Rethinking the role of LLMs: from central reasoning system to source of knowledge for symbolic AI systems

NeuroSym4MLLM

Guillaume Sarthou
Bastien Dussard

27 Janvier 2026

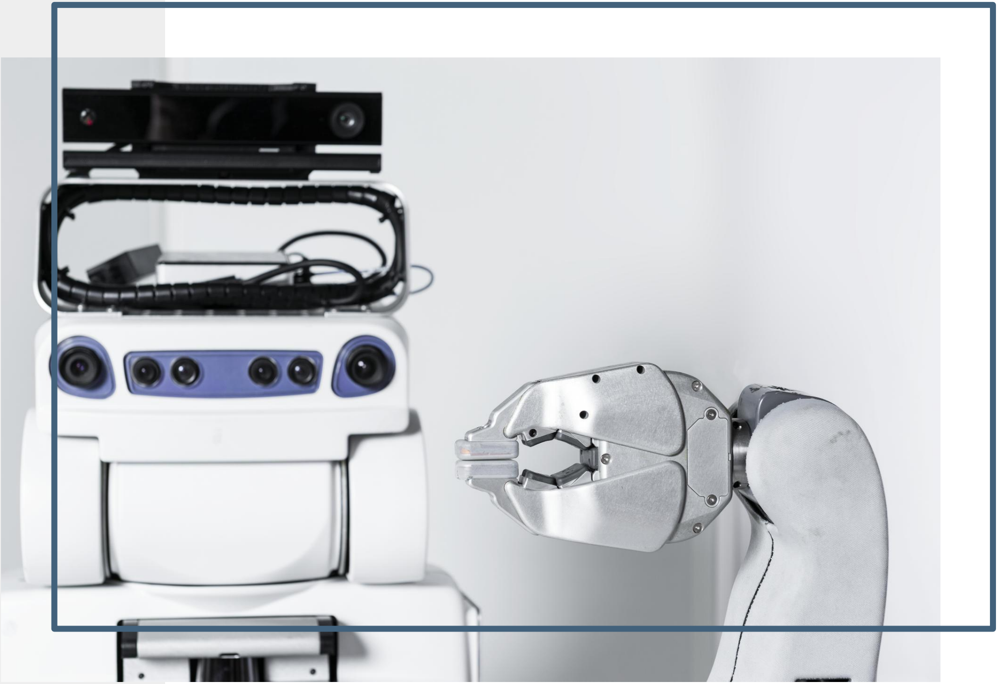

Université de Toulouse
LAAS CNRS

