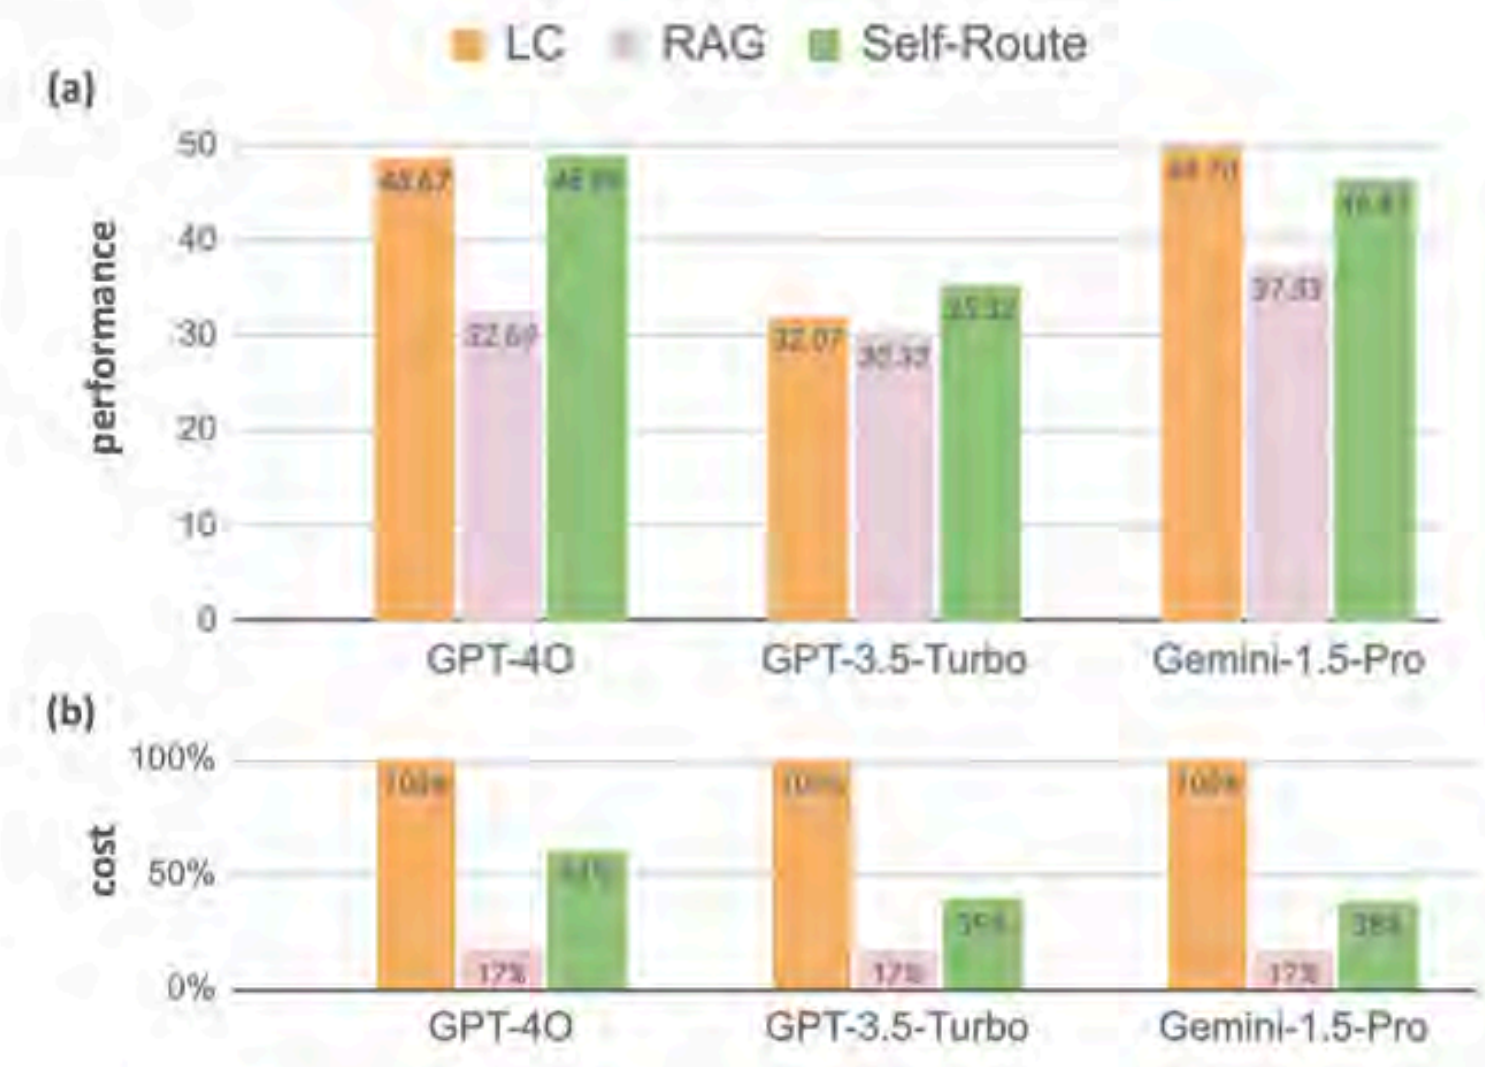
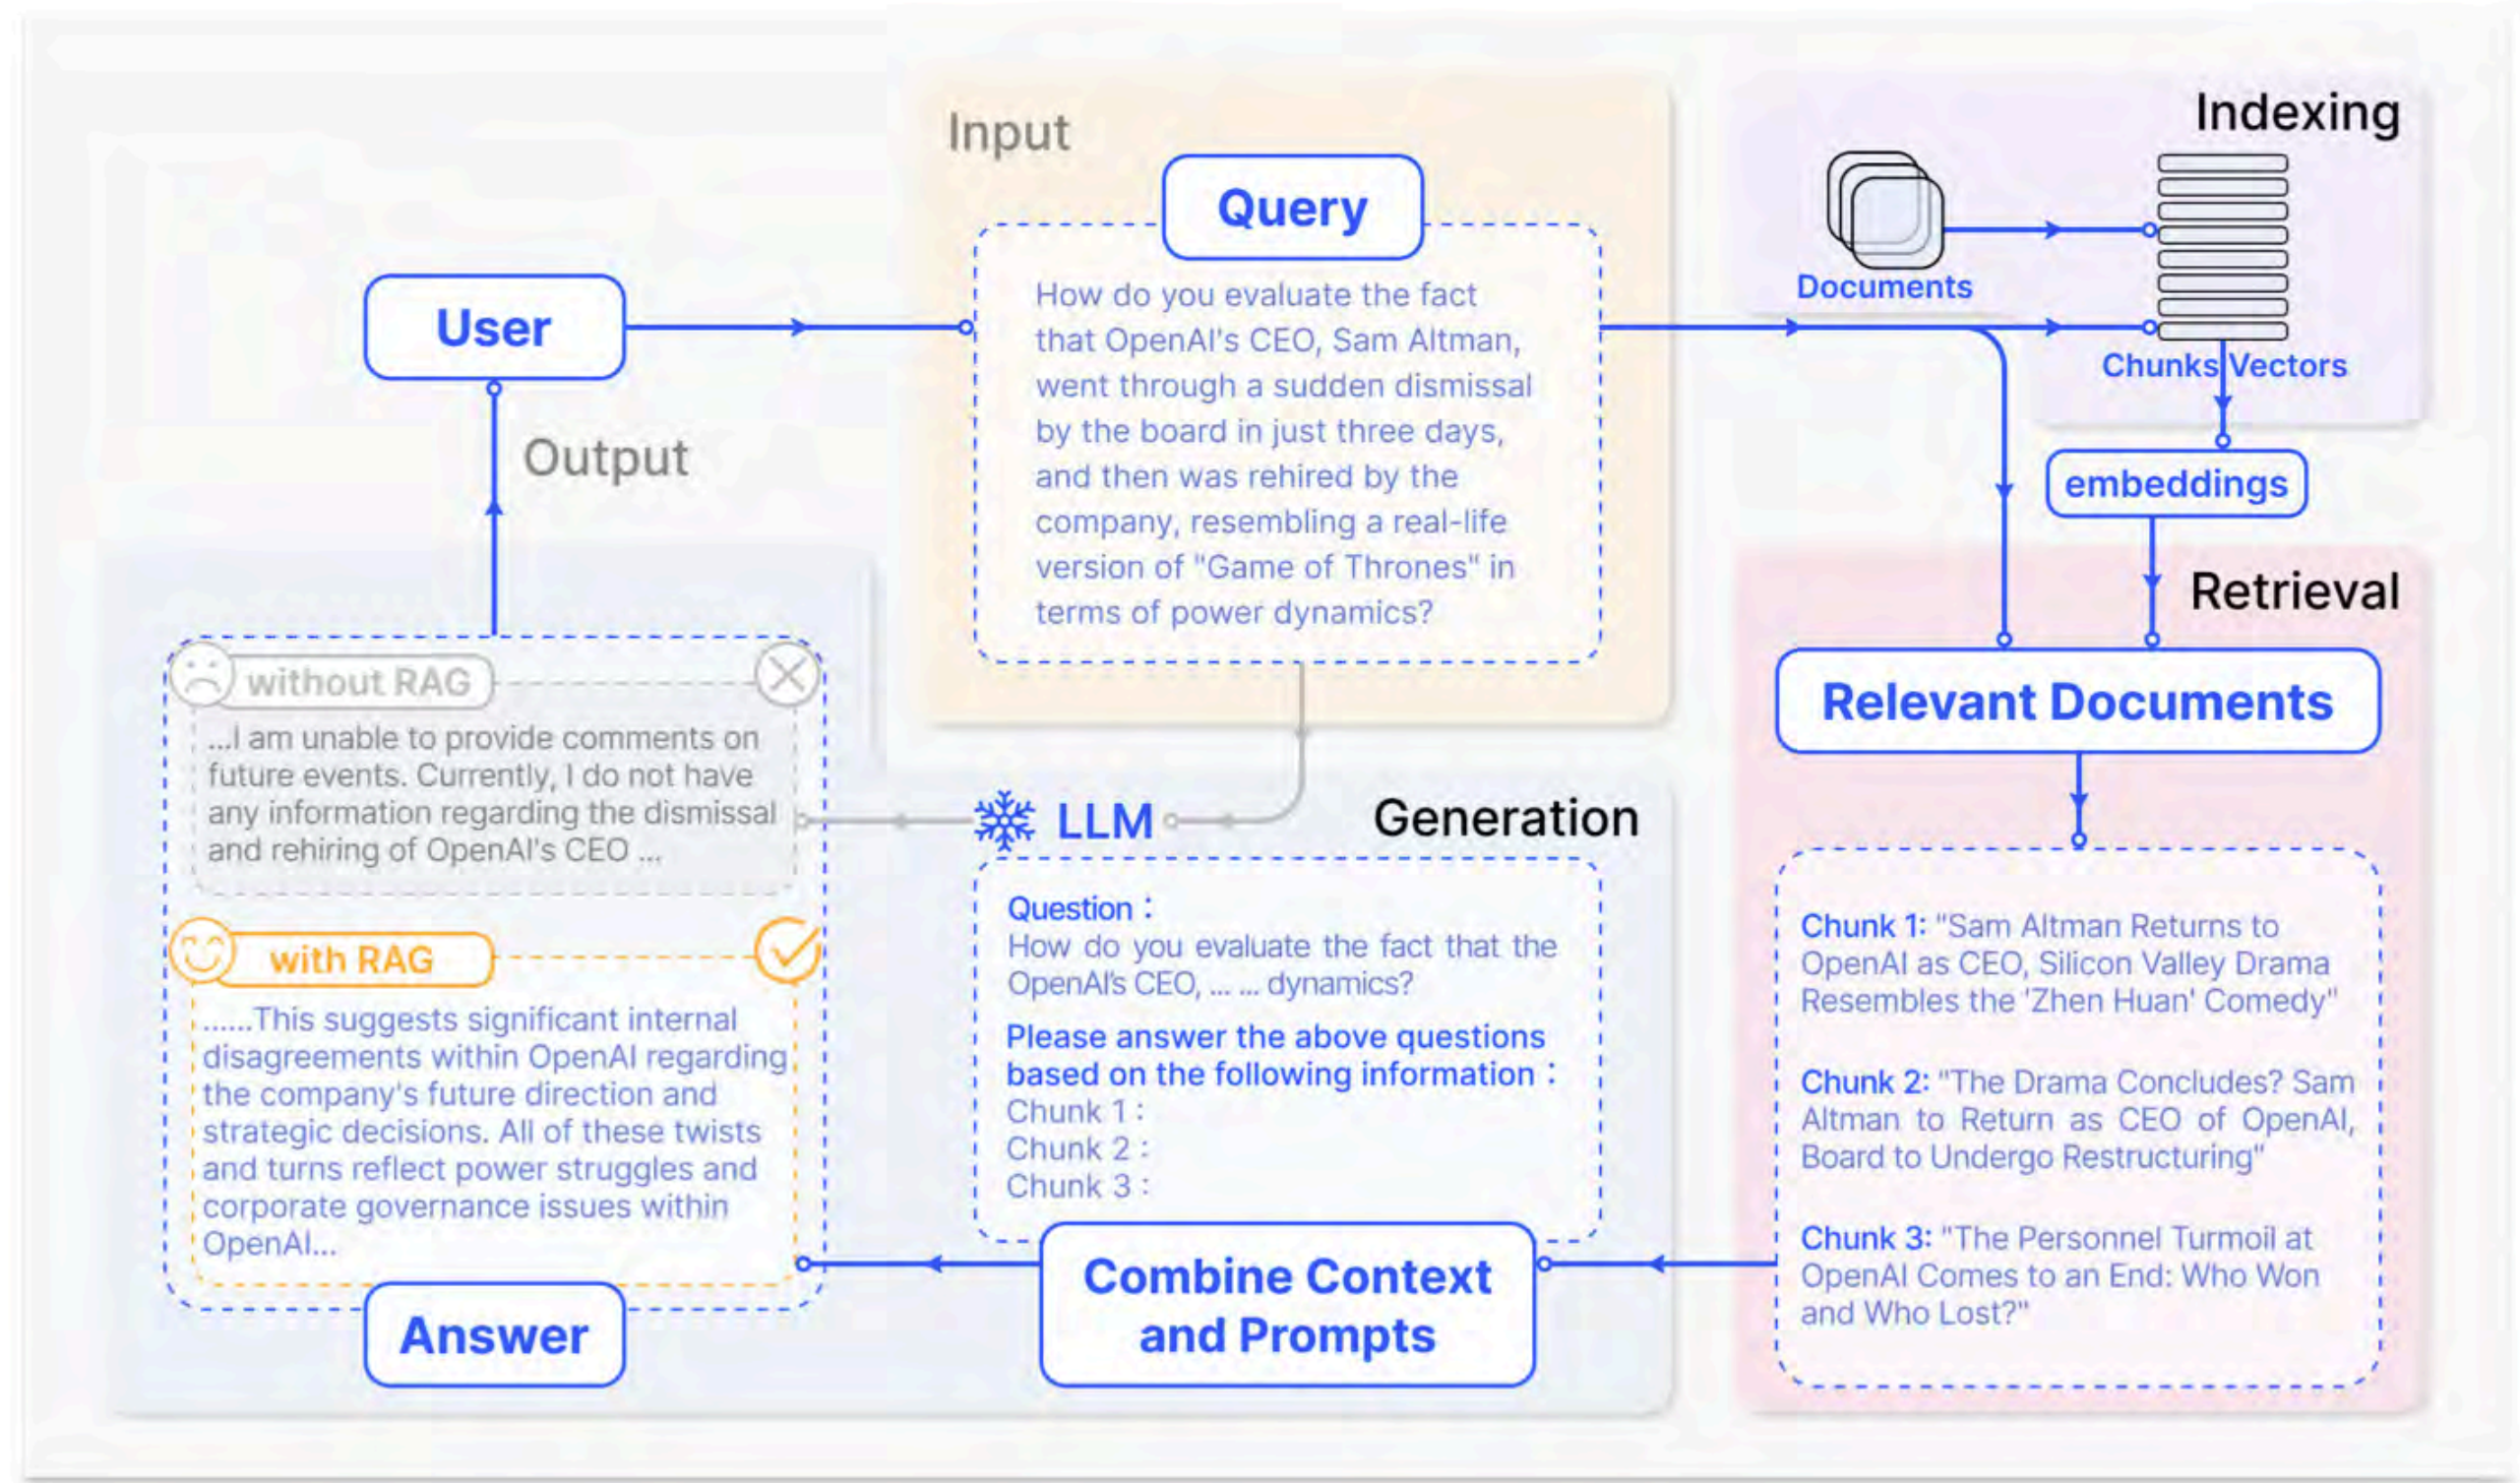


Figure 1: While long-context LLMs (LC) surpass RAG in long-context understanding, RAG is significantly more cost-efficient. Our approach, SELF-ROUTE, combining RAG and LC, achieves comparable performance to LC at a much lower cost.

“emergent” capabilities of LLMs

we only wanted to train next-token predictors, how did it learn...



coding



memory



in-context learning



grounding

“emergent” capabilities of LLMs

*we only wanted to train next-token
predictors, how did it learn...*



coding



memory

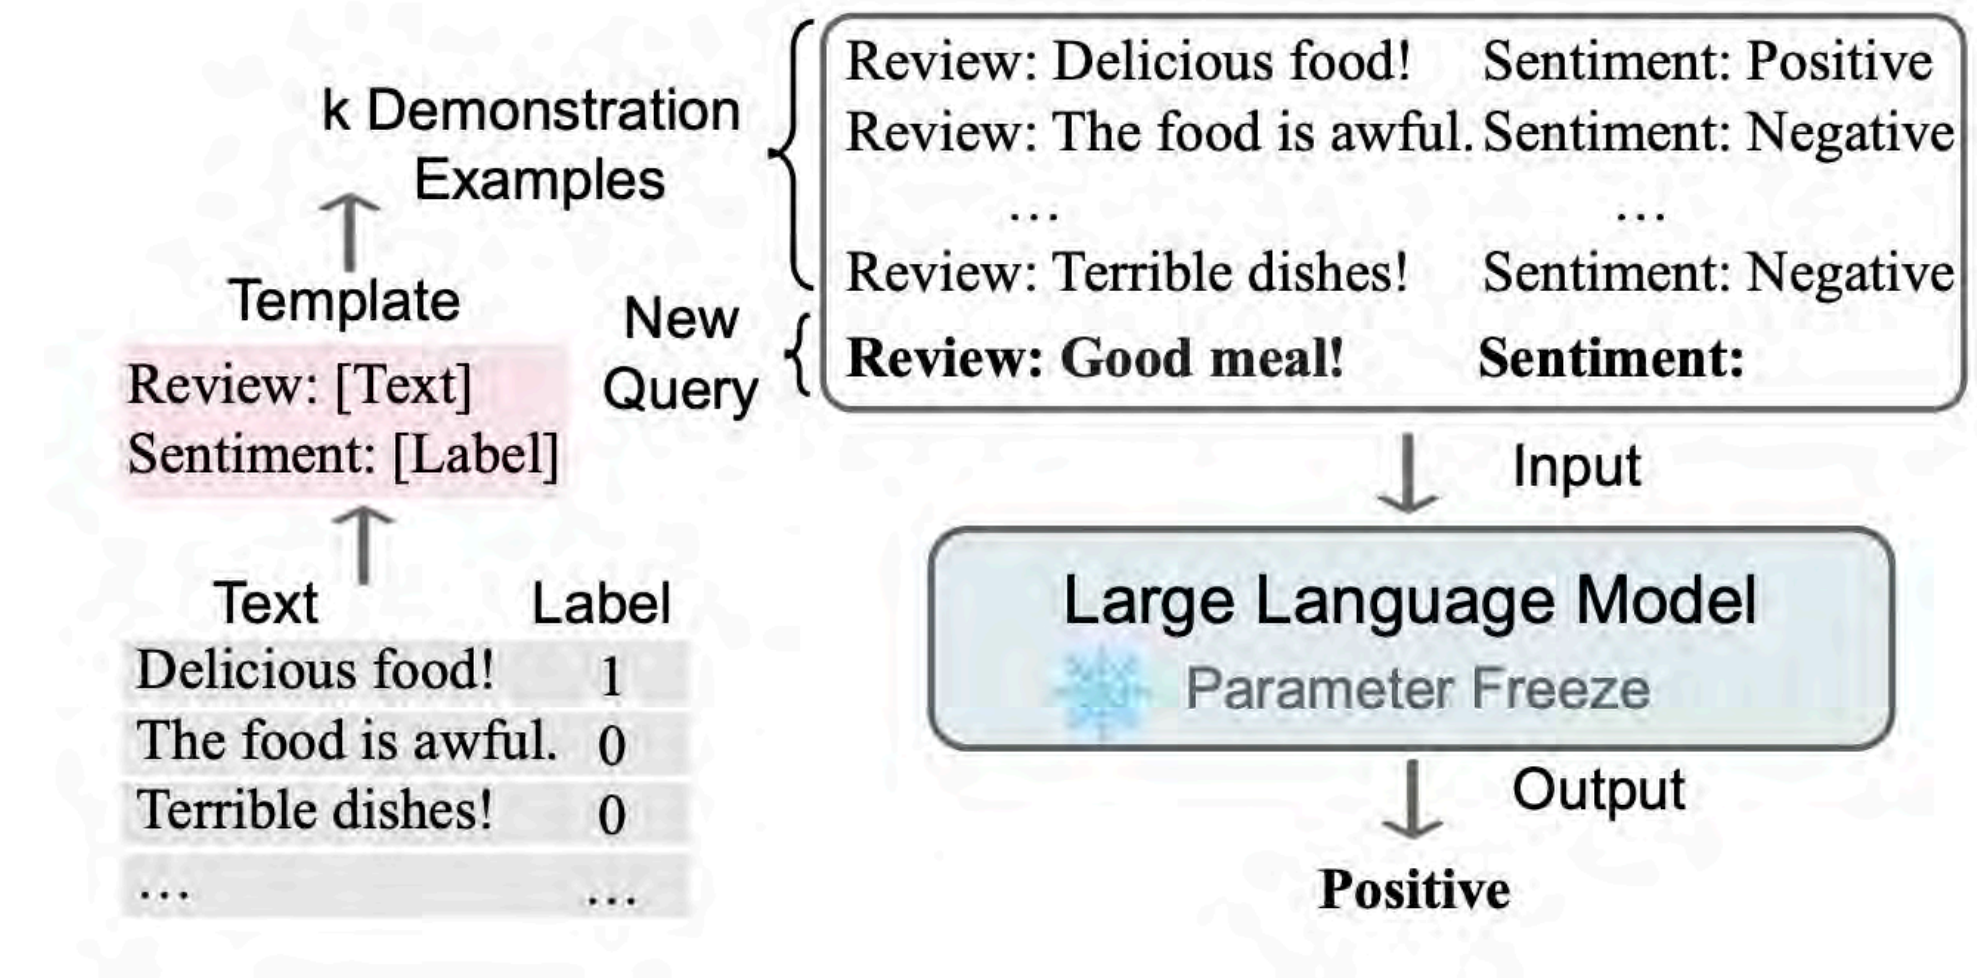
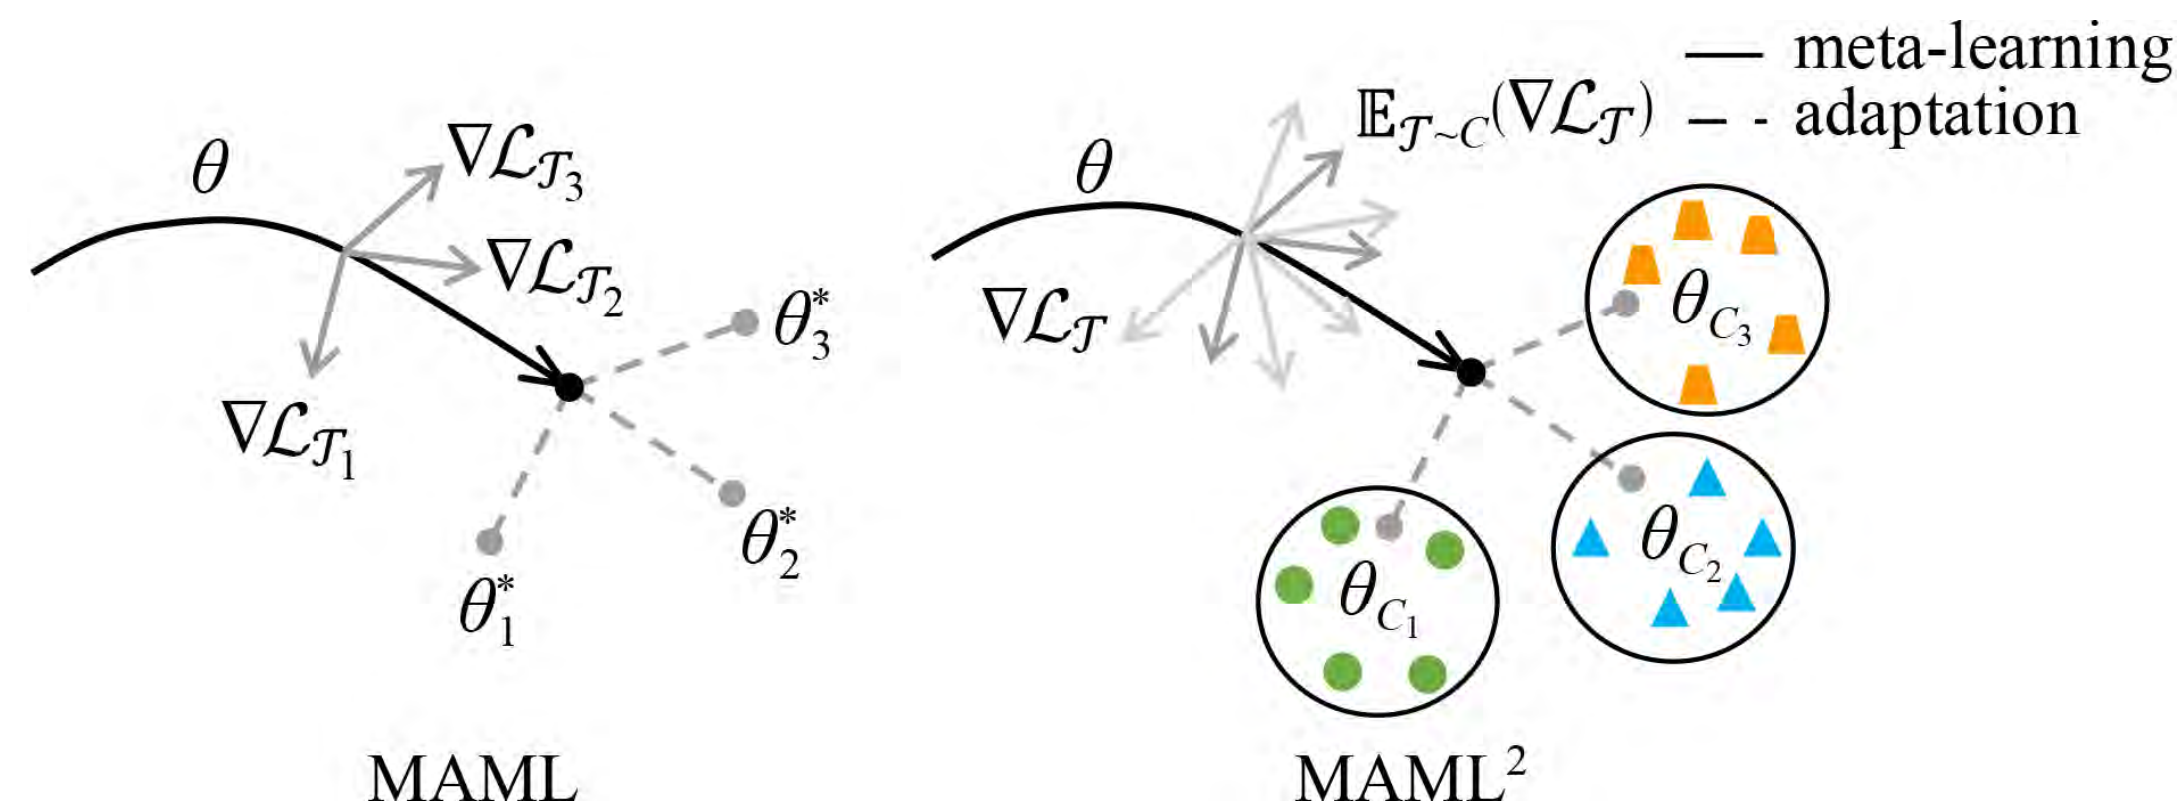


in-context learning



grounding

in-context learning



pre-LLM few-shot learning

gradient-based meta learning

(Finn et al, 2017)

ICL emerges

10M ~perfect recall

(Brown et al, 2020)

“emergent” capabilities of LLMs

*we only wanted to train next-token
predictors, how did it learn...*



coding



memory



in-context learning



grounding

“emergent” capabilities of LLMs

*we only wanted to train next-token
predictors, how did it learn...*



coding



memory

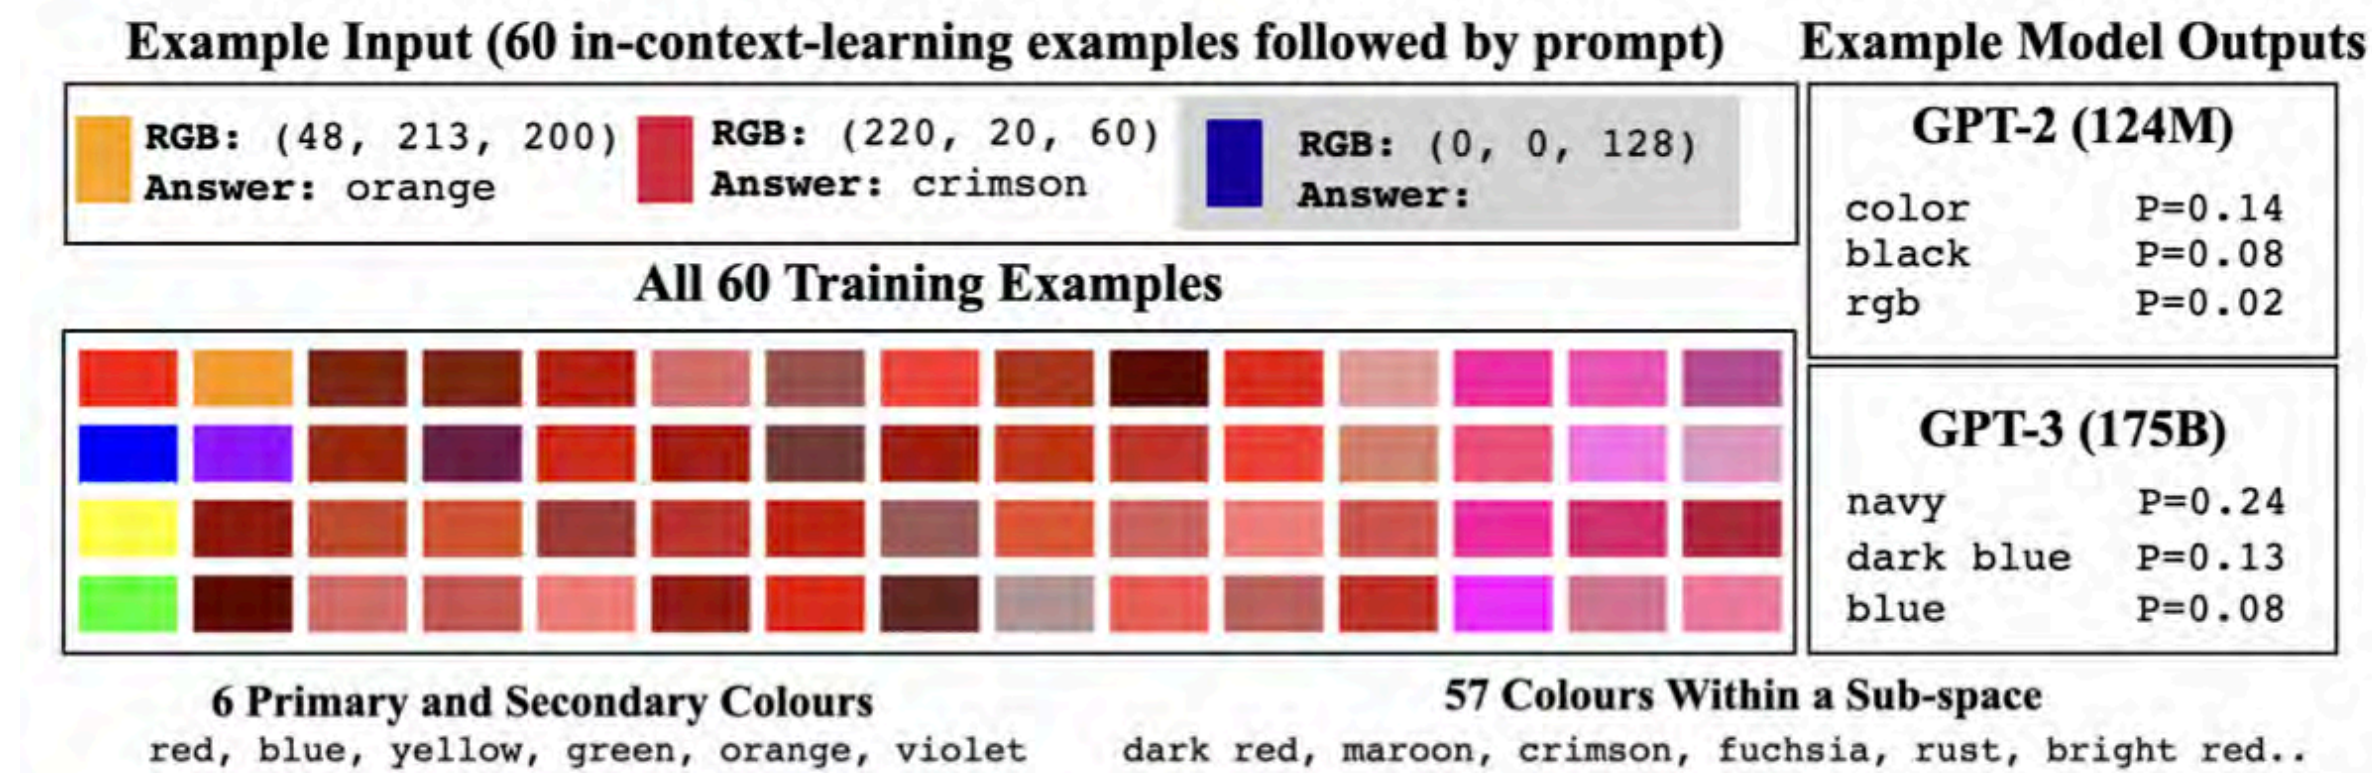


in-context learning



grounding

grounding



Mapping Language Models to Grounded Conceptual Spaces

non-multimodal LLMs

rich concept structure

(Patel et al, 2022)

Navigation Prompt:

You are on a busy road and you want to leave the main road and merge onto the ramp.



GPT-4V:

Observation: [...] The lane markings indicate that the current lane is ending or merging, as suggested by the dashed lines becoming solid and the arrow on the road pointing towards the merge. There is a sign indicating a hospital nearby, which may not directly affect driving decisions but suggests a possible increase in pedestrian and emergency vehicle activity.

Action: Reduce speed, prepare to merge right.

*On the Road with GPT-4V(ision):
Early Explorations
of Visual-Language Model on
Autonomous Driving*

GPT-4V

embodied design making

(Wen et al, 2023)

“emergent” capabilities of LLMs

*we only wanted to train next-token
predictors, how did it learn...*



coding



memory



in-context learning



grounding

two useful thinking tools



agent thinking

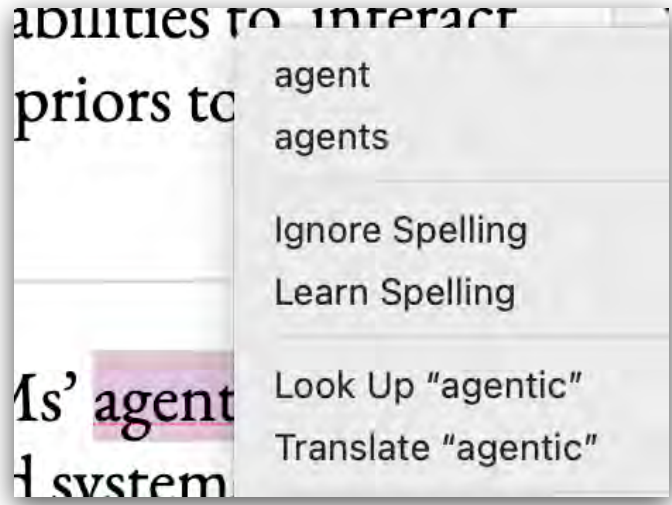
How to improve the models' capabilities to interact with the world? Are LLMs good priors to start with?



LLM Thinking

How to make full use of the LLMs' agentic capabilities? What algorithms and systems we should on top of them?

Yeah, I know. "Agentic" is weird.



what are agentic capabilities

*the capabilities that an agent needs to
have to interact with the world*



perception



planning



agency



learning

what are agentic capabilities

the capabilities that an agent needs to have to interact with the world



perception



agency

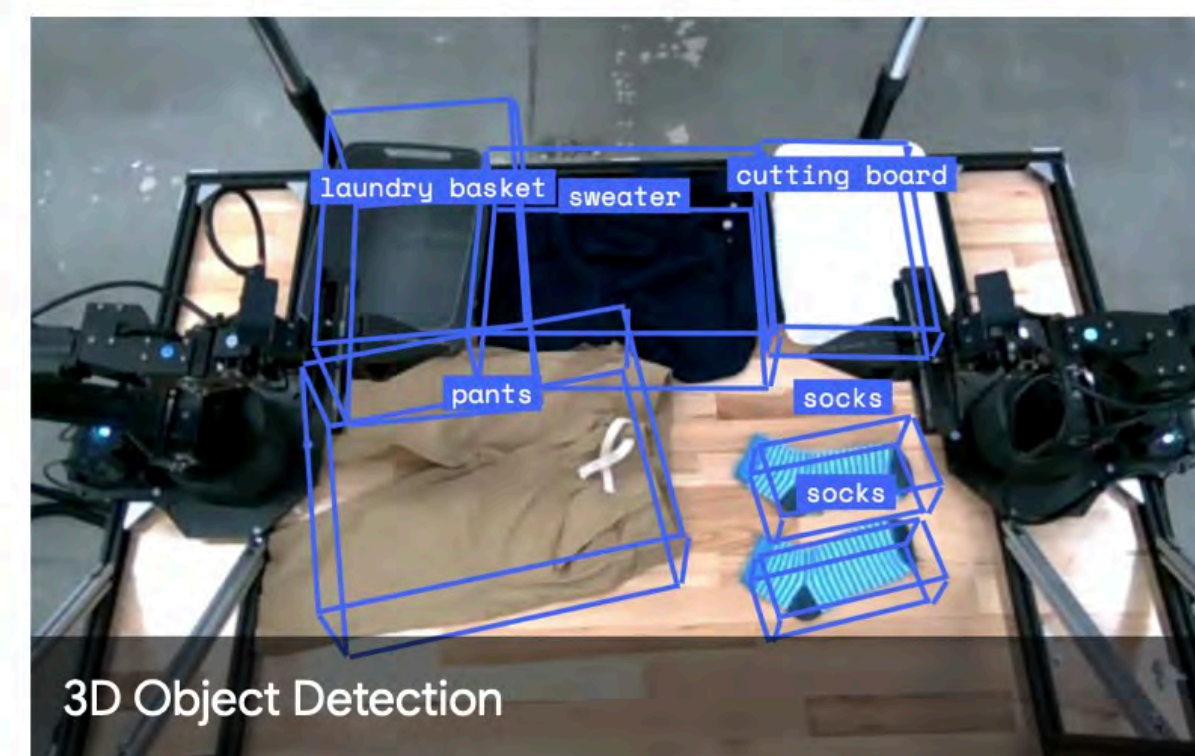
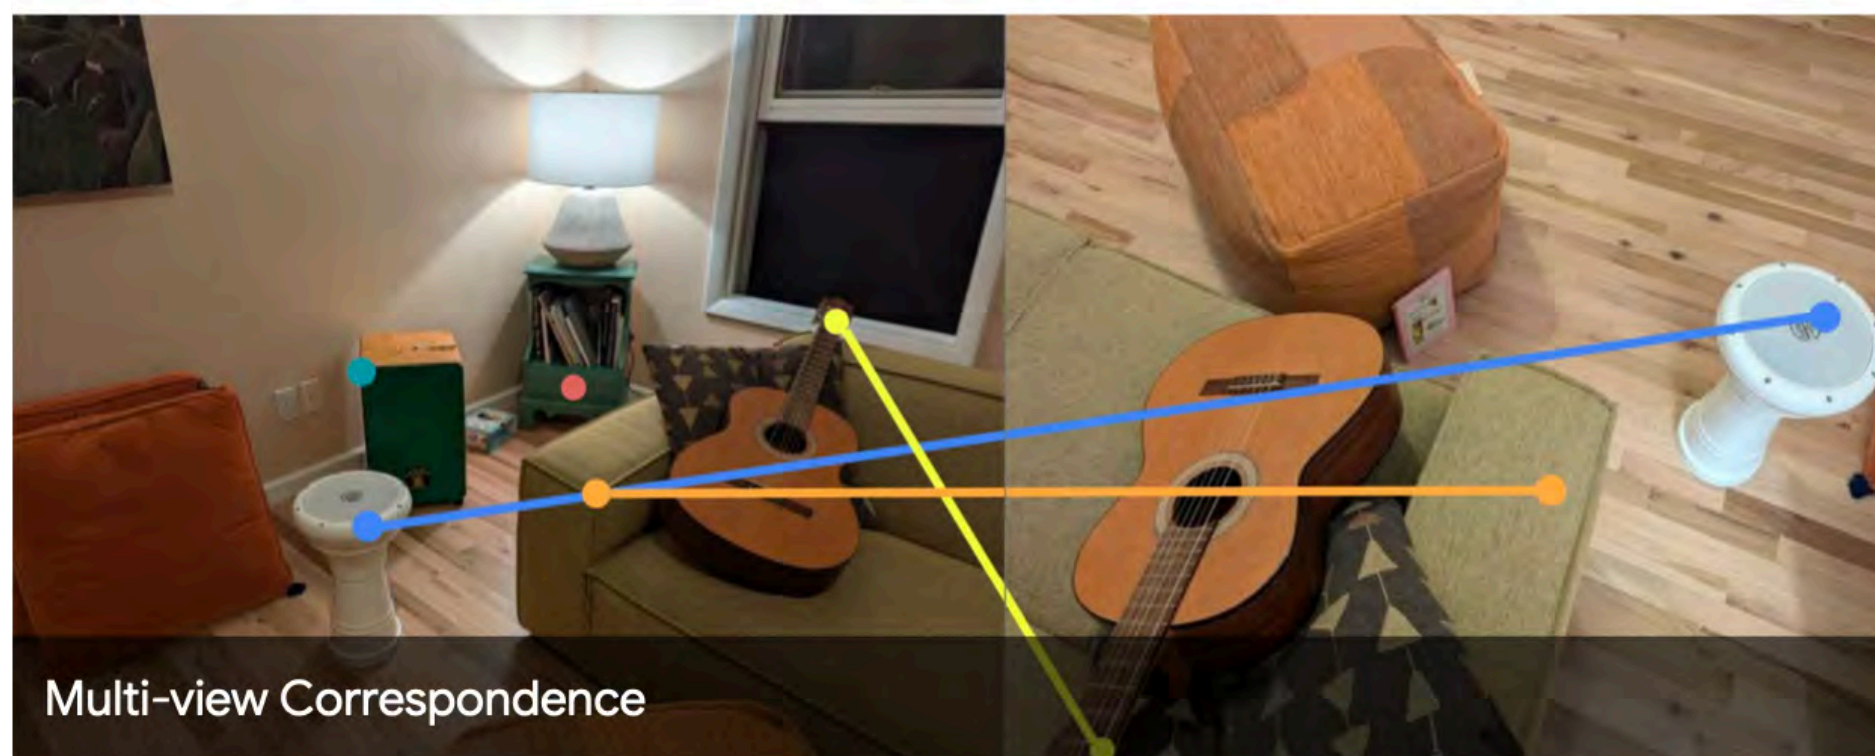
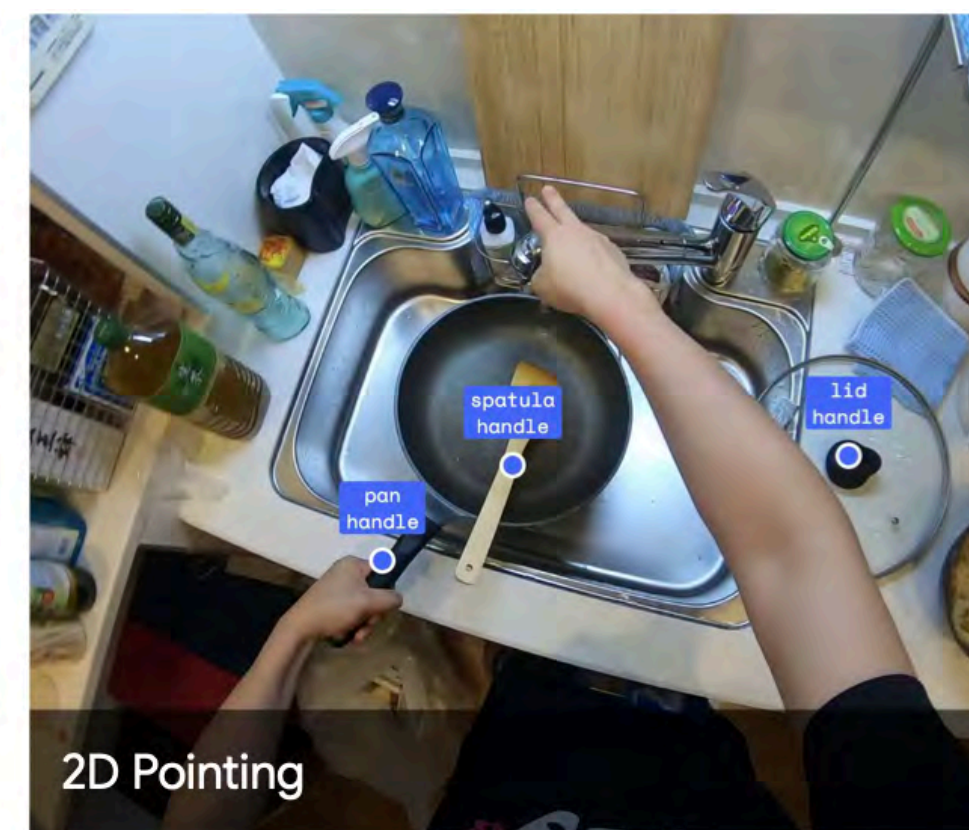
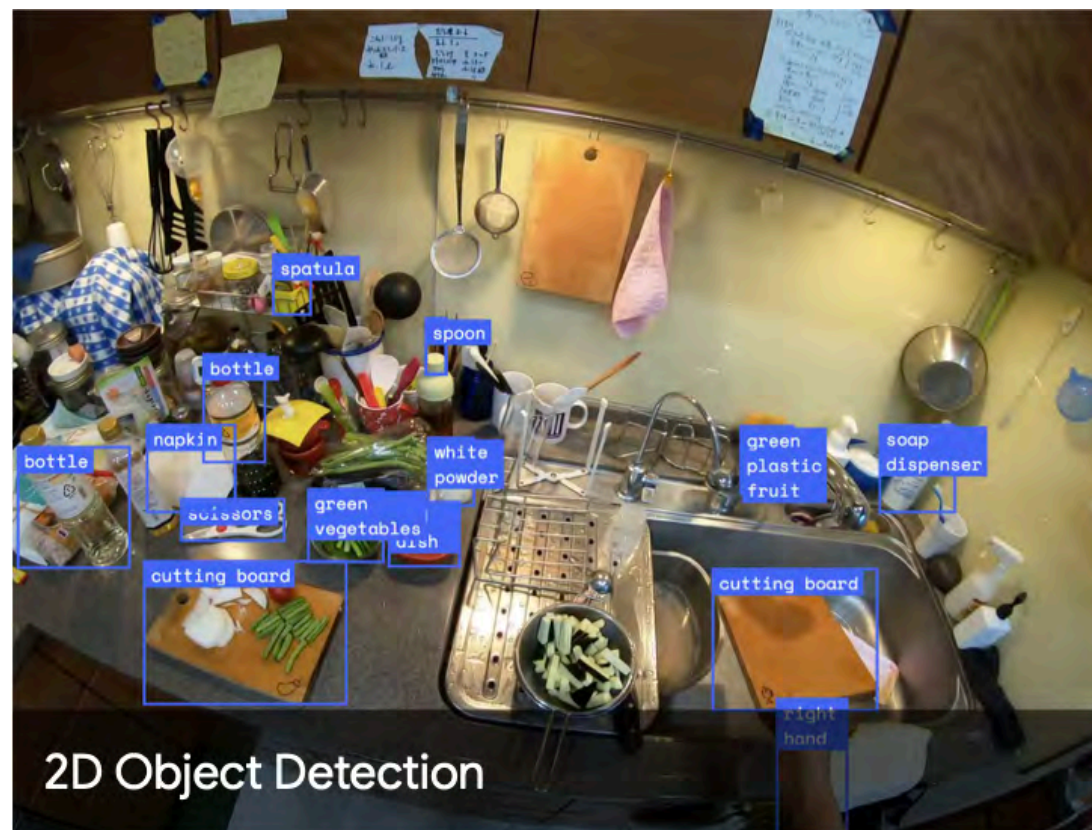


planning



learning

perception



**Gemini-
Robotics**
Understanding
the semantic
structure of
observation
(Google, 2025)

Gemini Robotics: Bringing AI into the Physical World

what are agentic capabilities

*the capabilities that an agent needs to
have to interact with the world*



perception



planning



agency



learning

what are agentic capabilities

the capabilities that an agent needs to have to interact with the world



perception



agency

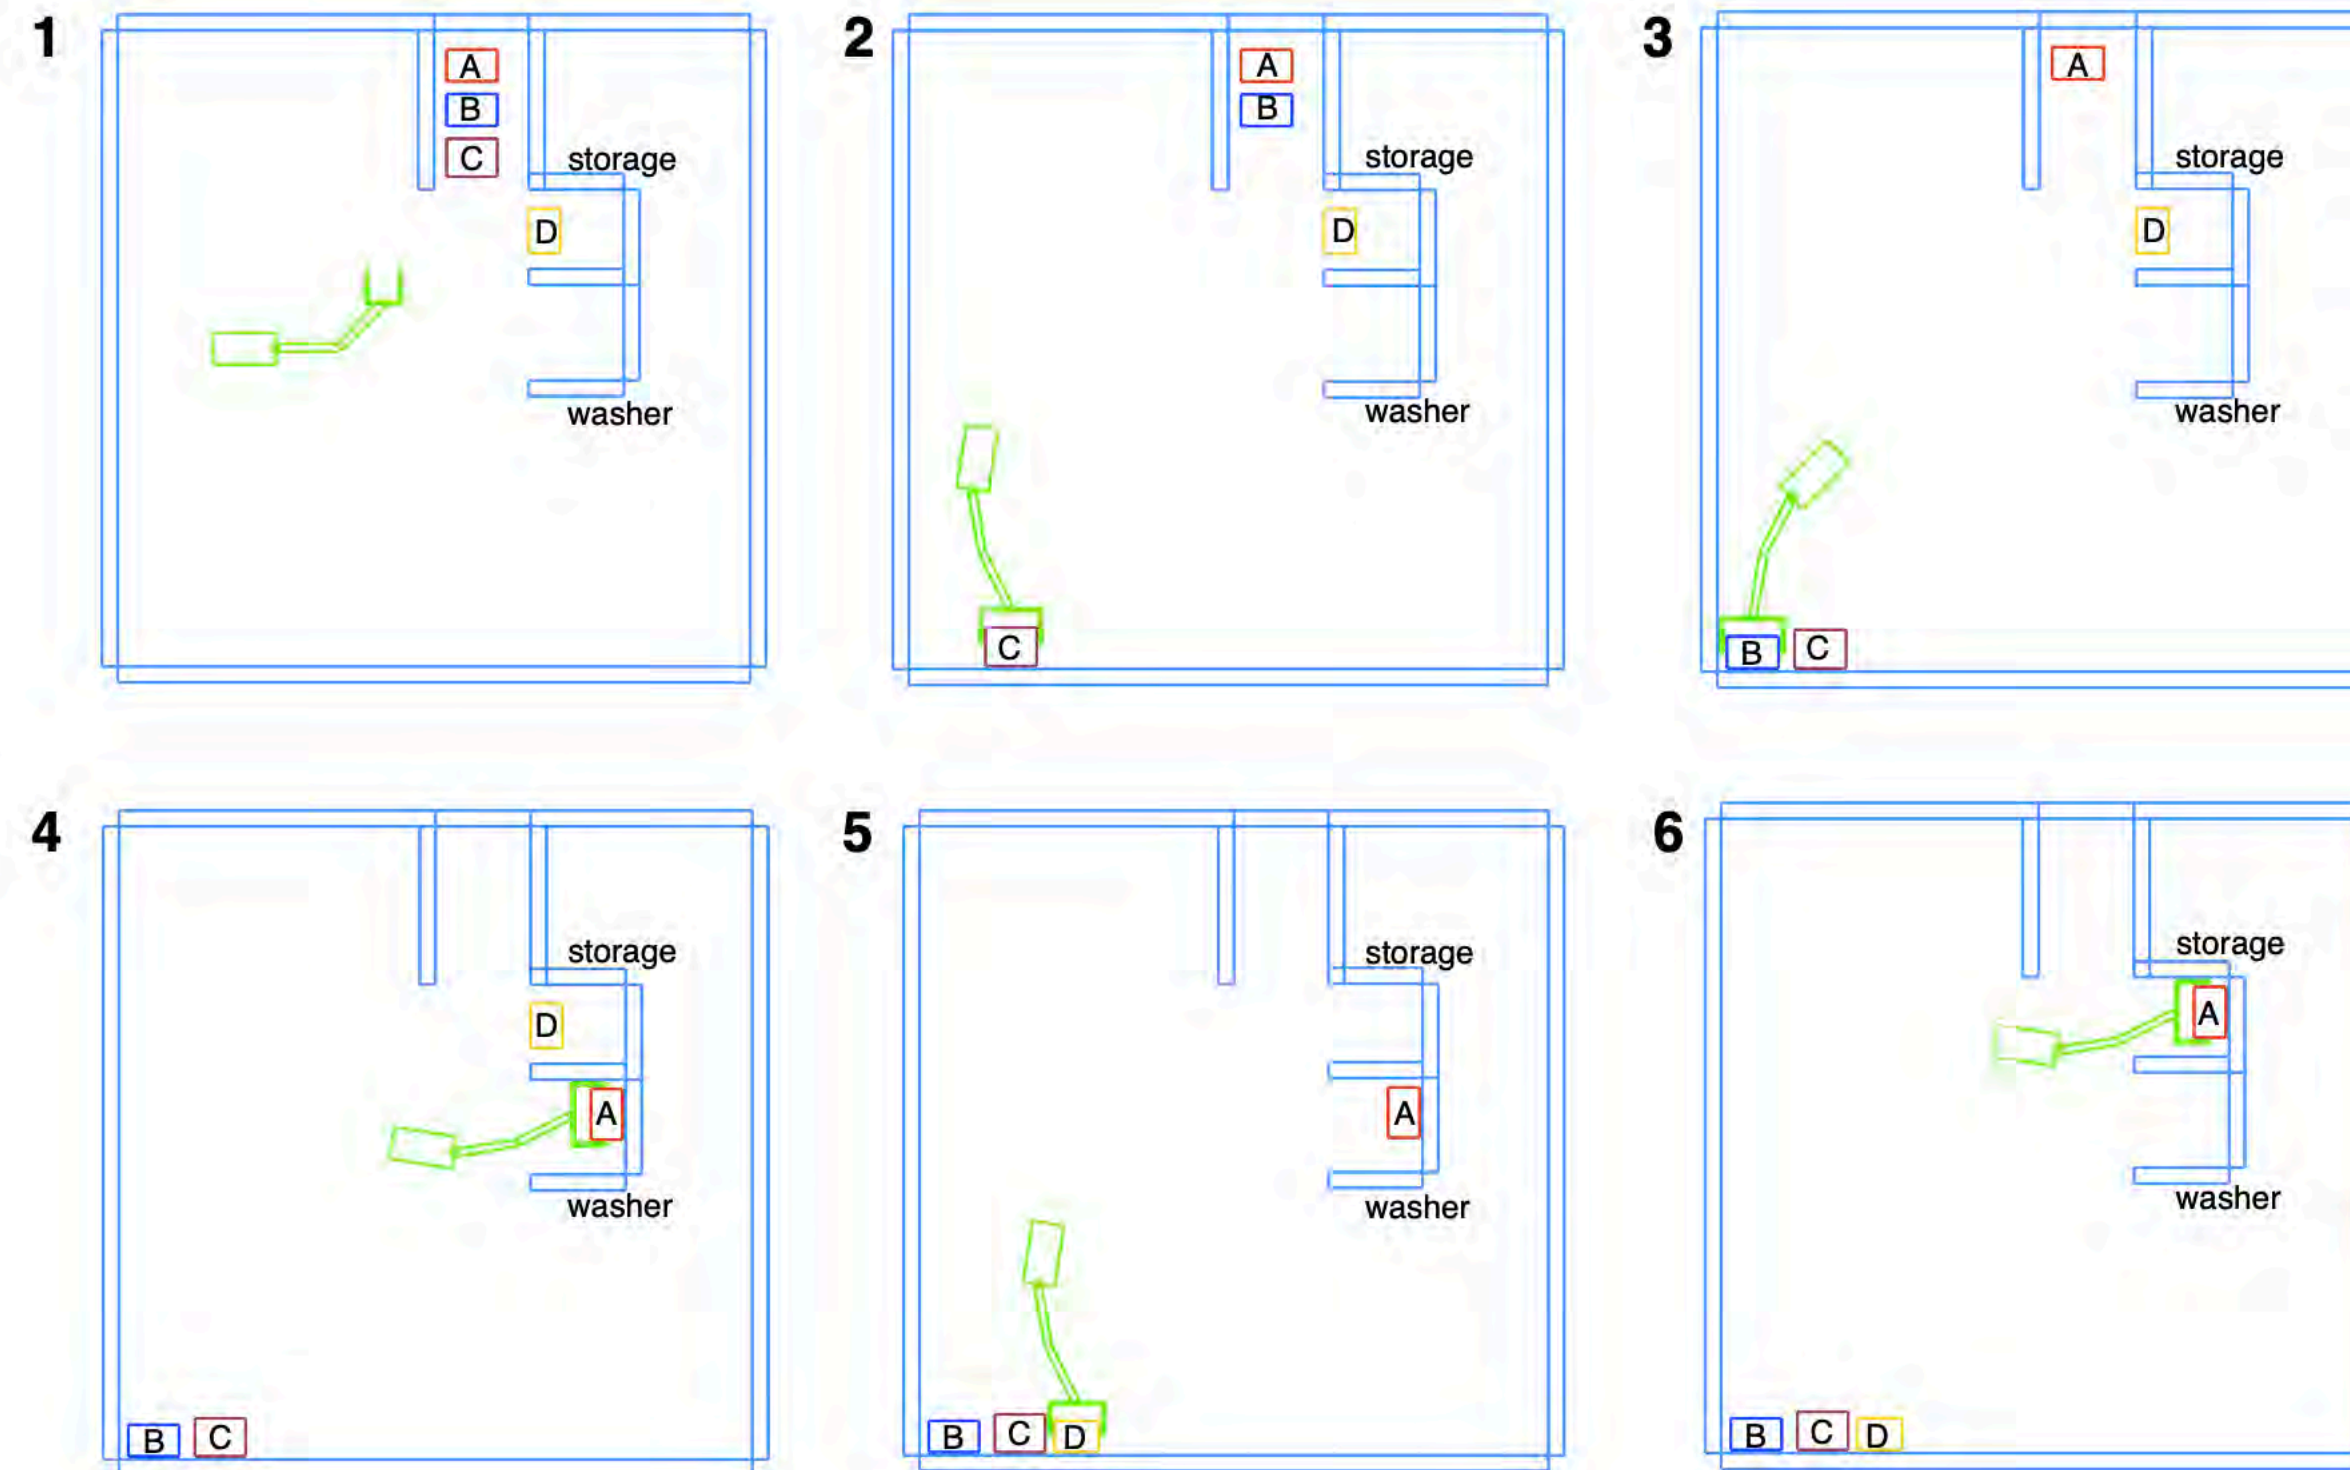


planning



learning

planning



Hierarchical task and motion planning in the now

Task and Motion Planning

Top-down decomposition w/ bottom-up constraints

(Kaelbling and Lozano-Pérez, 2011)

what are agentic capabilities

the capabilities that an agent needs to have to interact with the world



perception



agency



planning



learning

what are agentic capabilities

the capabilities that an agent needs to have to interact with the world



perception



planning

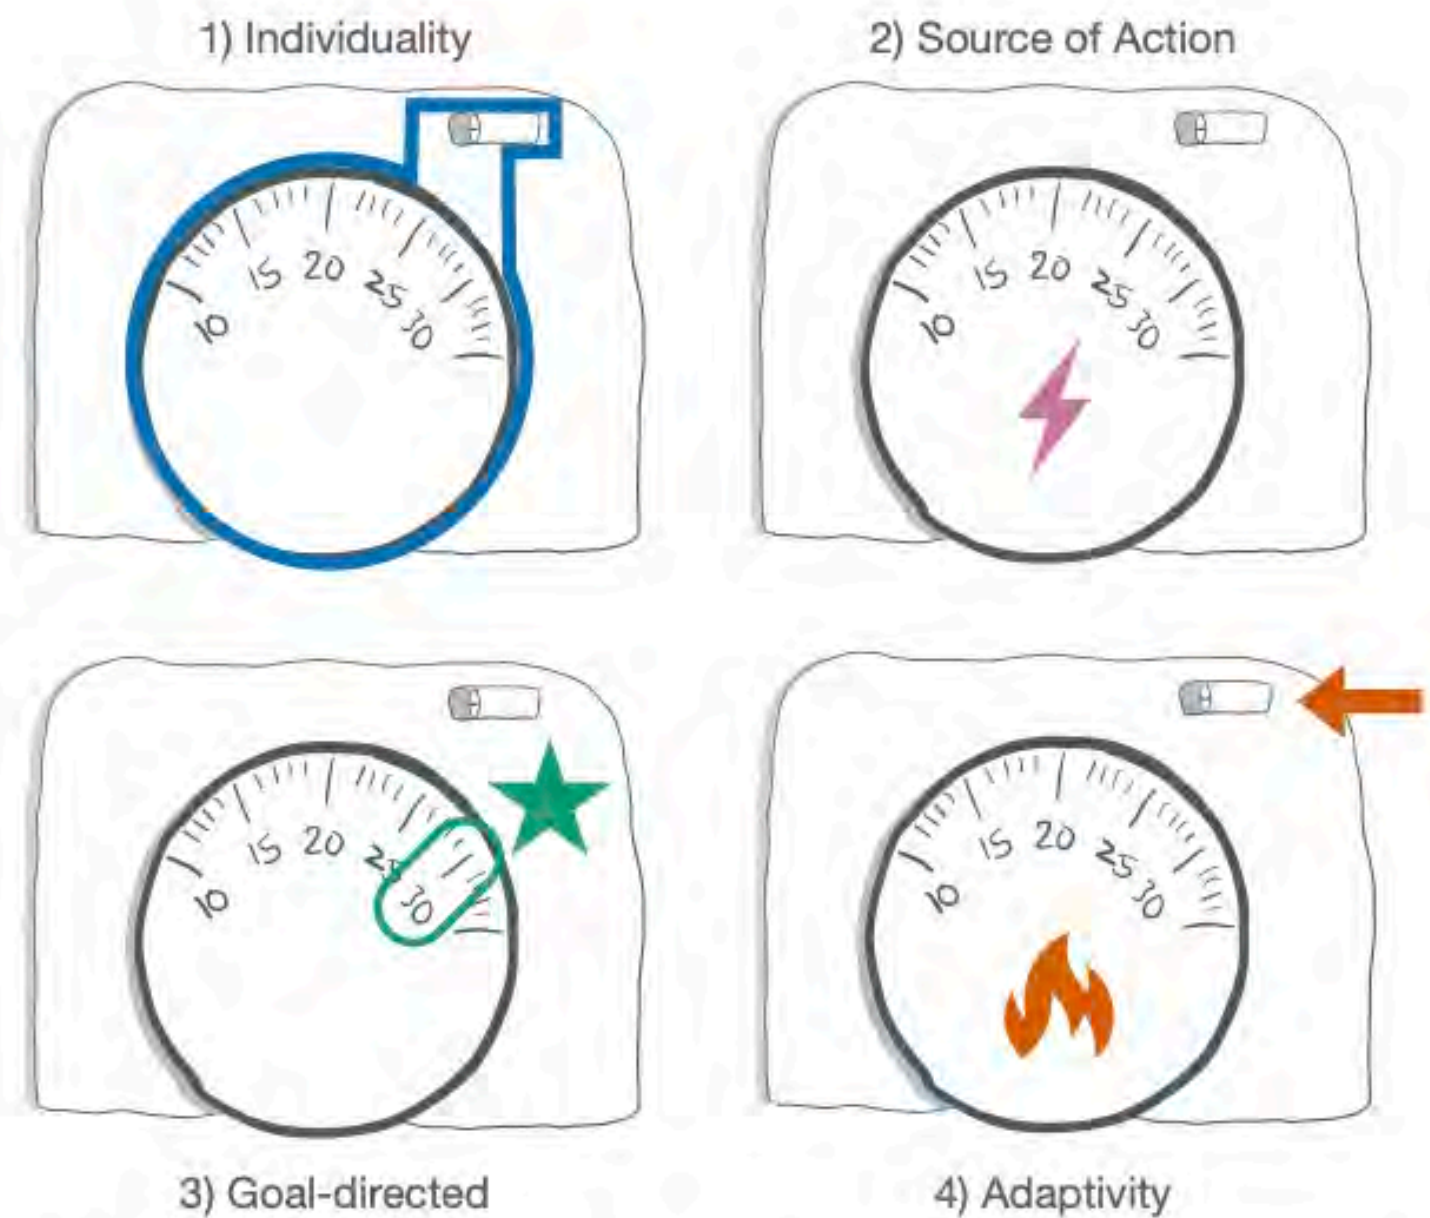


agency



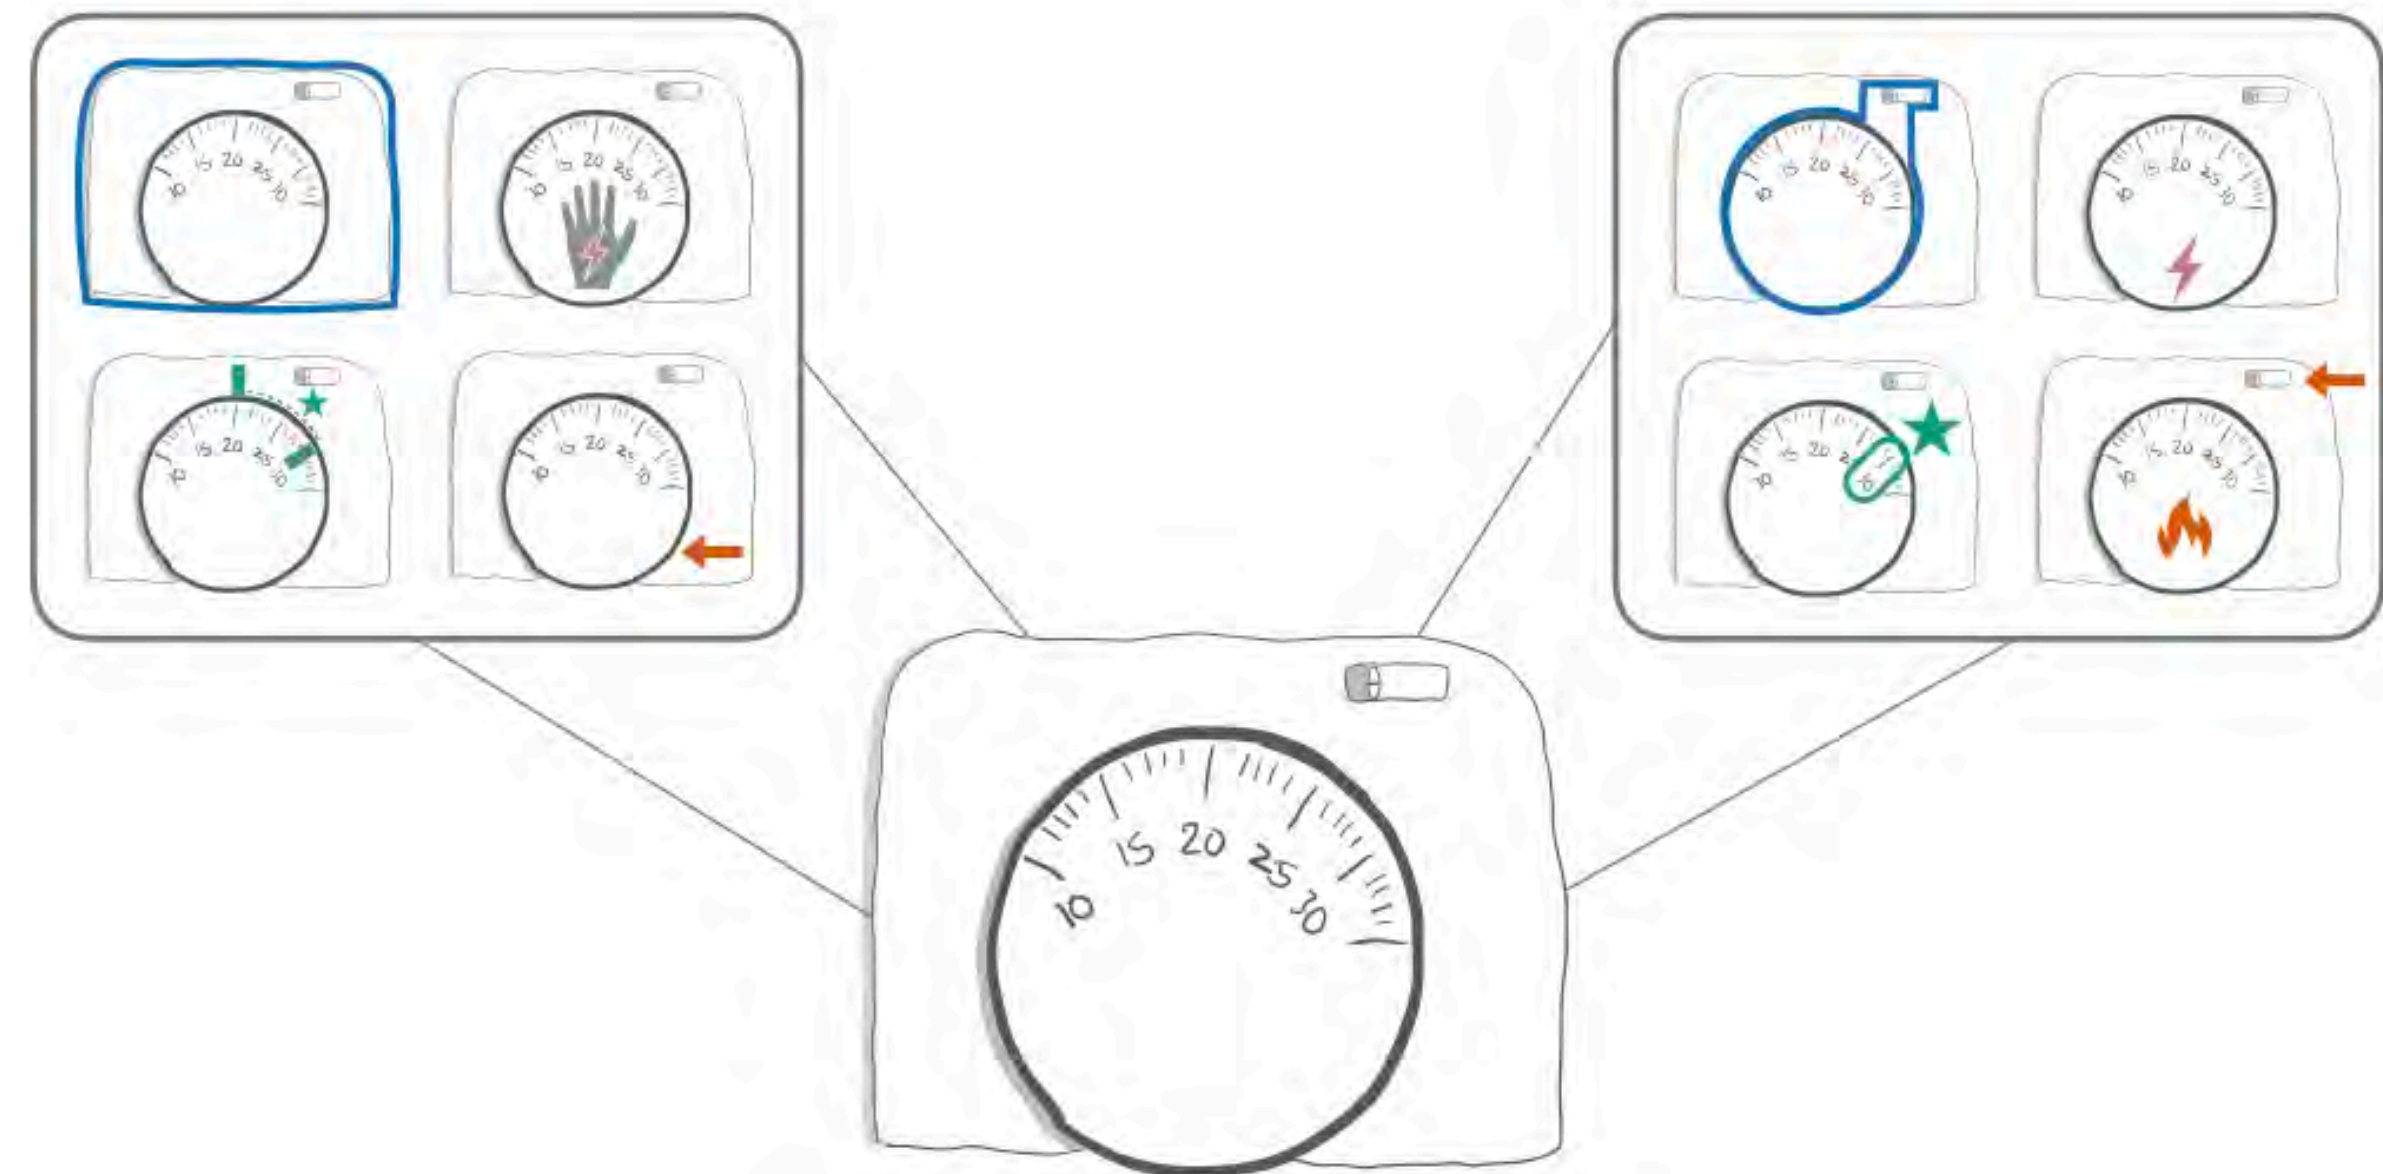
learning

agency — the essence but hard to define



(a) A Four-Part Account of Agency

Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action.



(b) Frame-Dependence

Agency Is Frame-Dependent

what are agentic capabilities

the capabilities that an agent needs to have to interact with the world



perception



planning



agency



learning

what are agentic capabilities

the capabilities that an agent needs to have to interact with the world



perception



planning

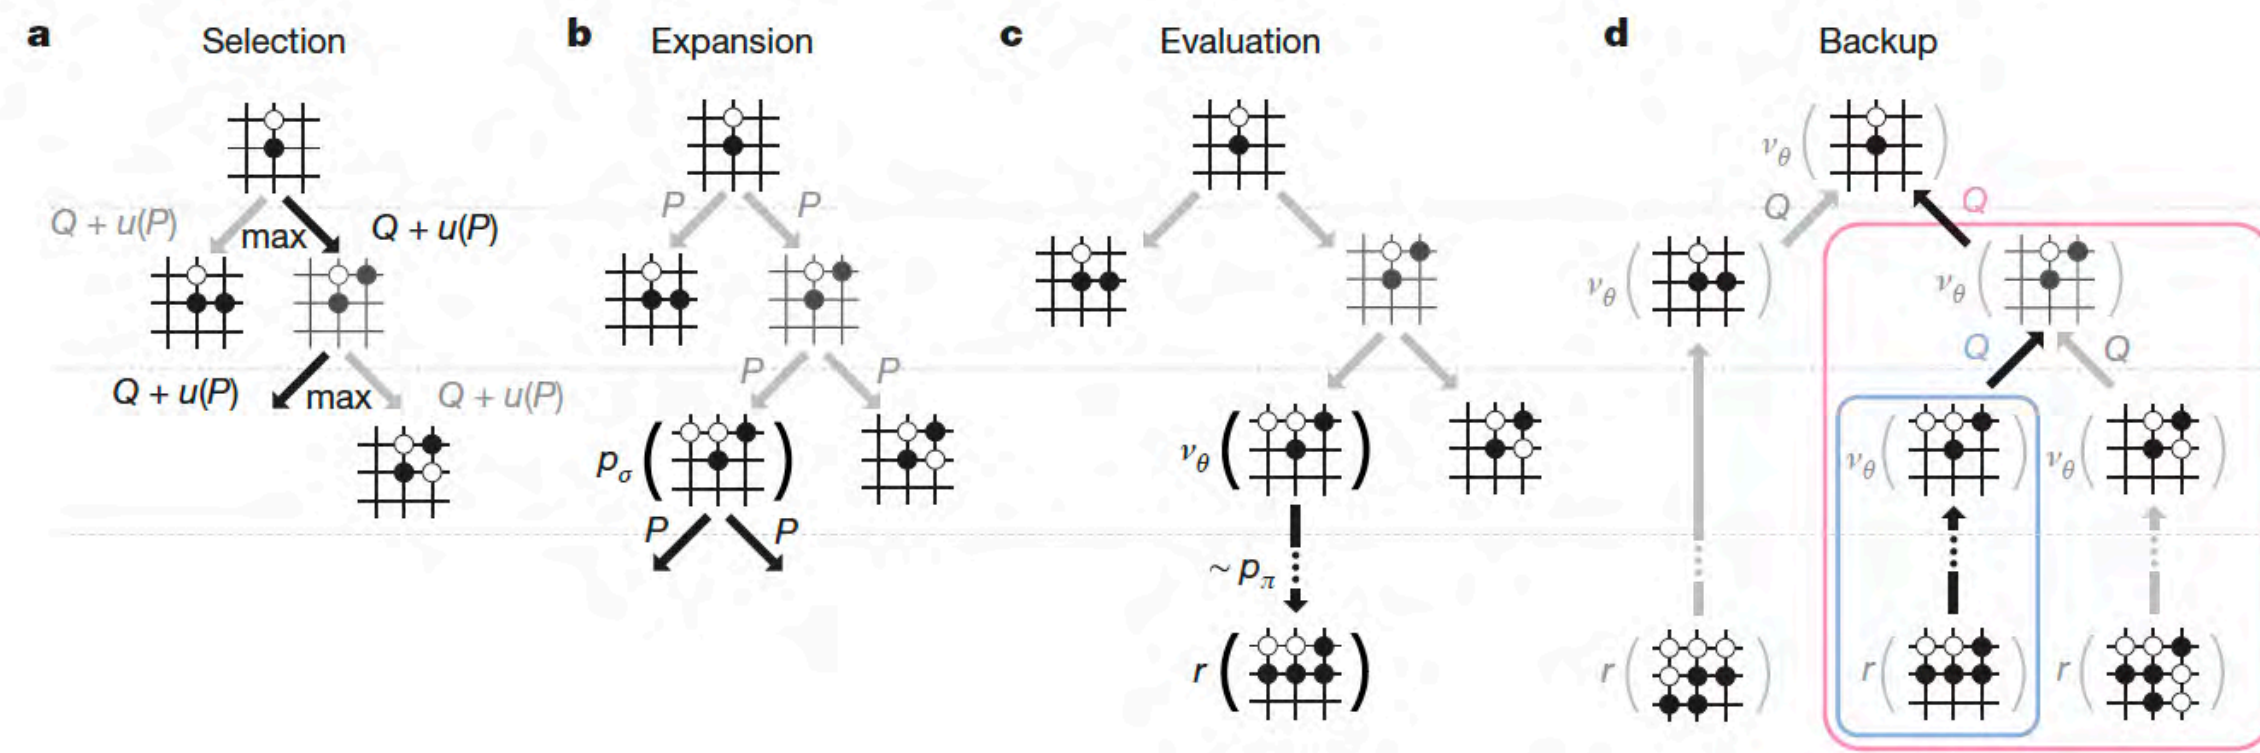


agency



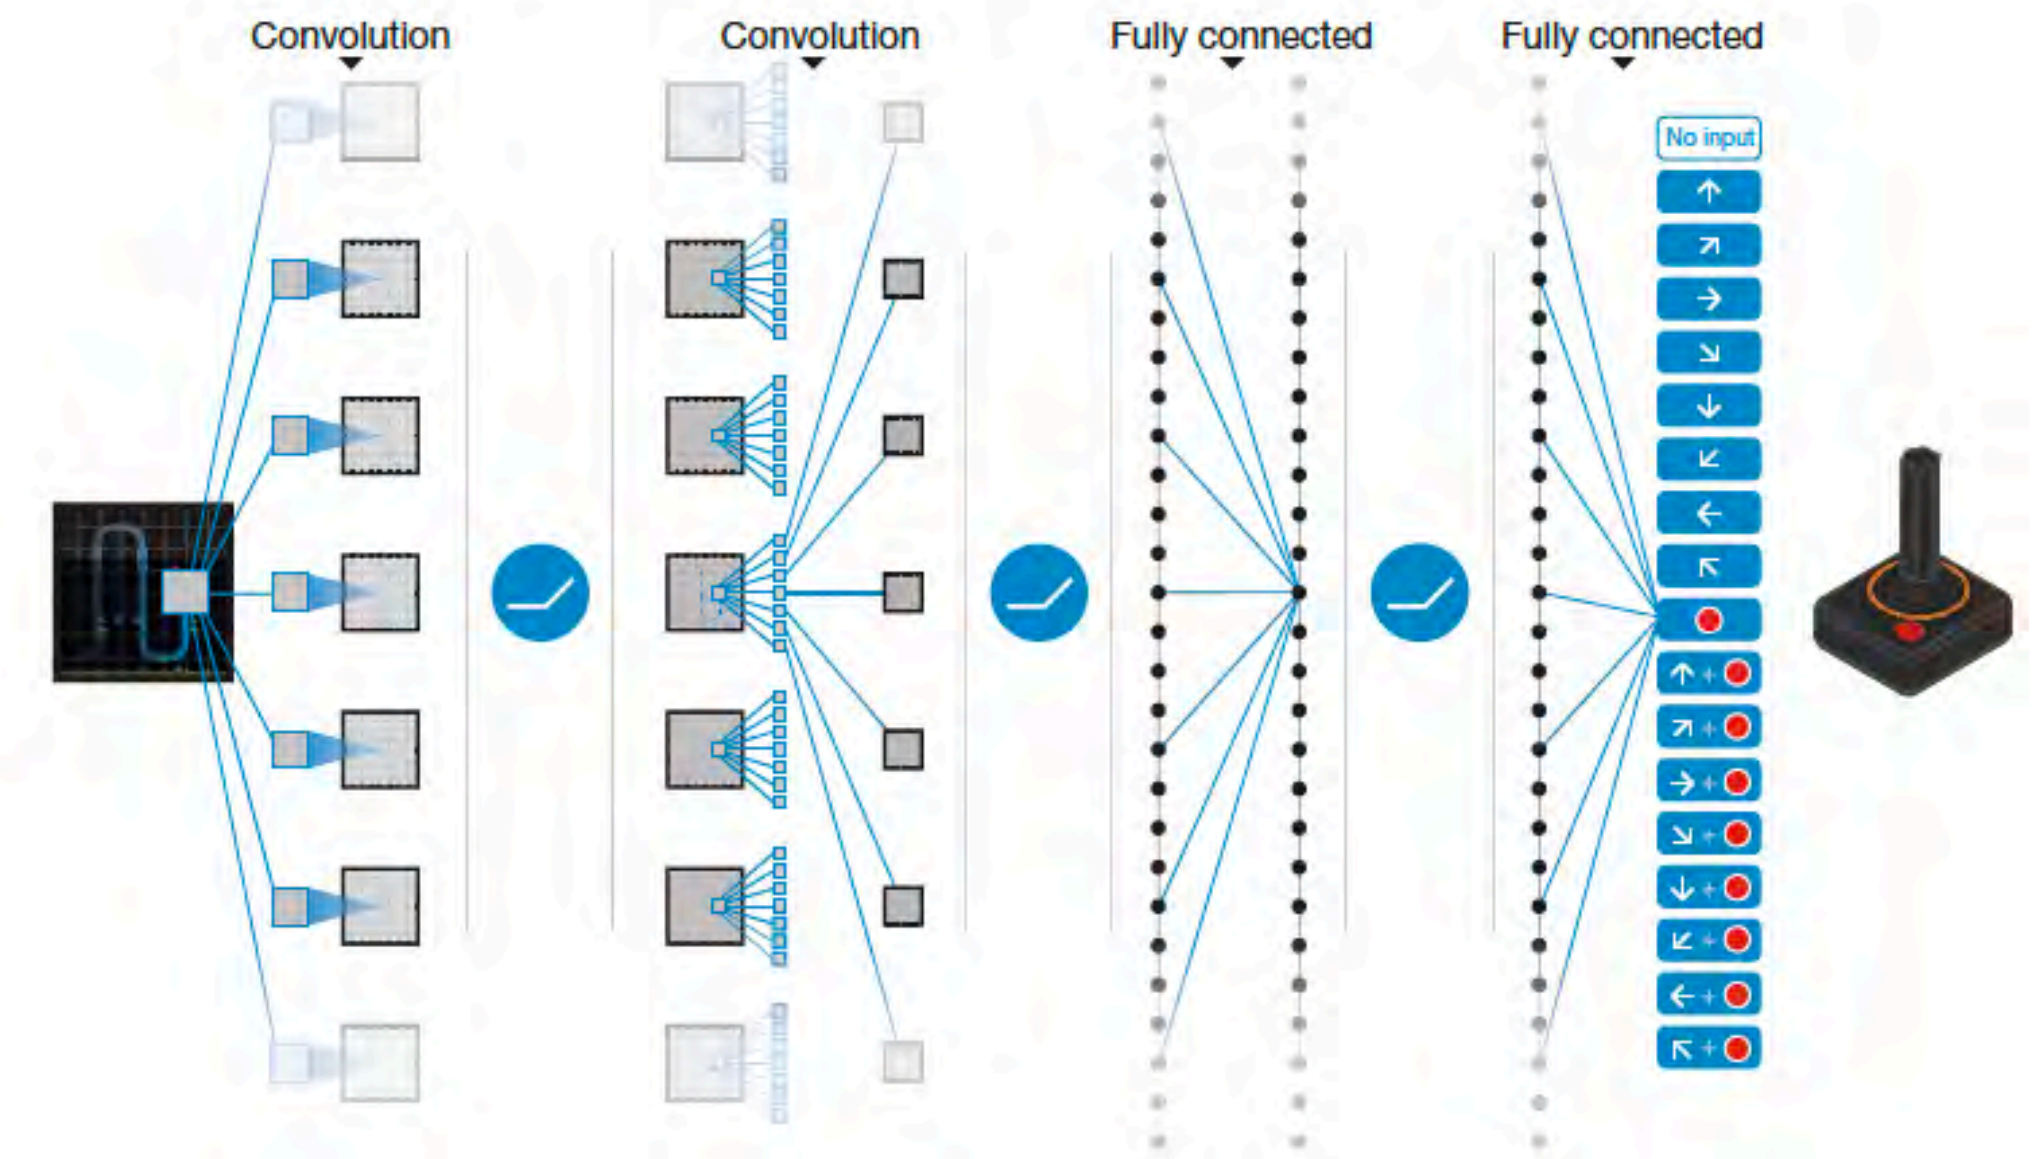
learning

learning



learning through search

Mastering the game of Go with deep neural networks and tree search



learning through RL

Human-level control through deep reinforcement learning

what are agentic capabilities

the capabilities that an agent needs to have to interact with the world



perception



planning



agency



learning

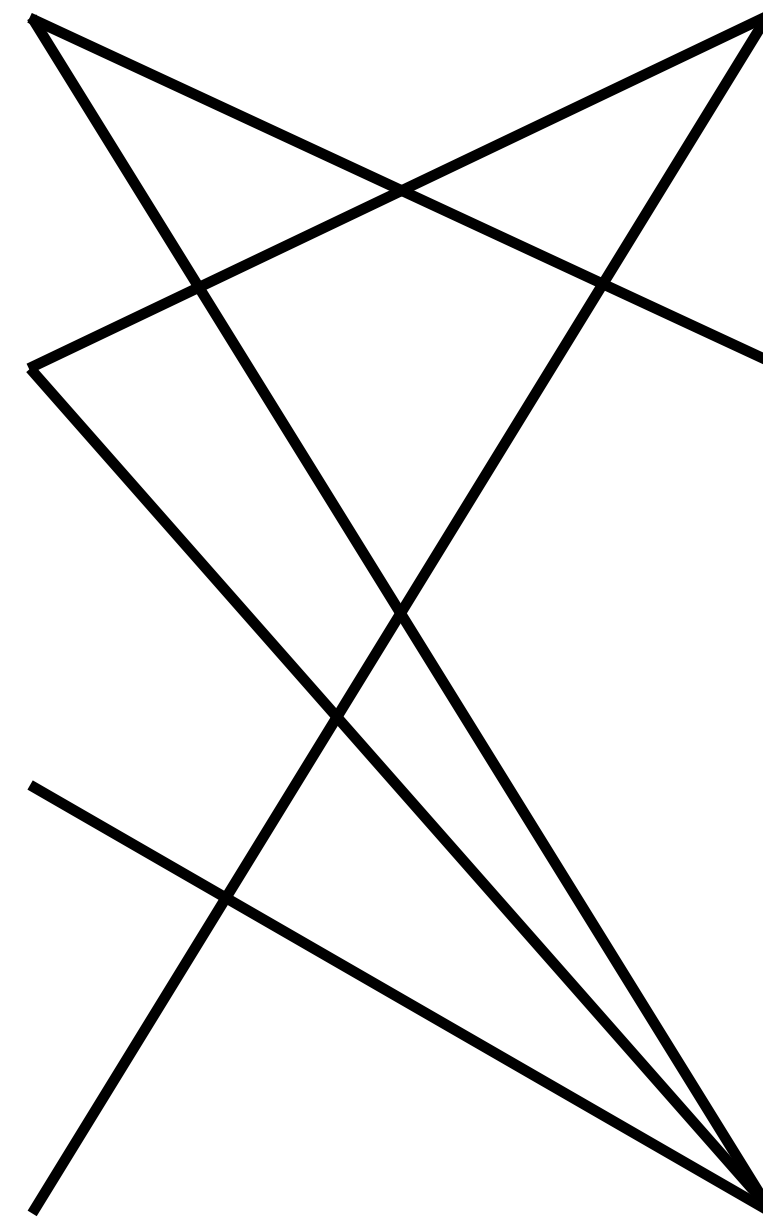
are LLMs good priors?

 **coding**

 **memory**

 **ICL**

 **grounding**



 **perception**

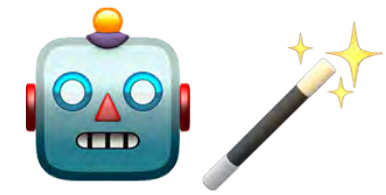
 **planning**

 **agency** ?

 **learning**

the **agentic** capabilities

short summary



agent thinking

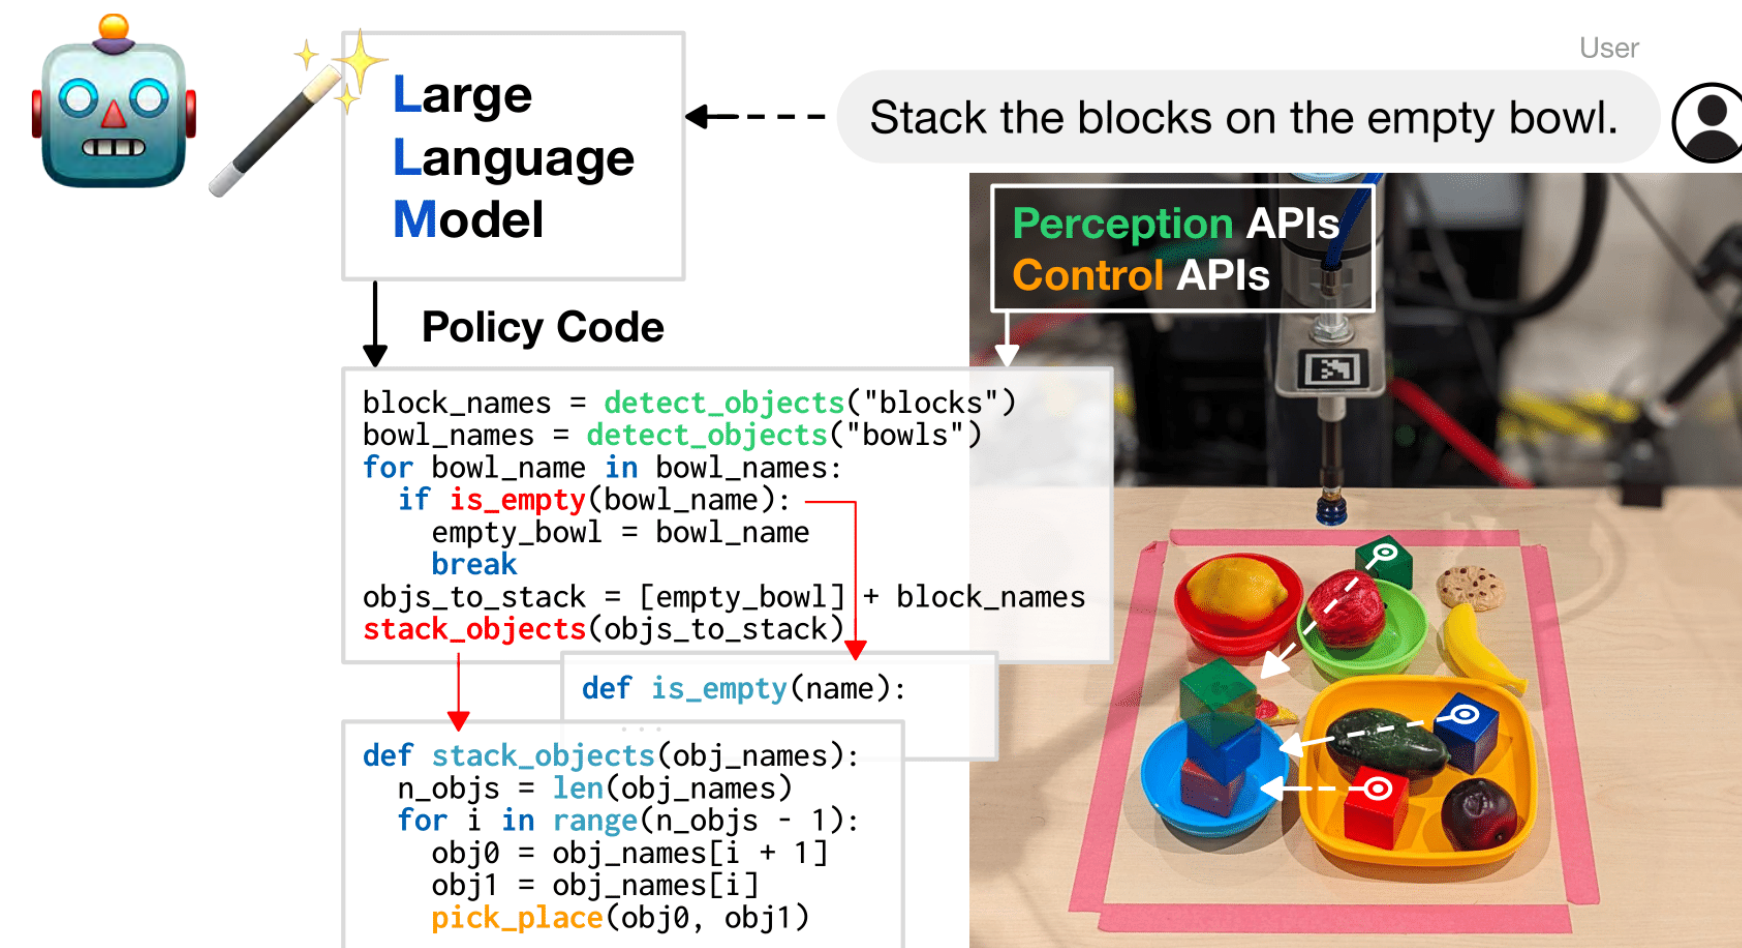
LLMs provide good priors for the capabilities that we consider as essential for agents, so we should start from there.



LLM Thinking

We could think from these agentic capabilities perspective when building agents.

coding for planning/learning

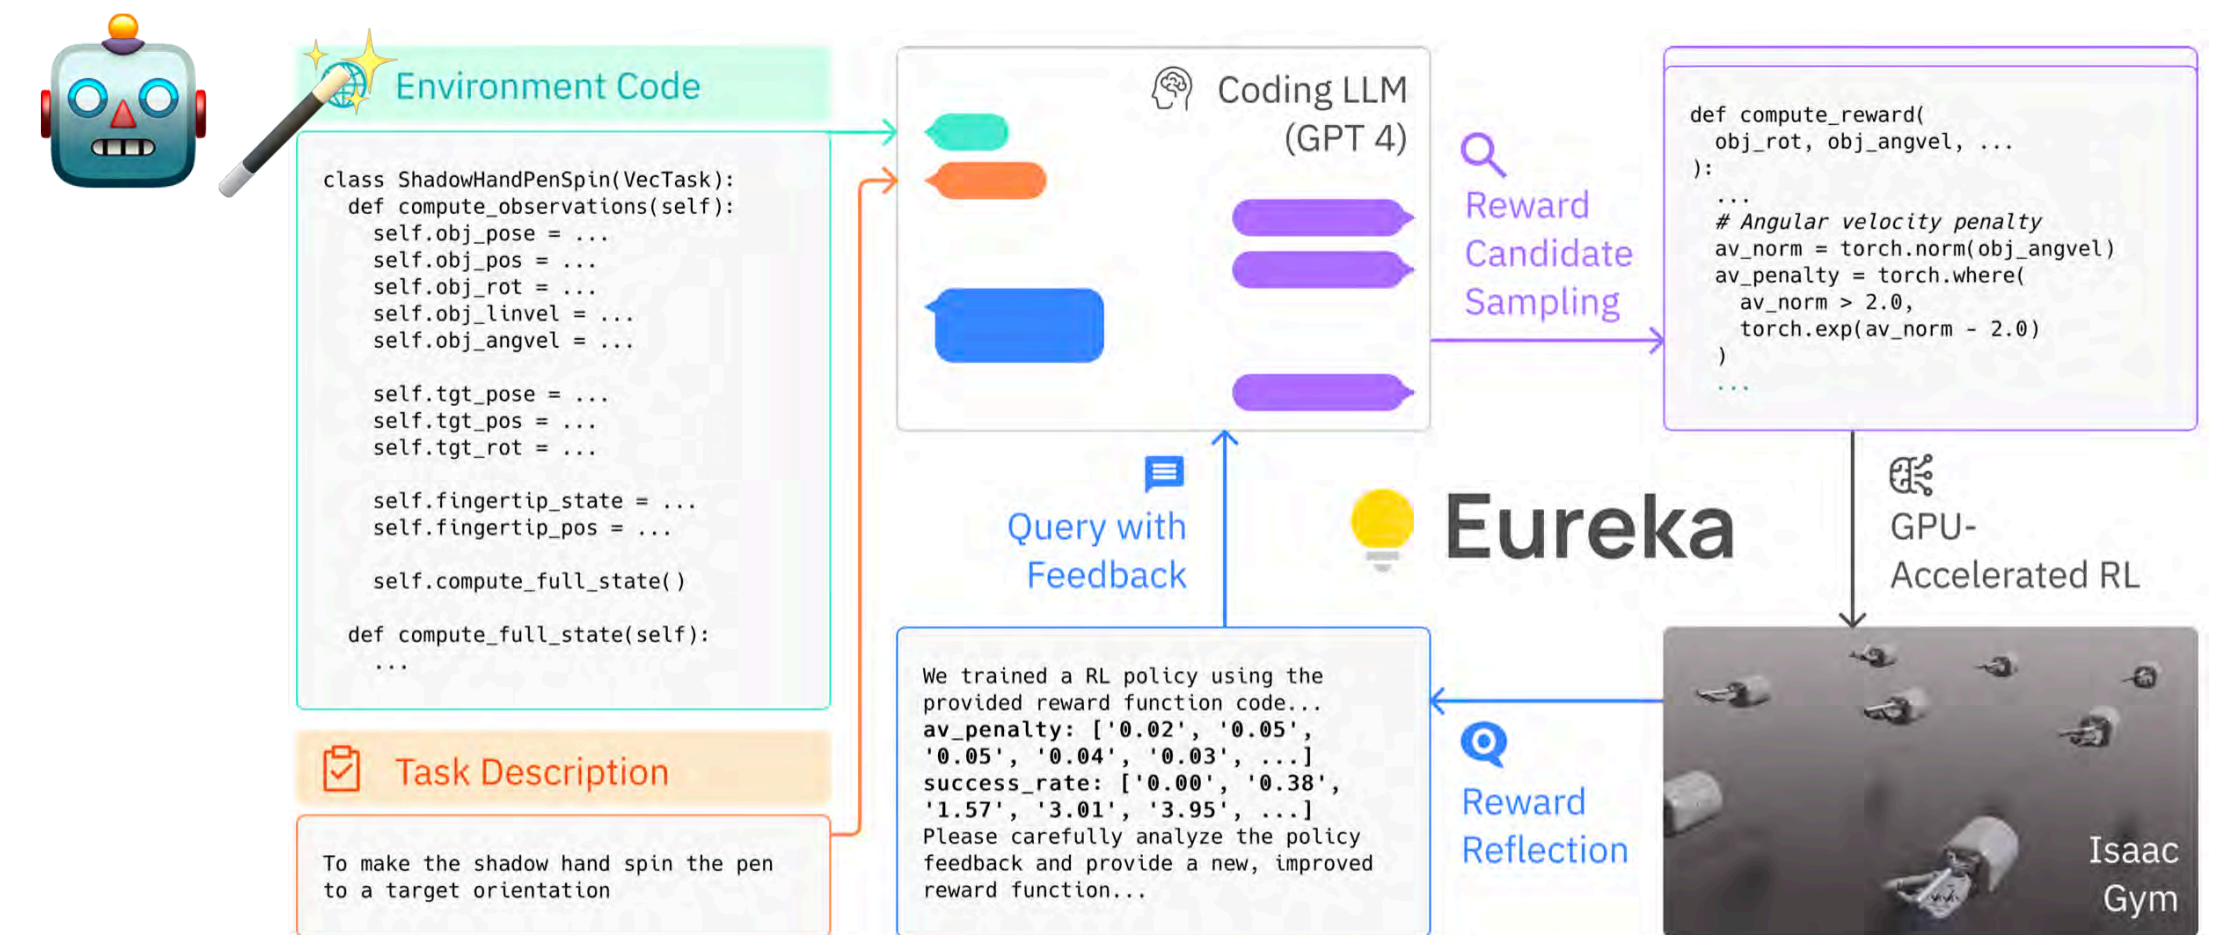


Code as Policies: Language Model Programs for Embodied Control

code as policy

control robot w/o eyes

(Liang et al, 2022)



Eureka: Human-Level Reward Design via Coding Large Language Models

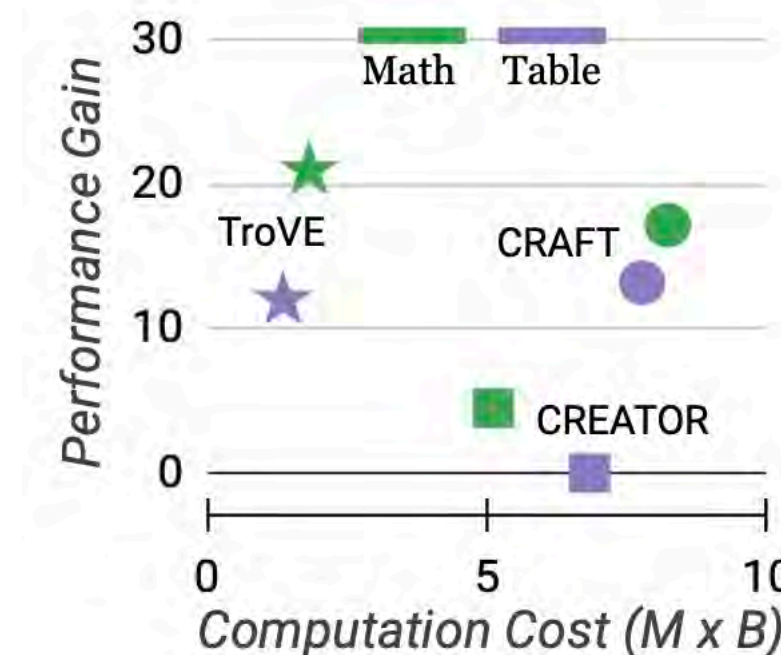
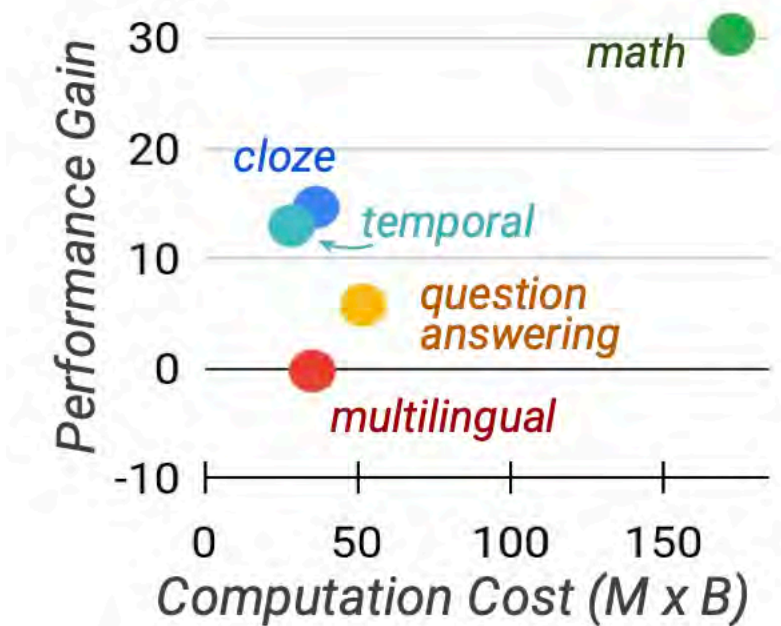
code as reward

prompt to code for QA

(Ma et al, 2024)

coding for planning

Benchmark	Tool Source	Example Curation	Domain (§4.1)	Executable
ToolBench ₁	existing dataset	adopted, human annotated	📁, 🌐	✓
ToolBench ₂	RapidAPI	model synthesized	📁, 🌐	✓
ToolQA	existing dataset	model synthesized	📁, 📖	✓
ToolAlpaca	PublicAPIs	model synthesized	📖, 📁, 🌐, 🎬	✗
API-Bank	PublicAPIs	human annotated	📁, 🌐	✓
MetaTool	OpenAI Plugins	model synthesized	📁, 🌐, 🎬	✗
Gorilla	HF, Torch, TF	model synthesized	👤	✗
HuggingGPT	HF	human annotated	👤	✗*
Task Bench	HF, PublicAPIs	model synthesized	👤, 🎬, 🌐	✗

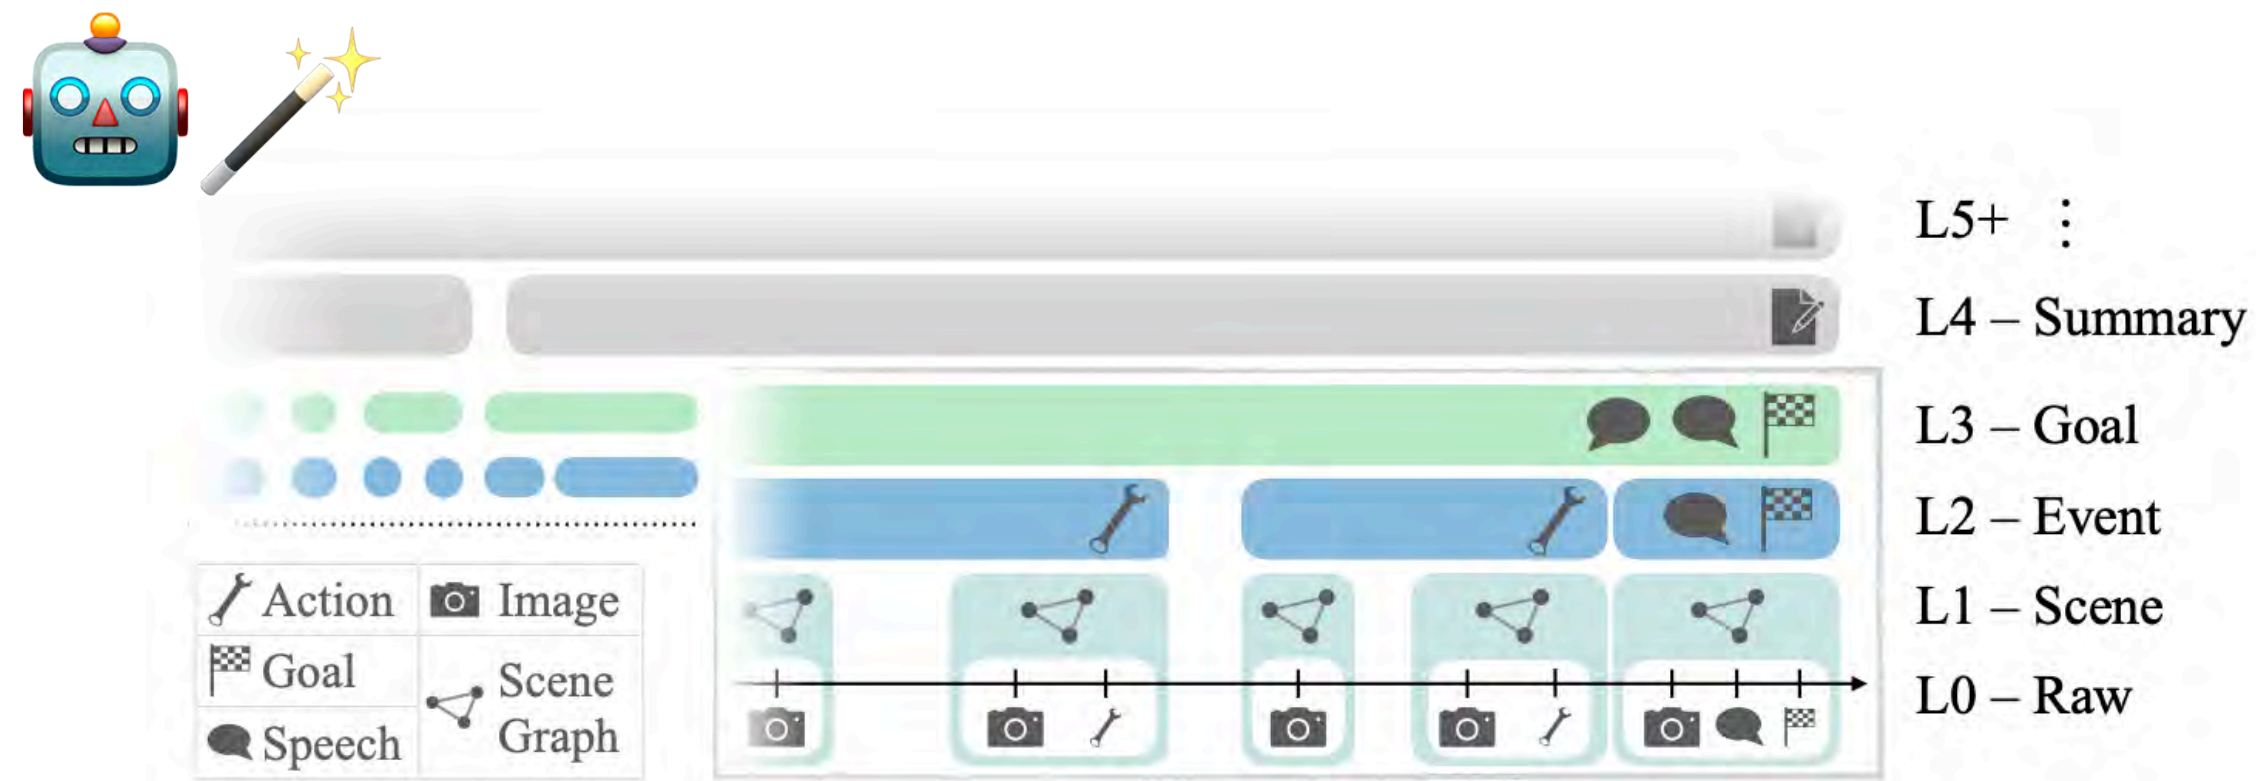


Tool use / make cases

when (not) to use tools

(Wang et al, 2024)

memory for perception

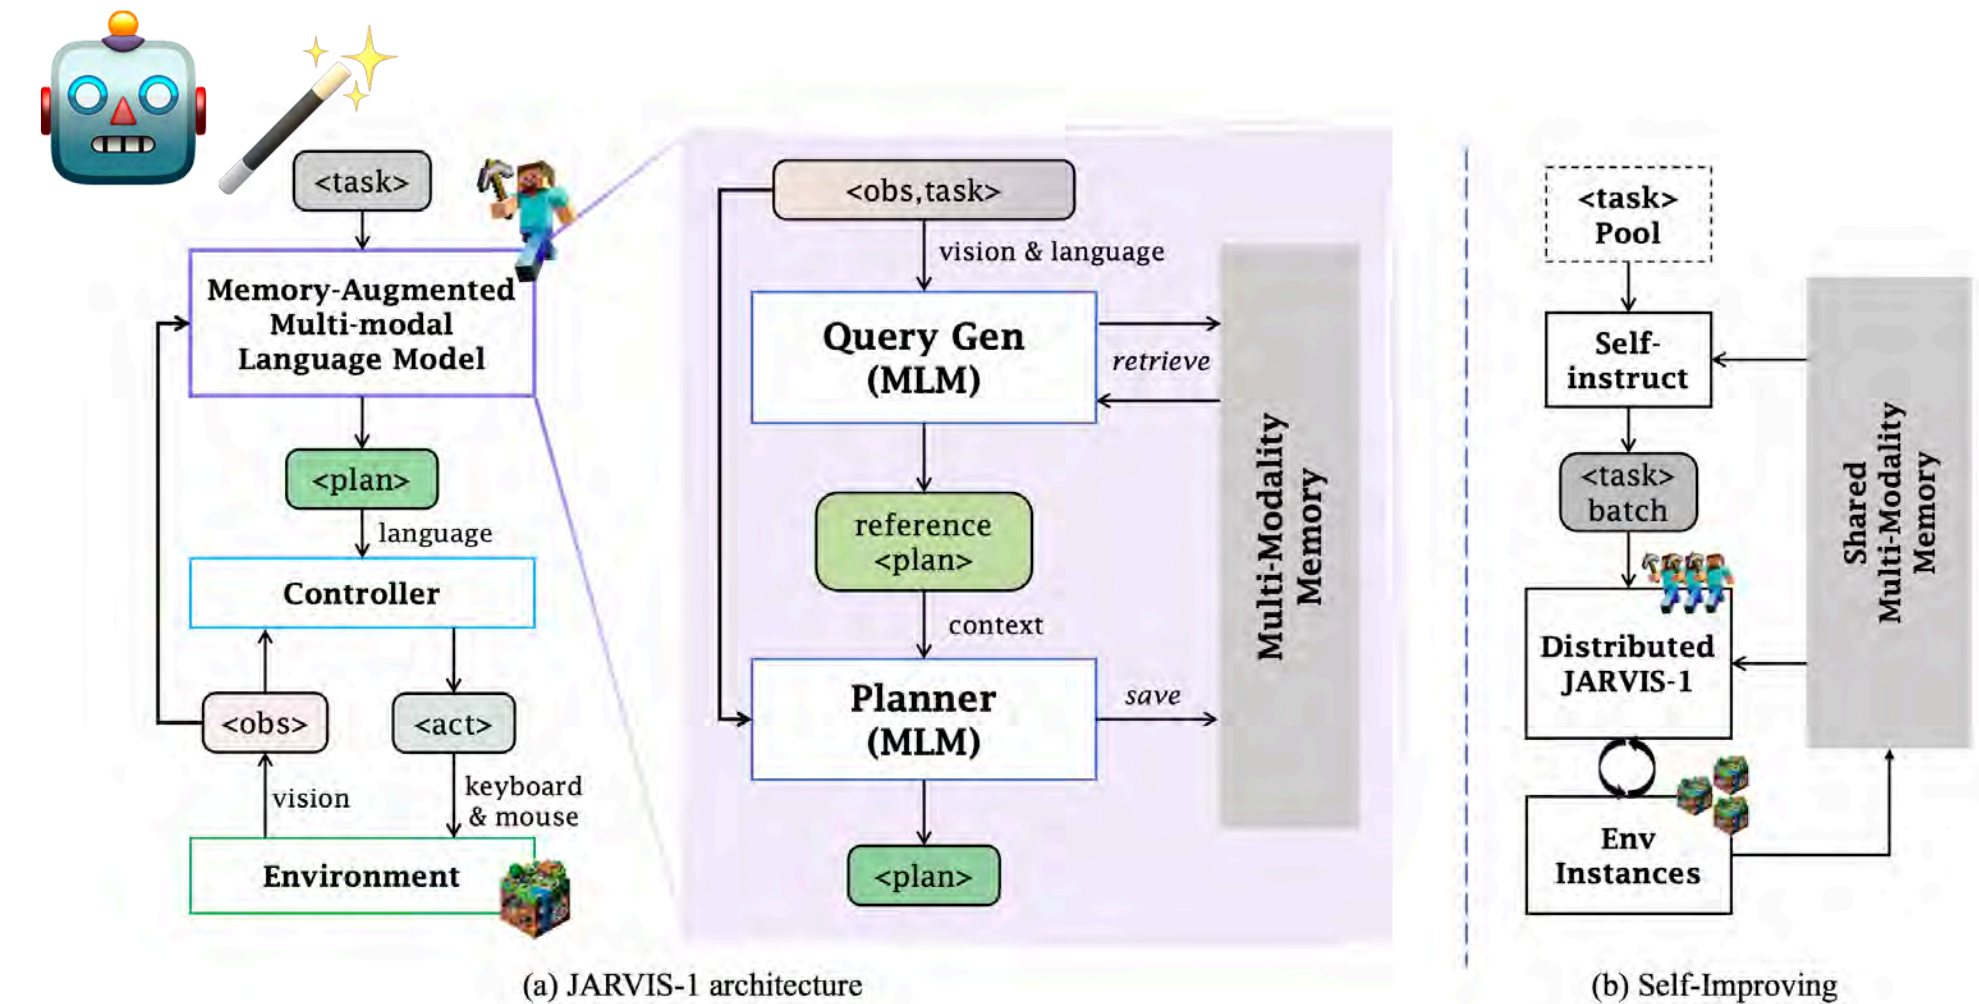


Episodic Memory Verbalization using Downscaled Hierarchical Representations of Life-Long Robot Experience

hierarchical memory

long memory robotics QA

(Bärmann et al, 2024)



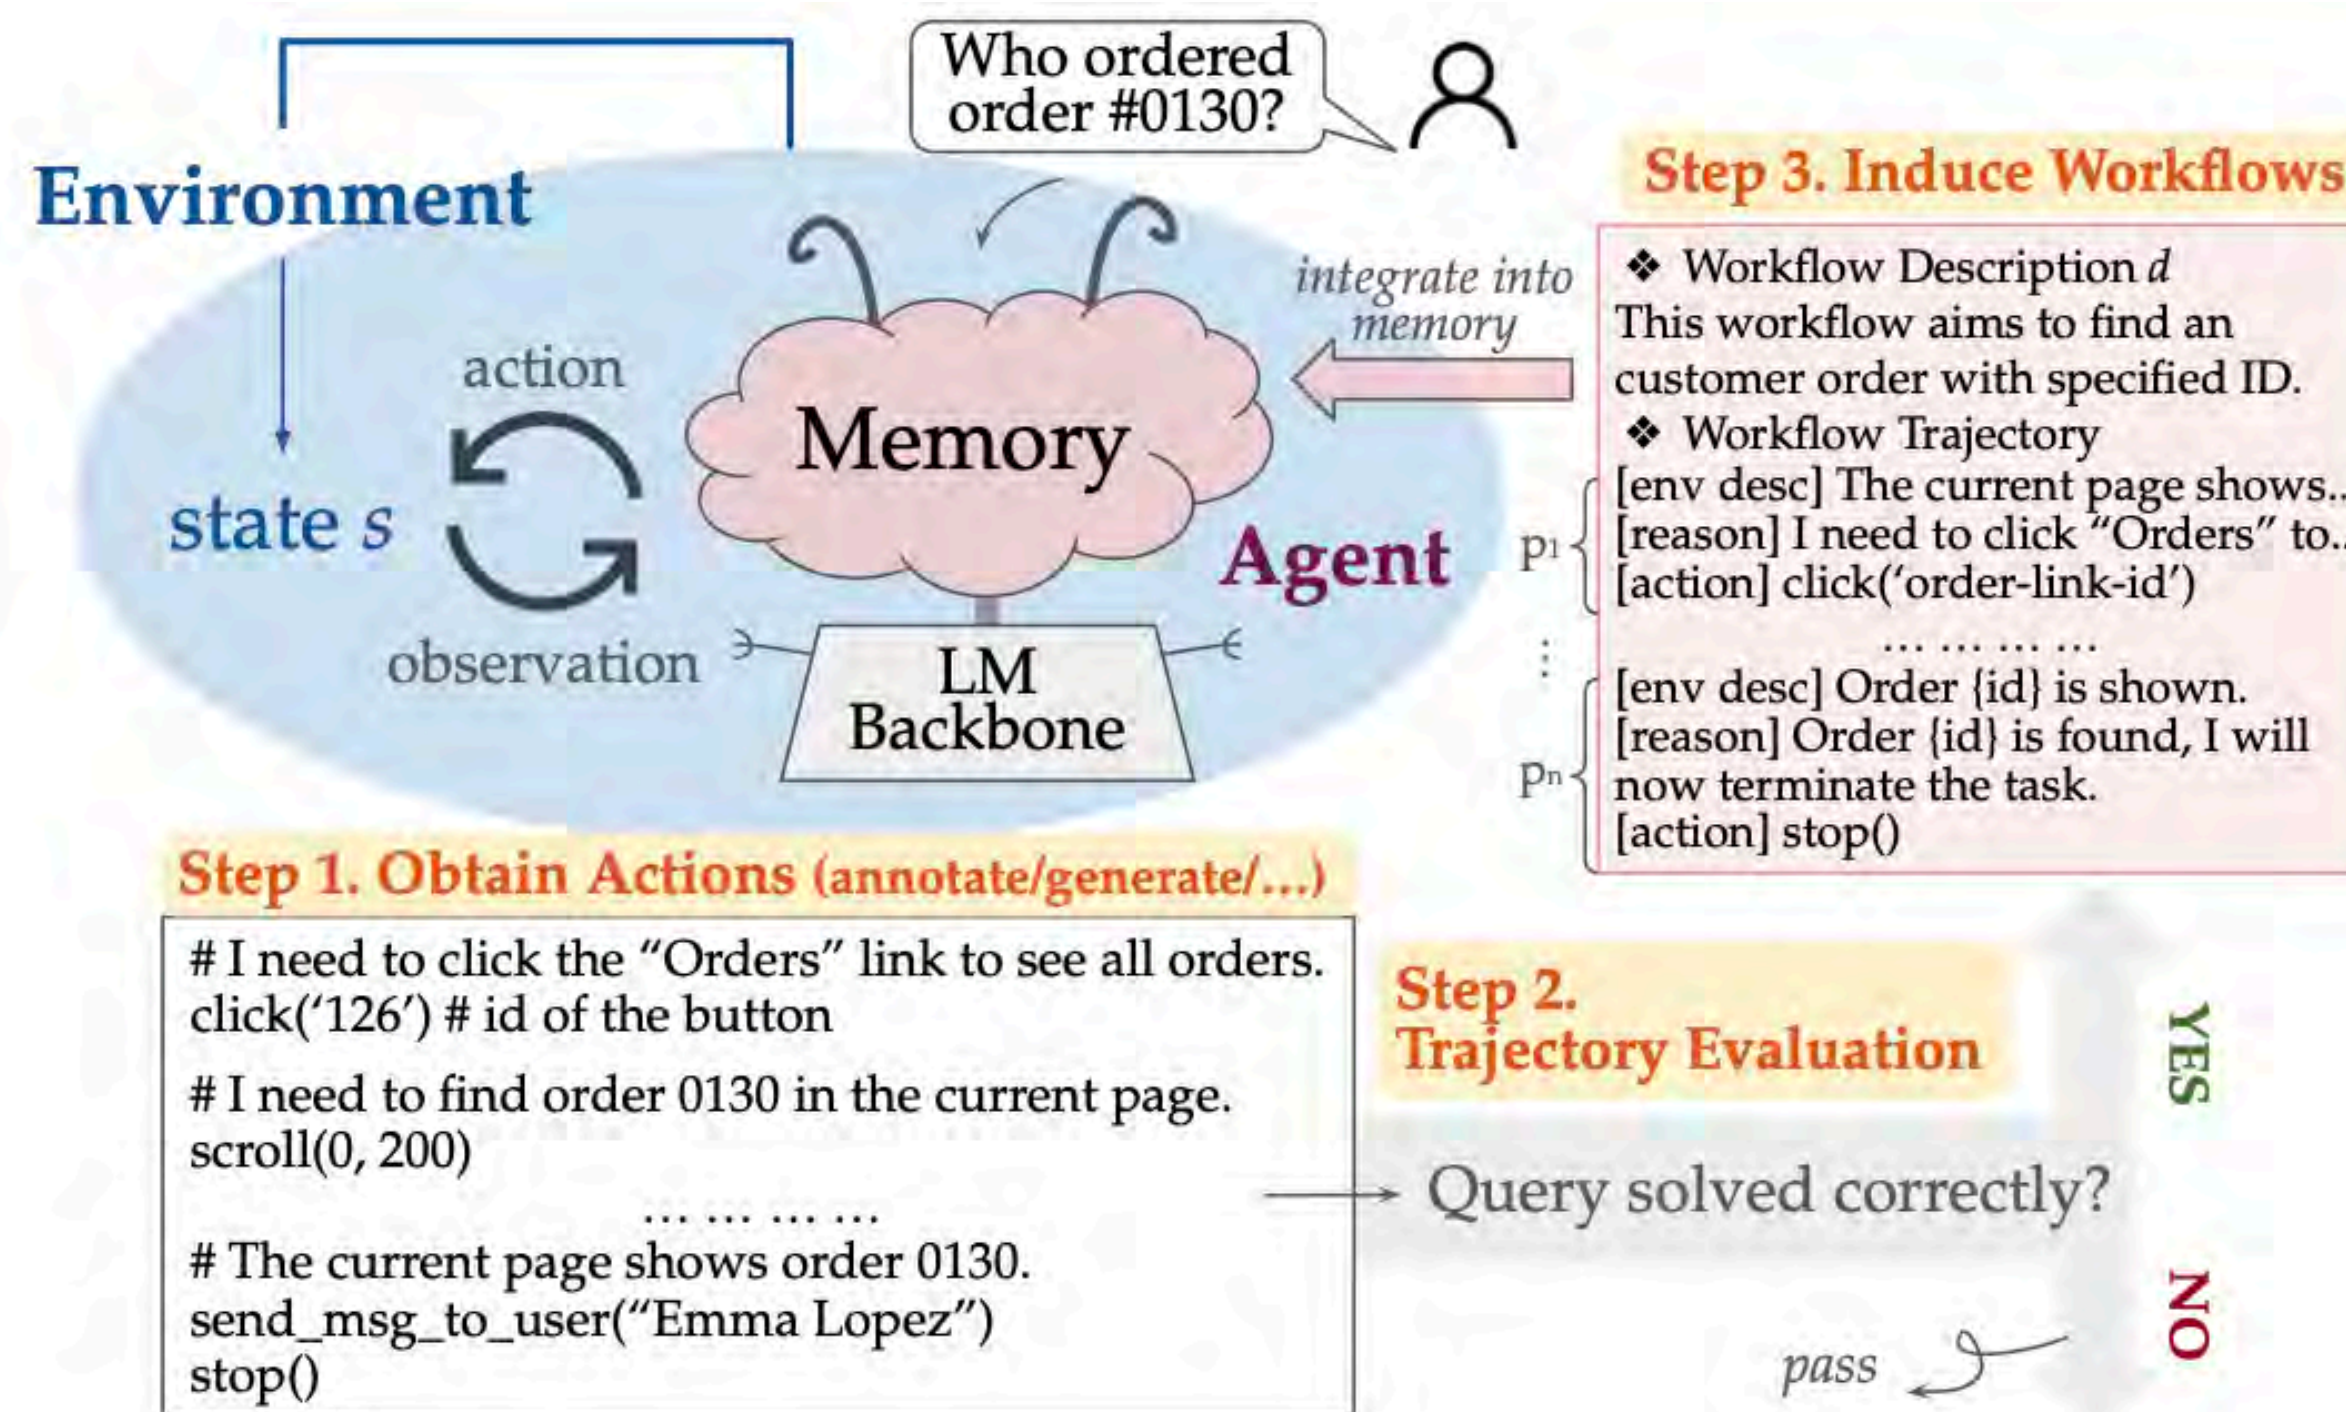
JARVIS-1: Open-World Multi-task Agents with Memory-Augmented Multimodal Language Models

self-improving agents

memory augmented MLM

(Wang, 2024)

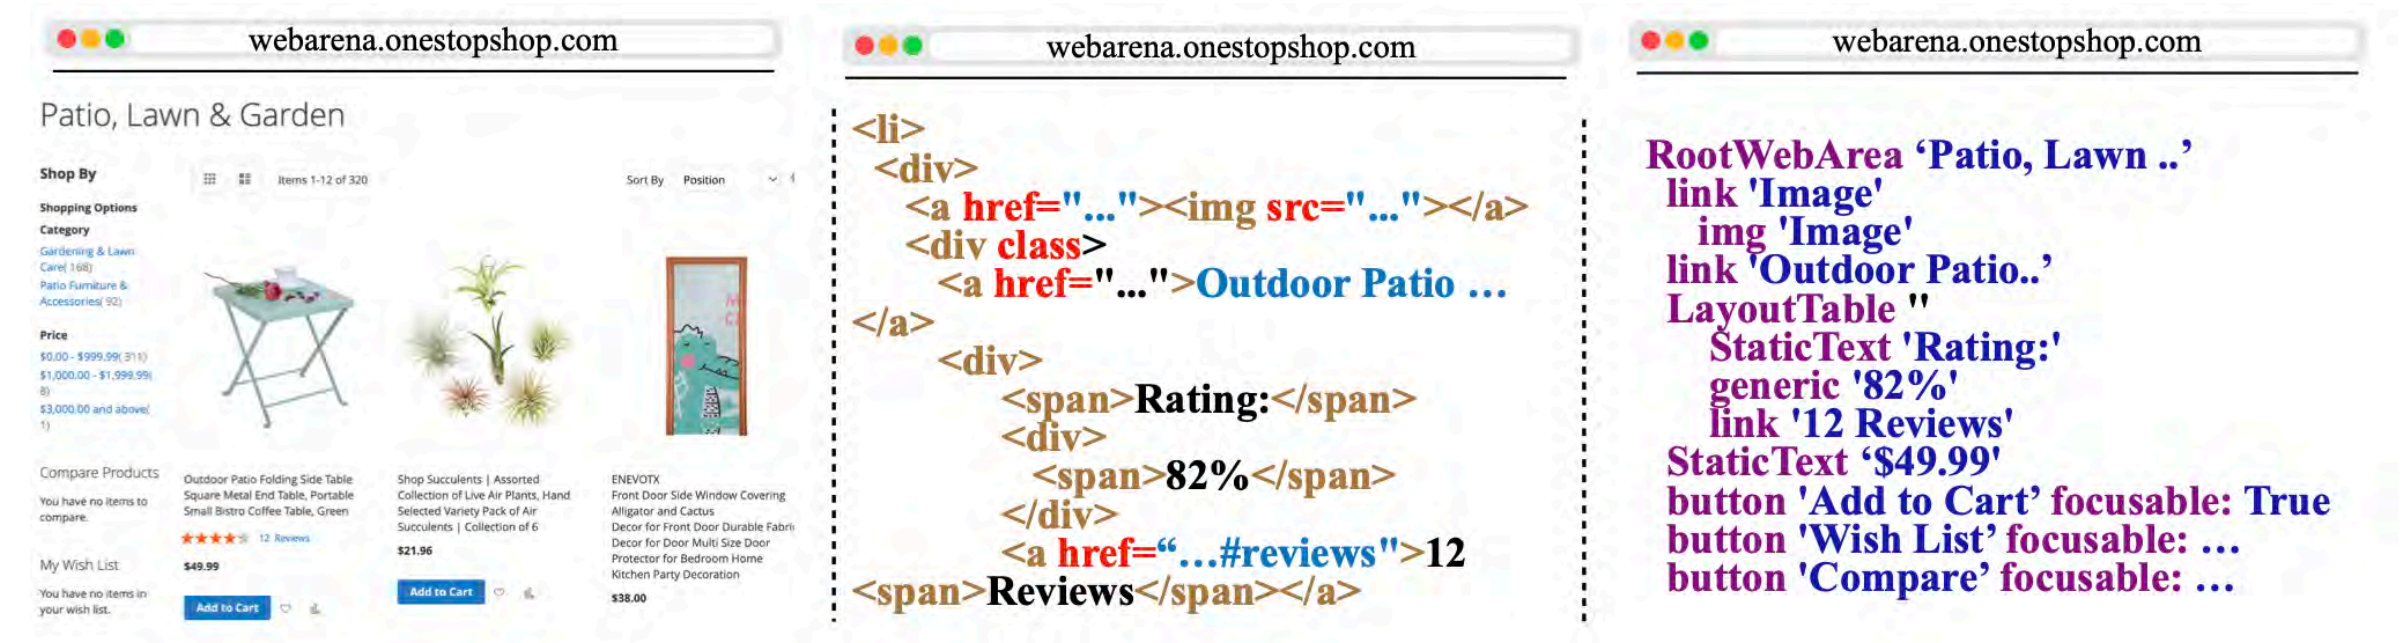
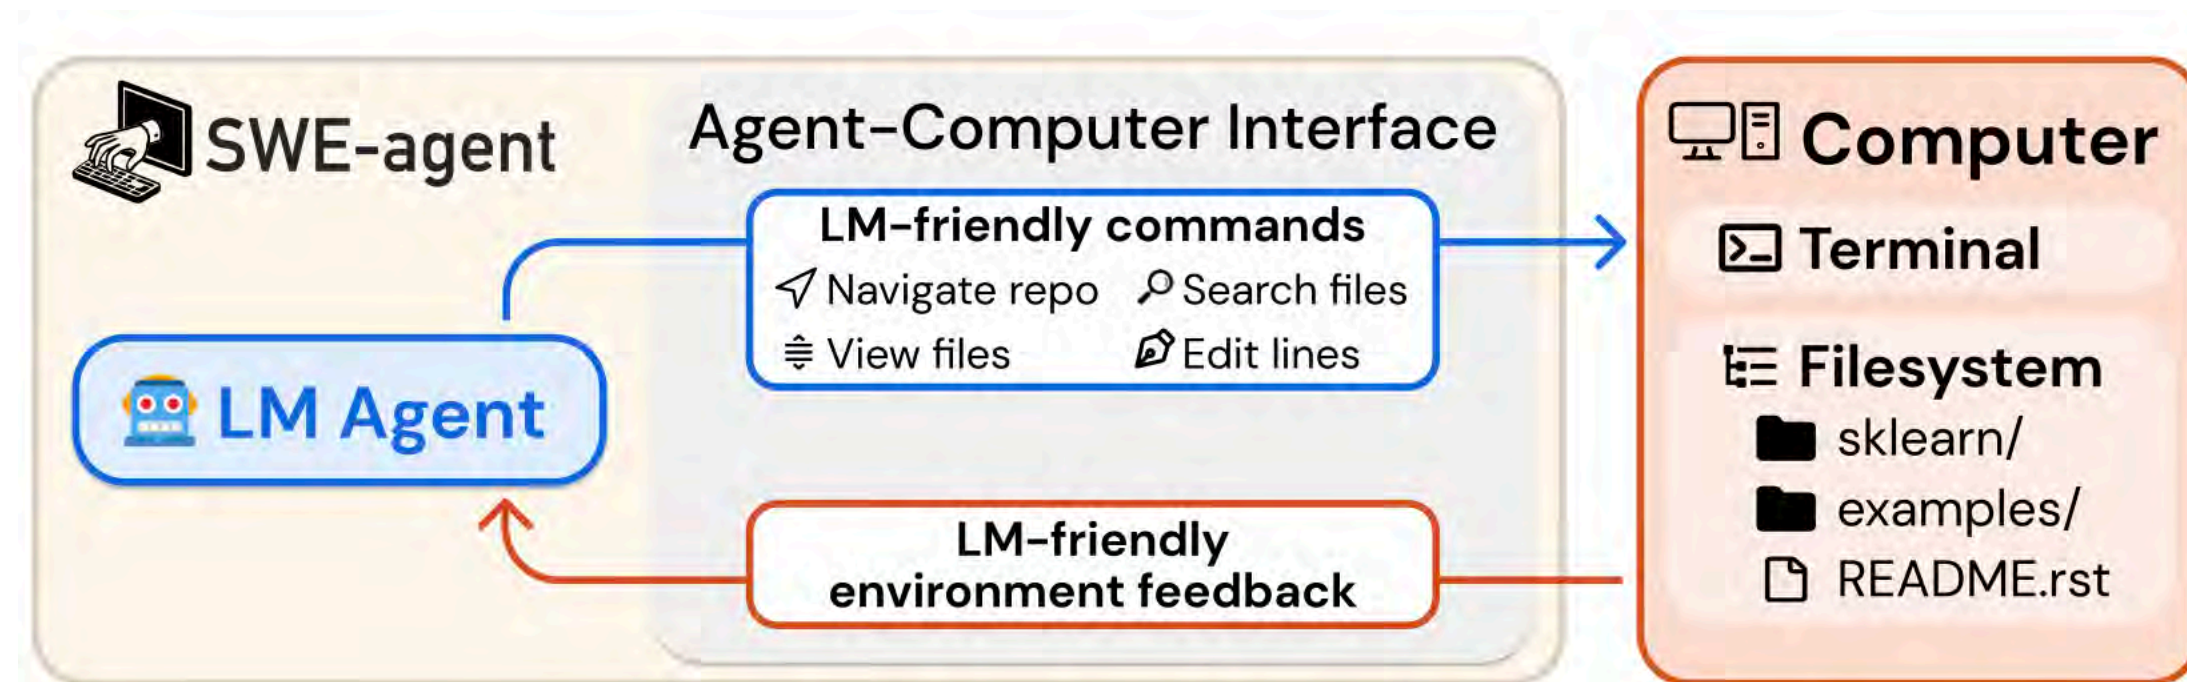
memory & ICL for learning



**agent workflow
memory
summarization +
positive experience
replay**

(Wang et al, 2024)

perception w/ LLM-favored input



  *SWE-agent: Agent-Computer Interfaces Enable Automated Software Engineering*

  *WebArena: A Realistic Web Environment for Building Autonomous Agents*

agent-computer interface

tools/feedback/guardrails

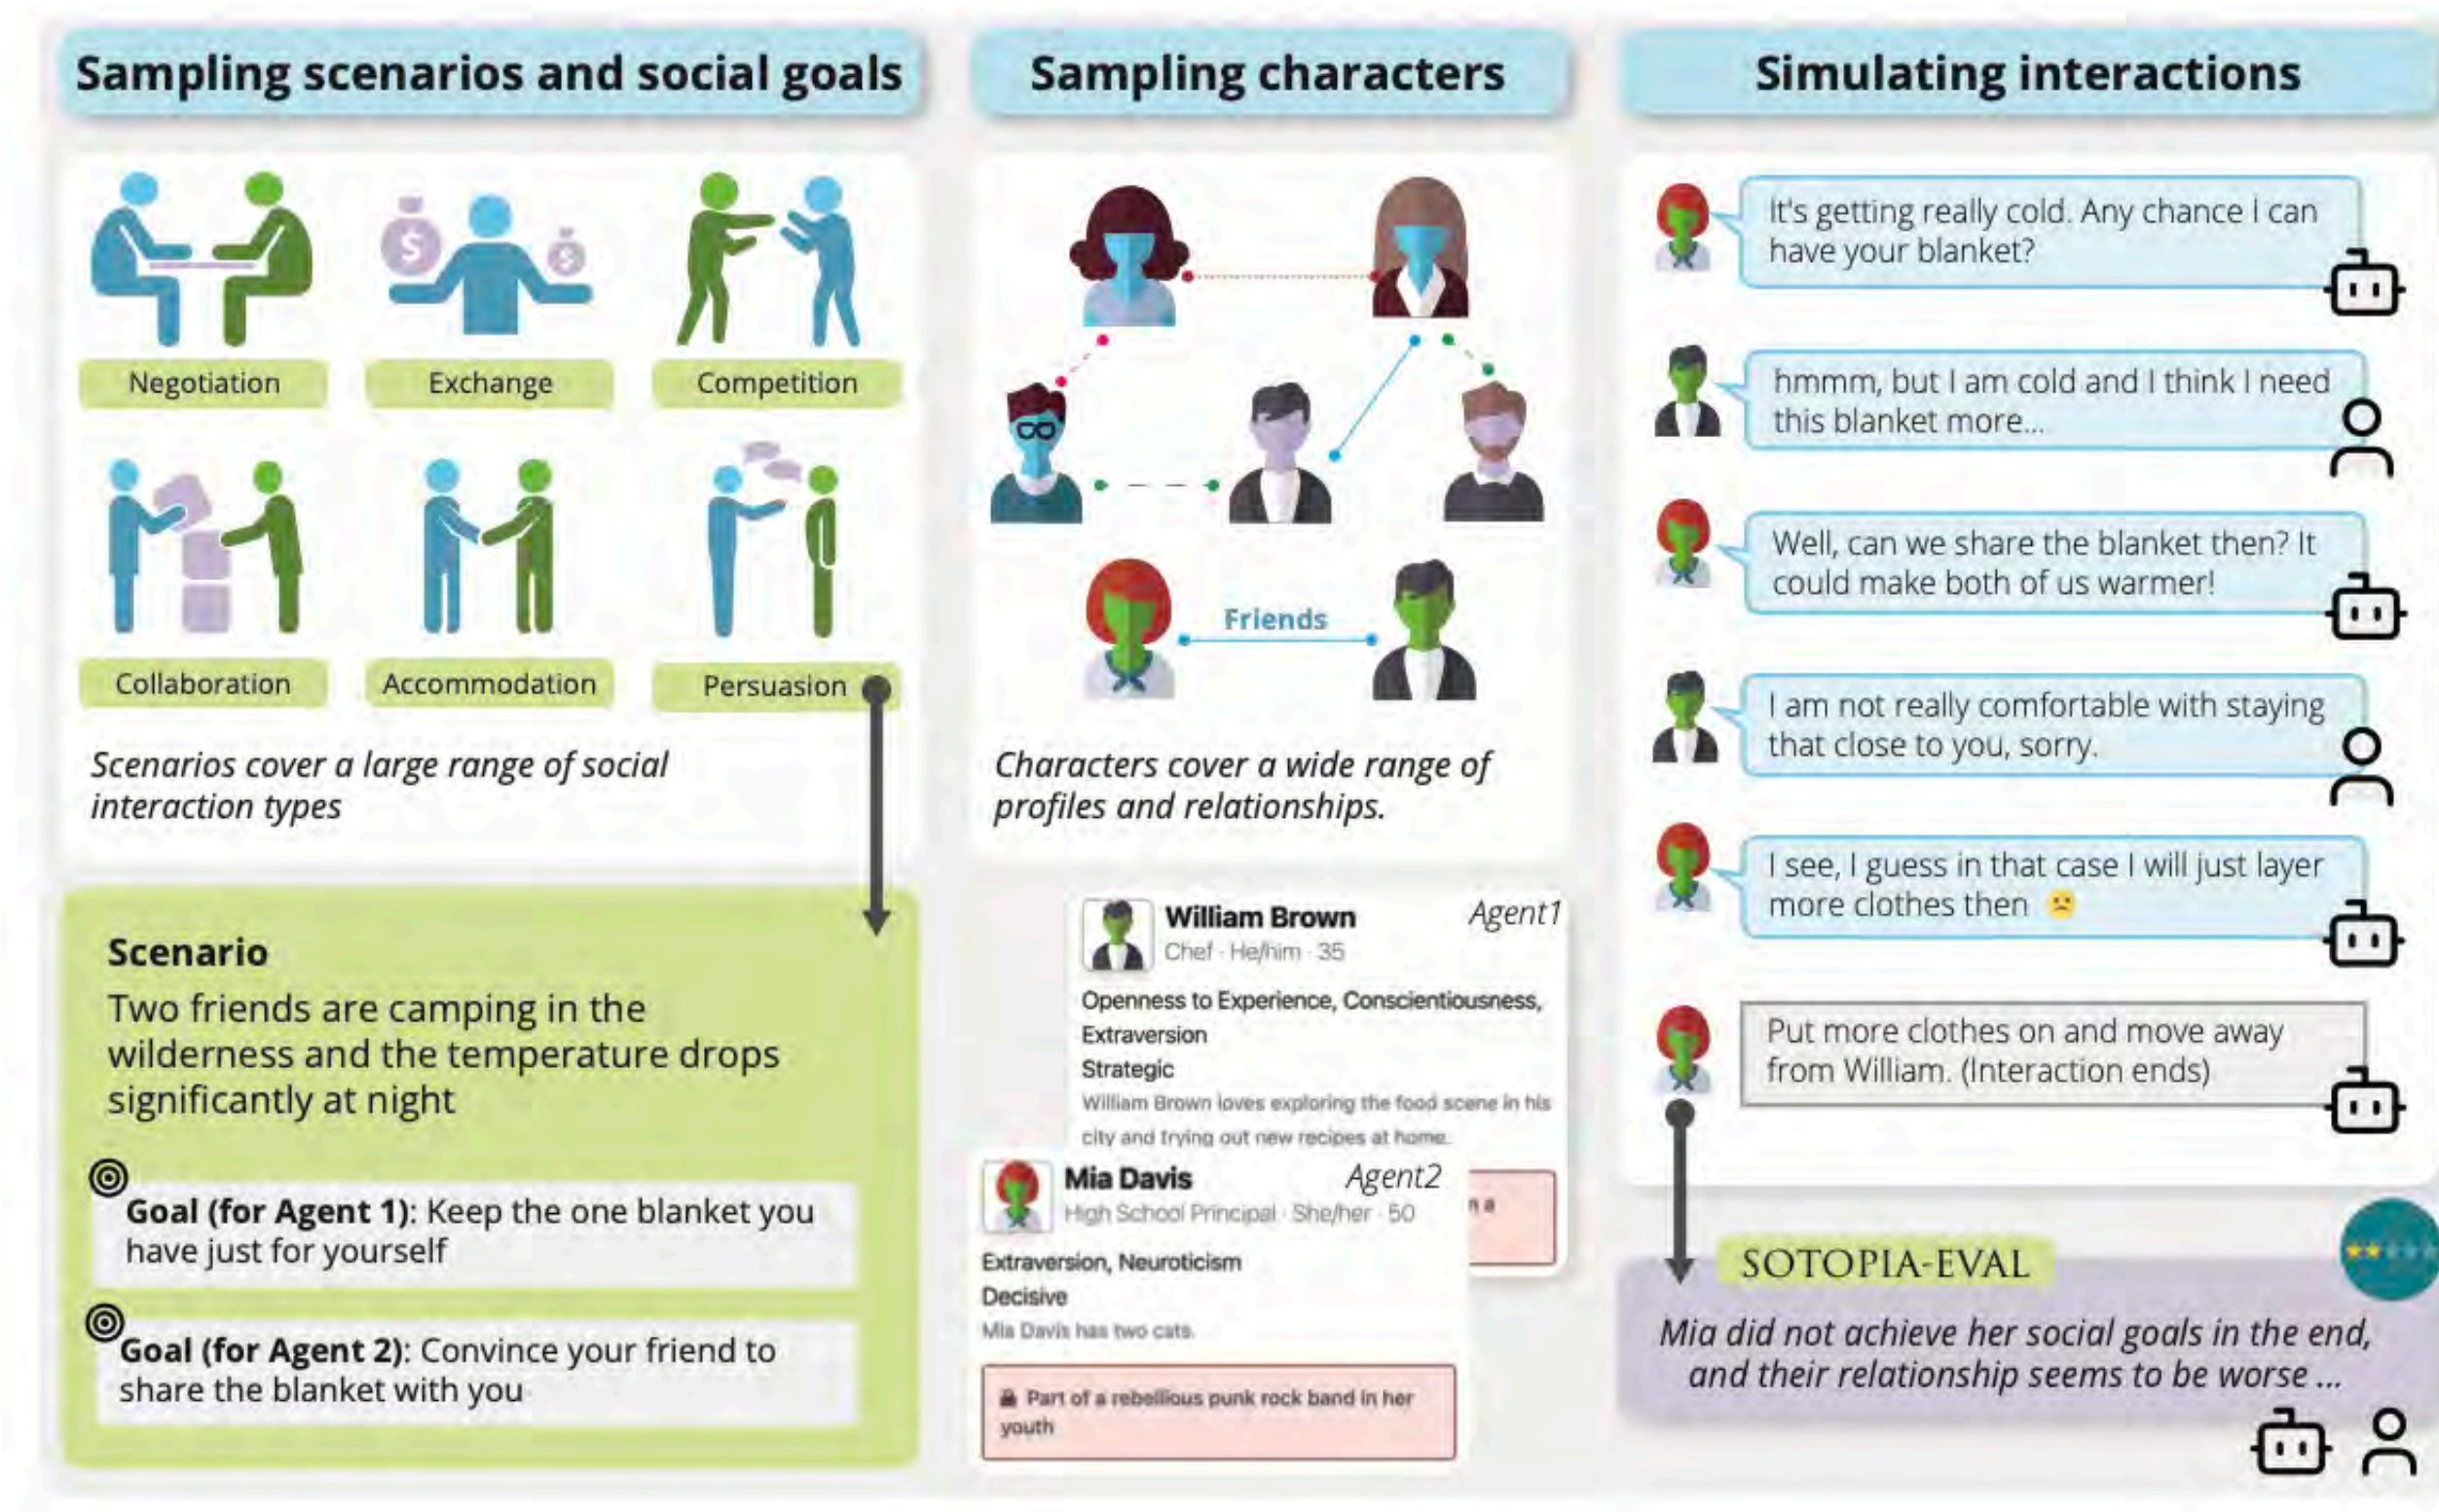
(Yang et al, 2024)

web navigation

accessibility tree

(Zhou et al, 2024)

agency



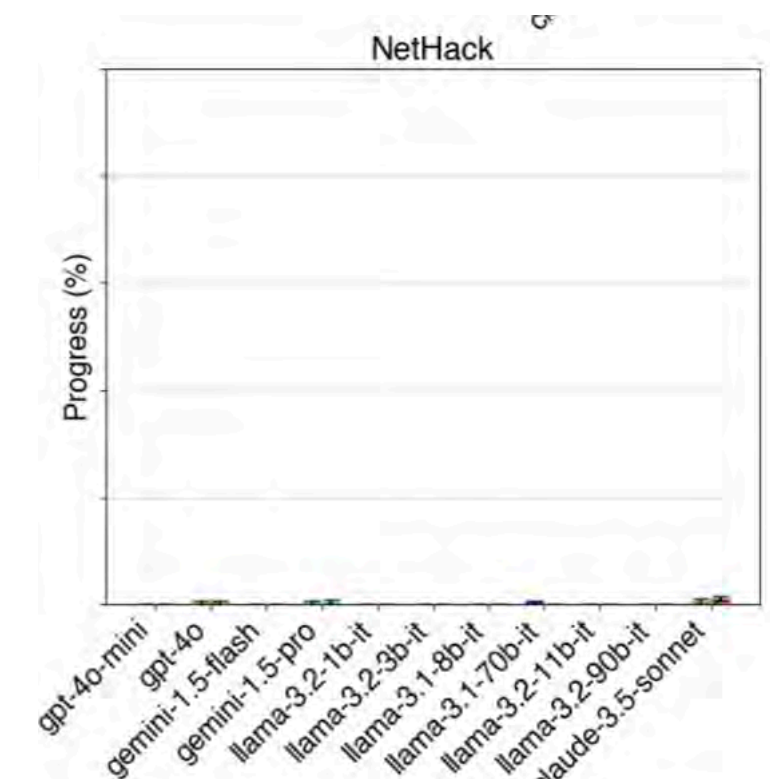
social intelligence
sources of action are
explicit (goals) +
implicit (norms)

🧠 ✨ individuality

(Zhou et al, 2024)

Sotopia: Interactive evaluation for social intelligence in language agents

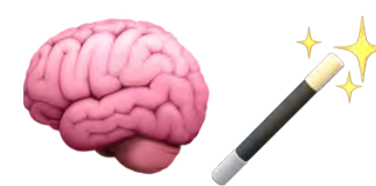
planning — looking ahead



10k-100k
turns

The NetHack Learning Environment

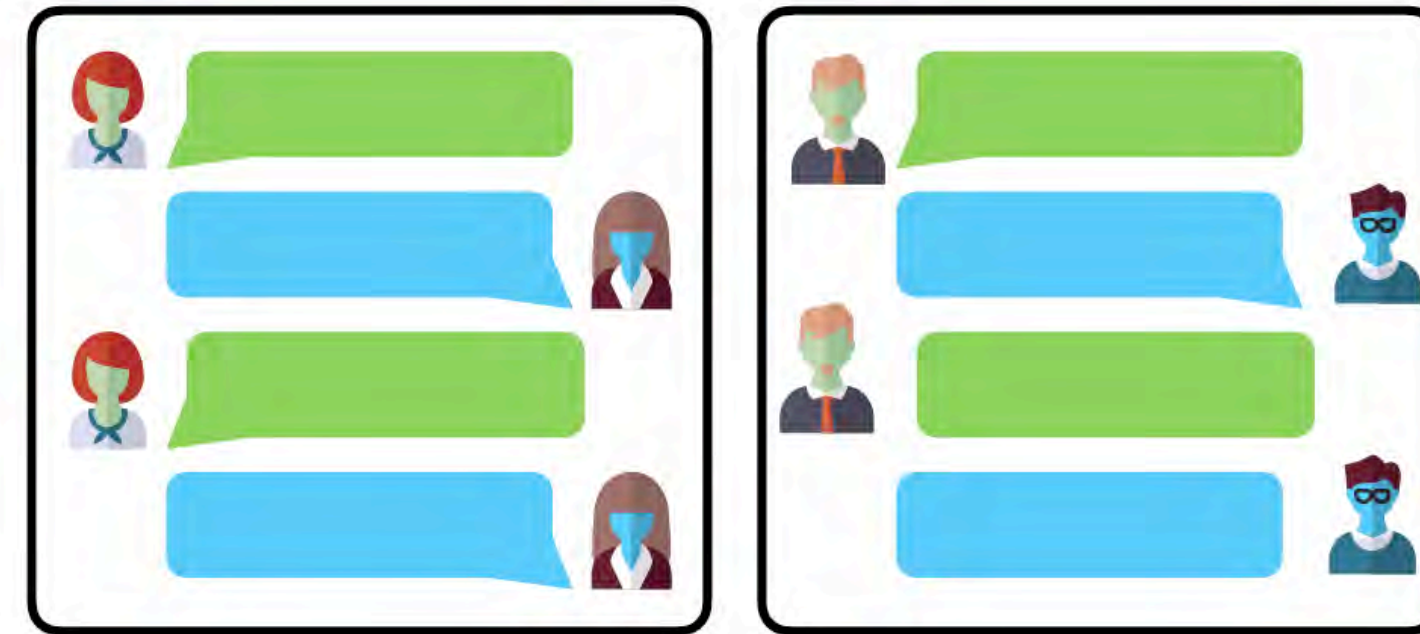
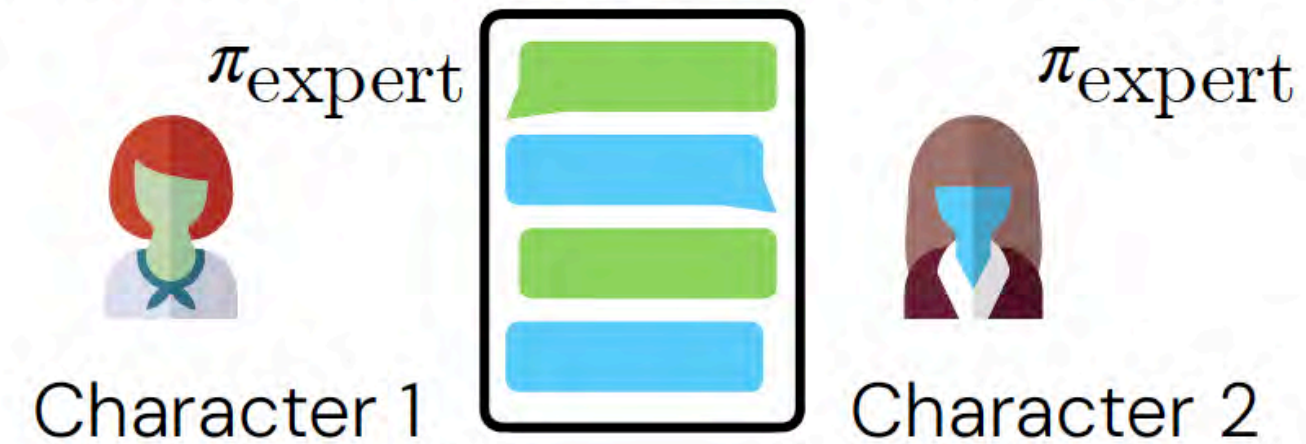
BALROG: Benchmarking Agentic LLM and VLM Reasoning On Games



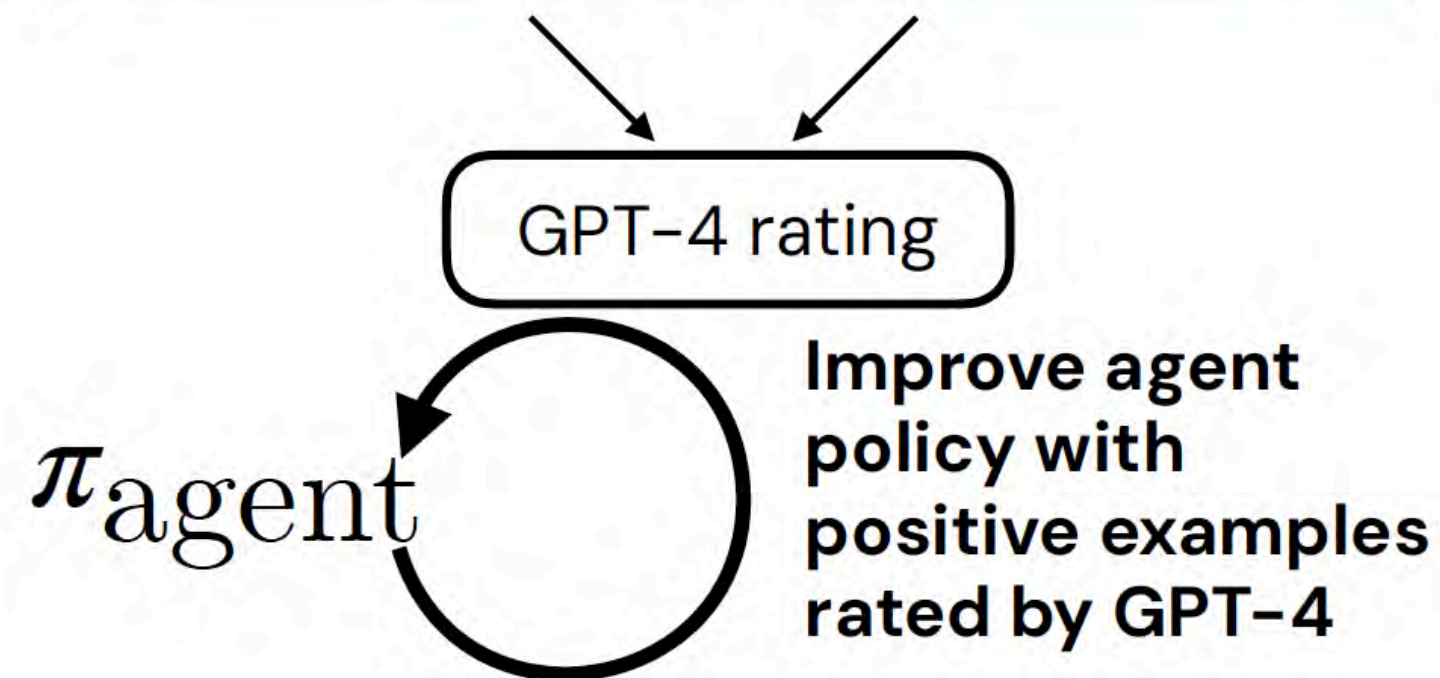
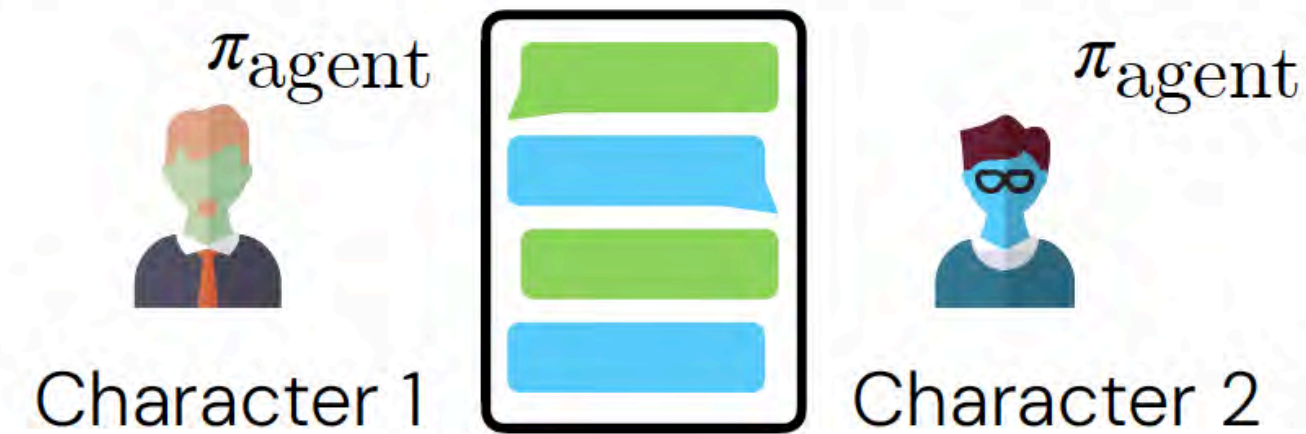
LLMs are good long context are they good at super long horizon?

learning

Collect data for *Behavior Cloning*



Collect data for *Self-Reinforcement*



**simple but robust
recipe**

**BC + SR (filtered BC)
(Wang et al, 2024)**

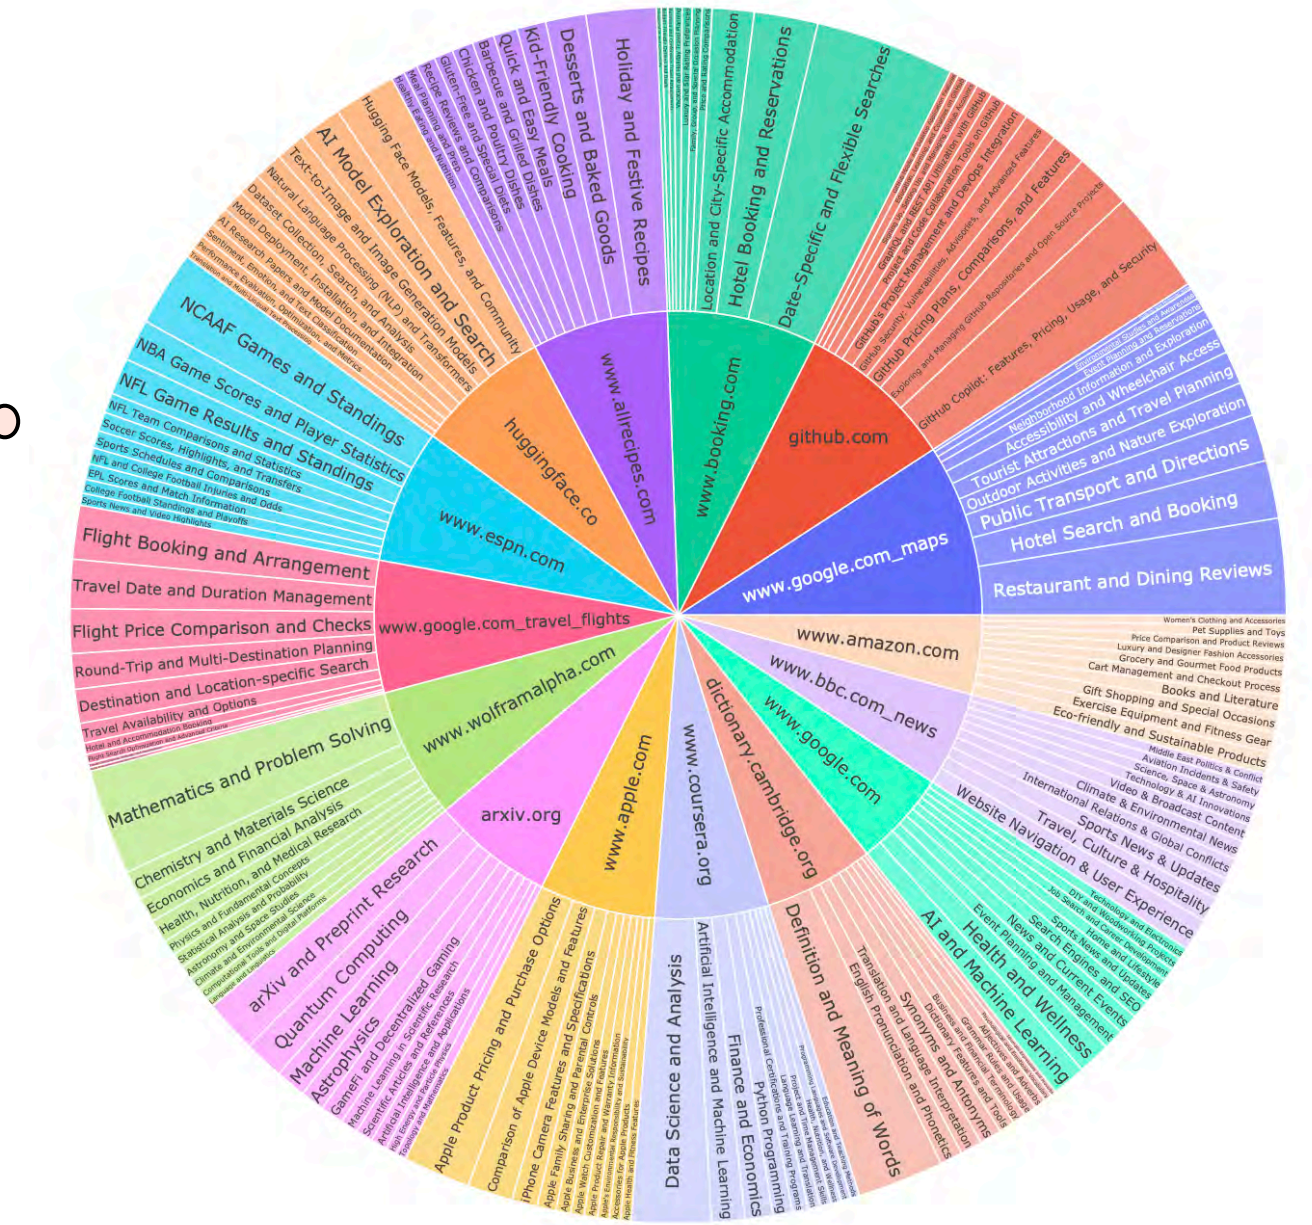
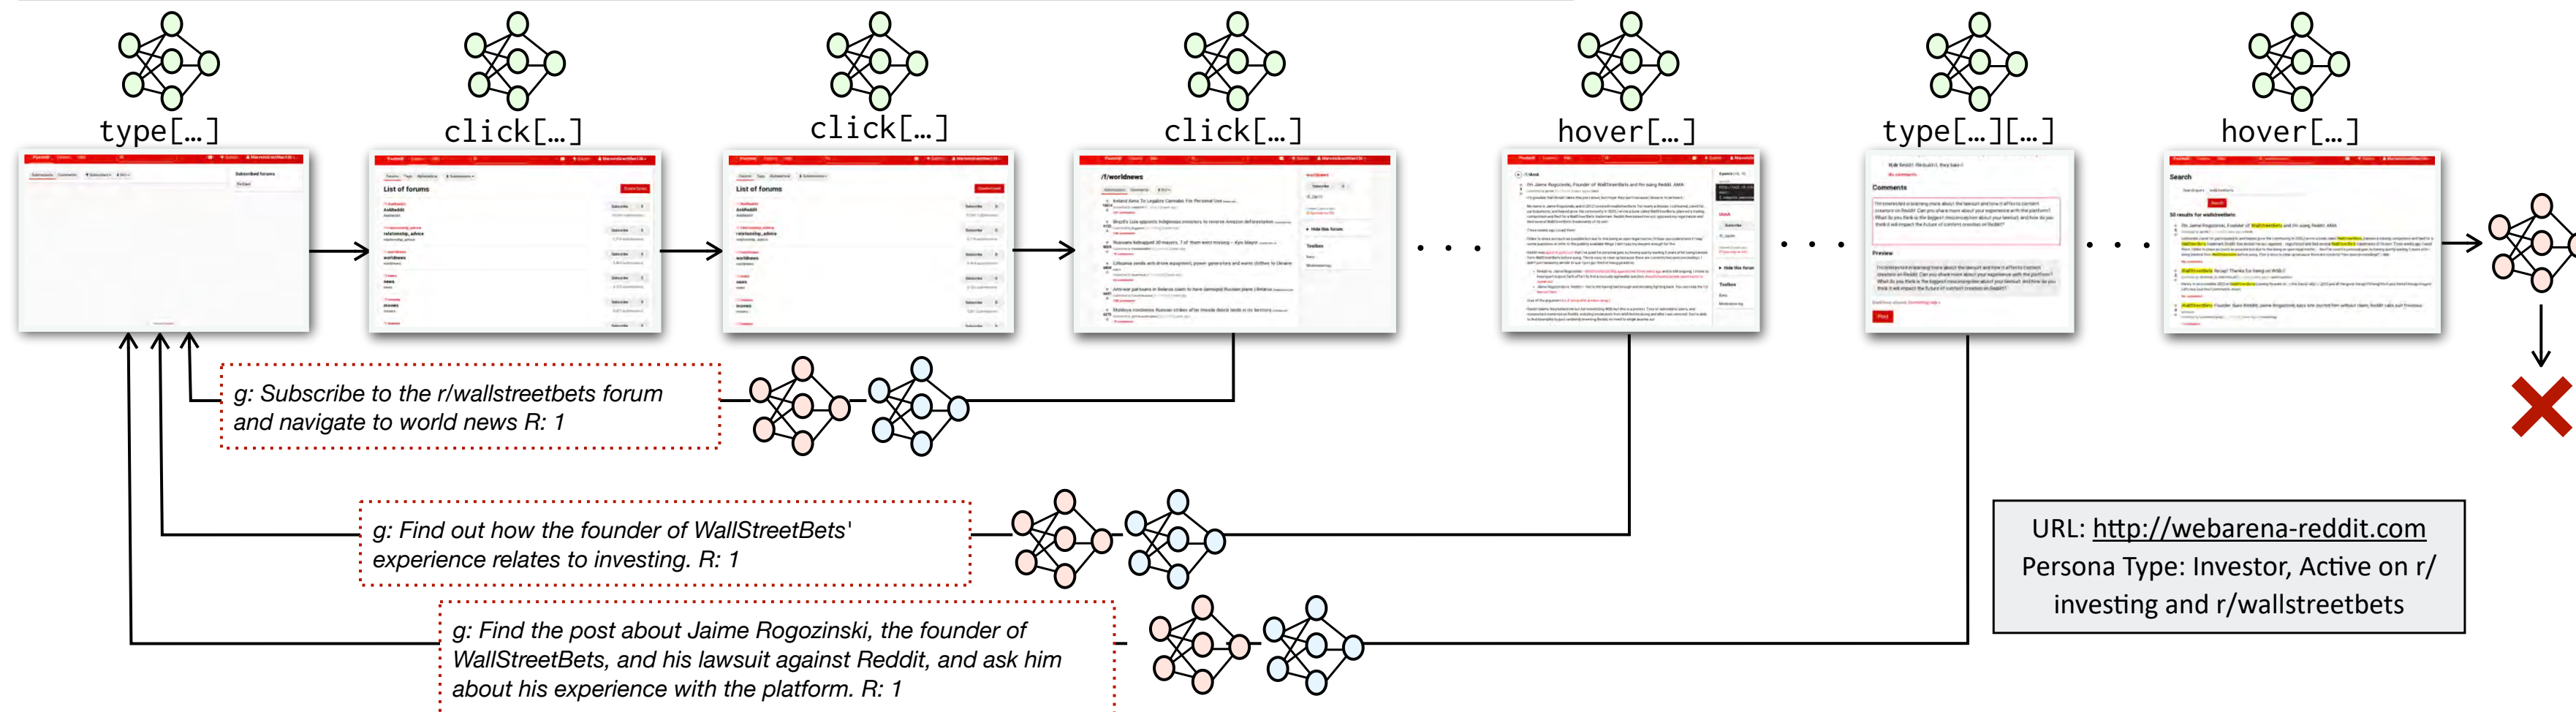
Sotopia- π : Interactive Learning of Socially Intelligent Language Agents.



SR only reinforces existing good behavior, won't work without a good prior.

learning & agency

1. Interact with a website with a structured exploration policy + trajectory labeler



NNetNav: Unsupervised Learning of Browser Agents Through Environment Interaction in the Wild

exploration + hindsight labeling


(Murty et al, 2025)

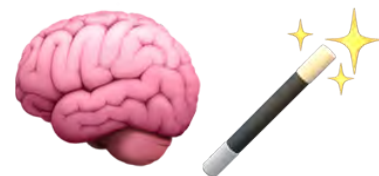
Persona driven diversity

short summary




agent thinking

When studying the aspects of agents, consider the strengths and weaknesses of LLMs, i.e. using .



LLM Thinking

When building on top of LLMs capabilities, consider the agentic aspect of them. Do they contribute to planning, learning, perception, or agency? .

bonus: another thinking tool



agent thinking



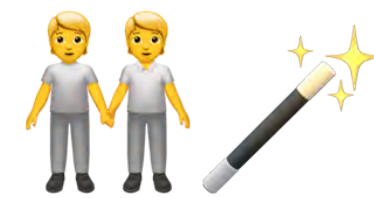
LLM Thinking





human thinking

human thinking

As a homework, review the previous examples, are there *safety, reliability, privacy* or other concerns that a human user might have?



human thinking

Are  and  helpful in mitigating these concerns?

What do people want from AI agents? Reliability? Safety? Privacy? Social Norm? Social Intelligence? Sense of control?

this lecture is heavily influenced by

- Graham Neubig (CMU): <https://youtu.be/a3SjRsqV9ZA>
- Hongyi Li (李宏毅, NTU): <https://youtu.be/M2Yg1kwPpts> (in Mandarin)
- Prithviraj Ammanabrolu (UCSD): <https://pearls-lab.github.io/ai-agents-course/index.html>

Please check them out.

thanks!

questions?

you can also reach me at <https://zhuhao.me>