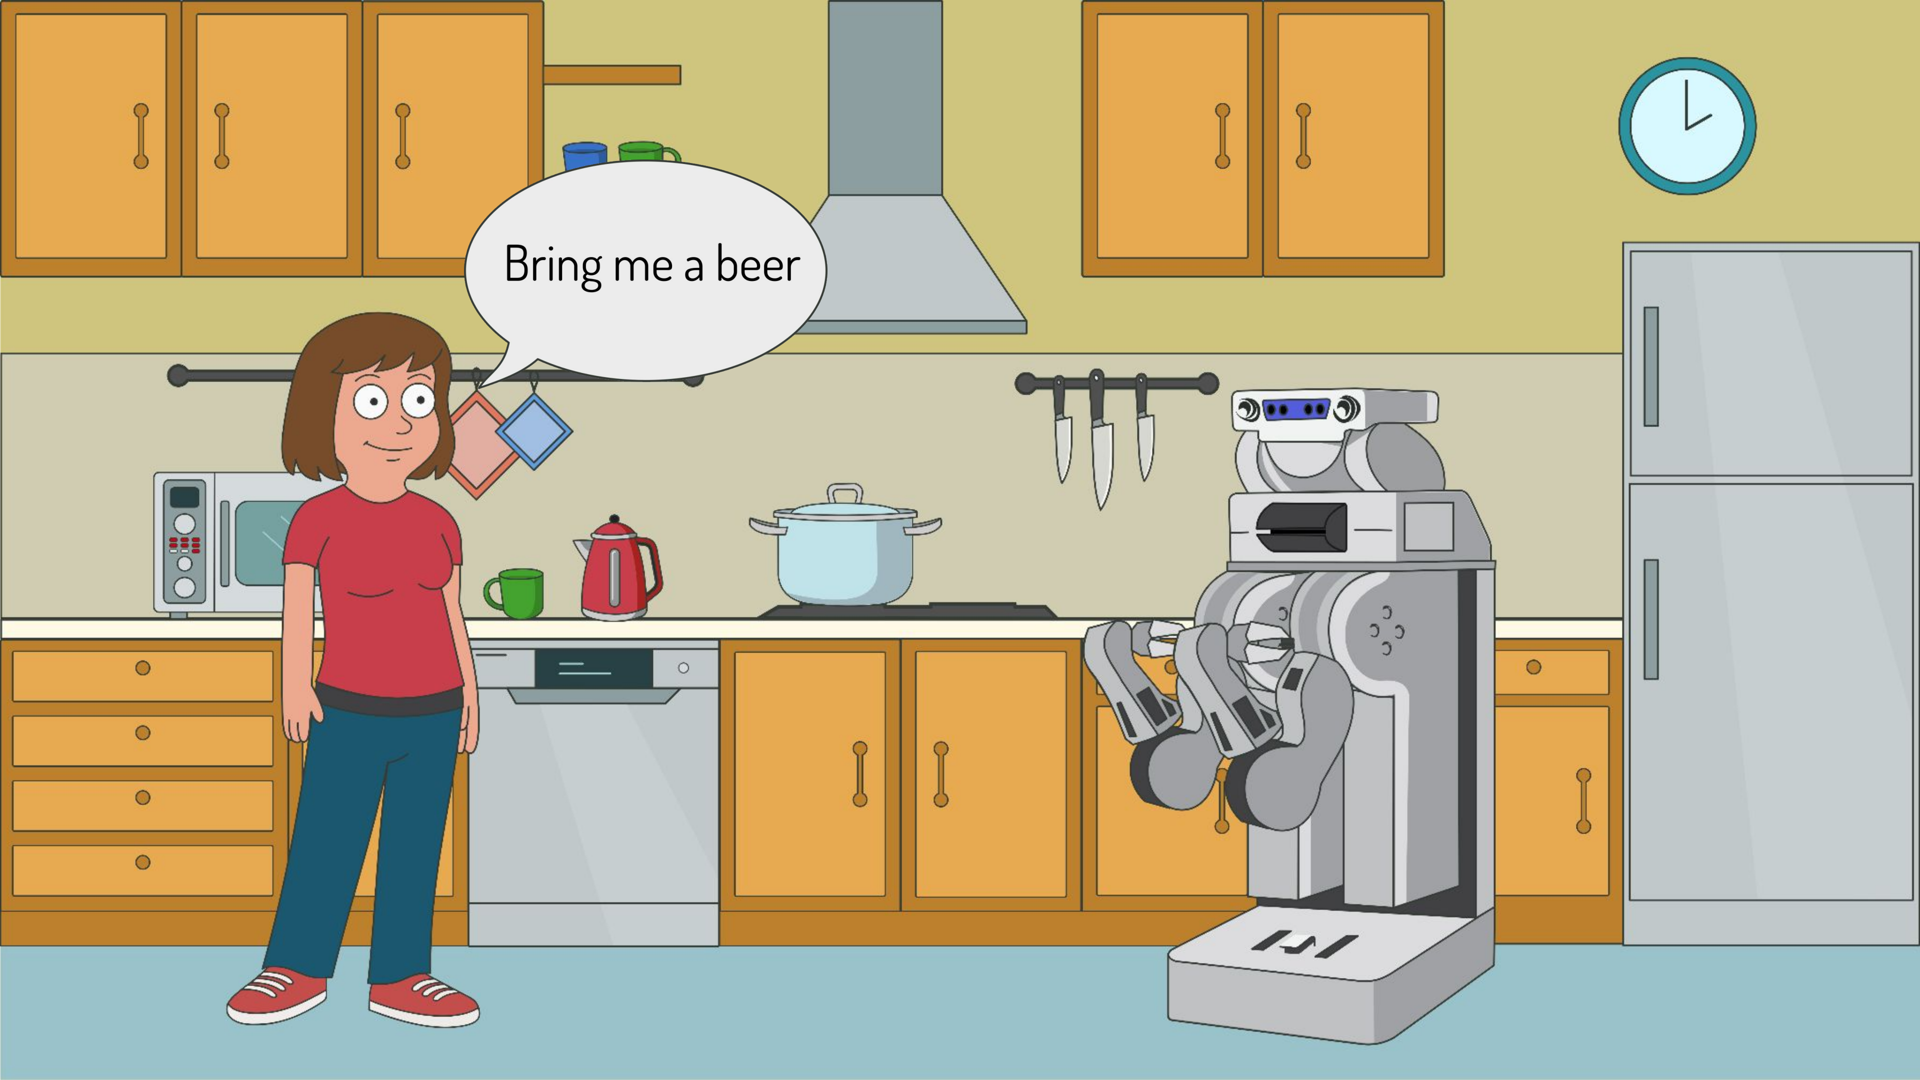
Bring me a beer

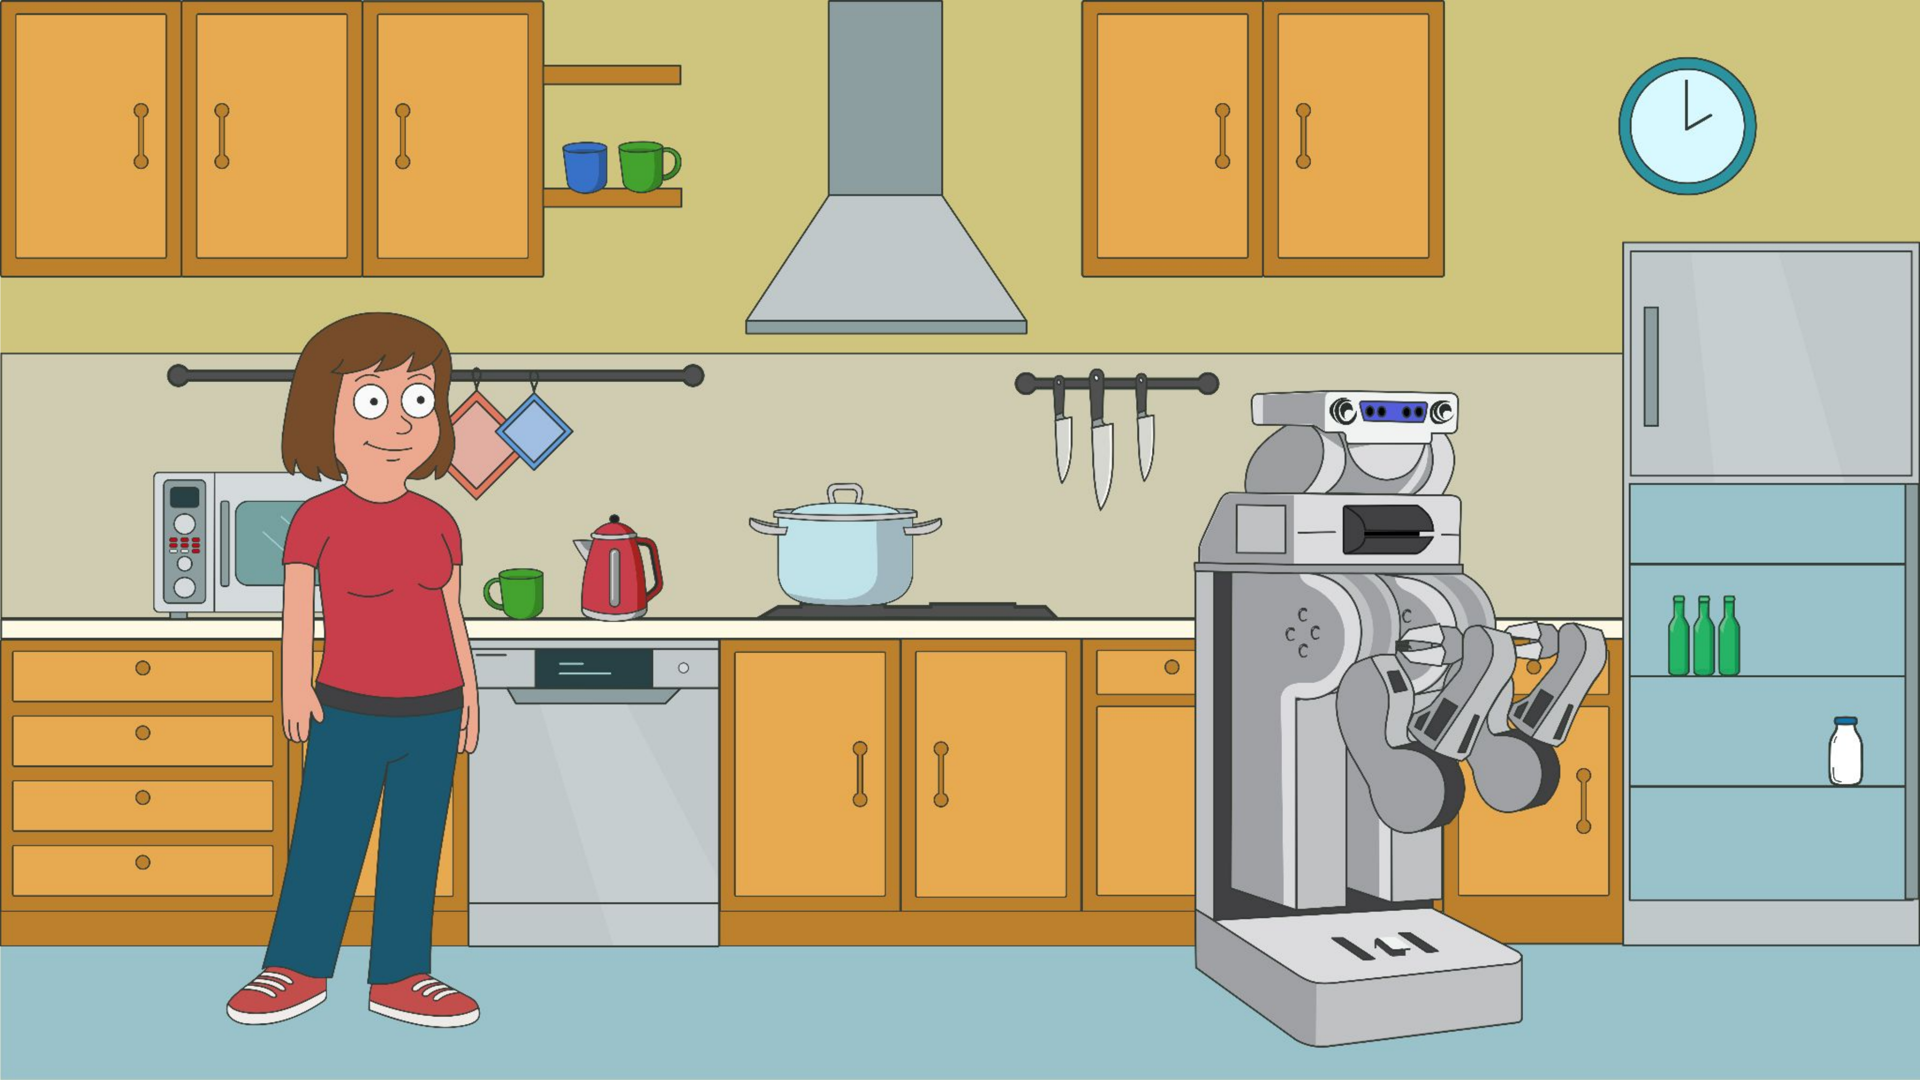

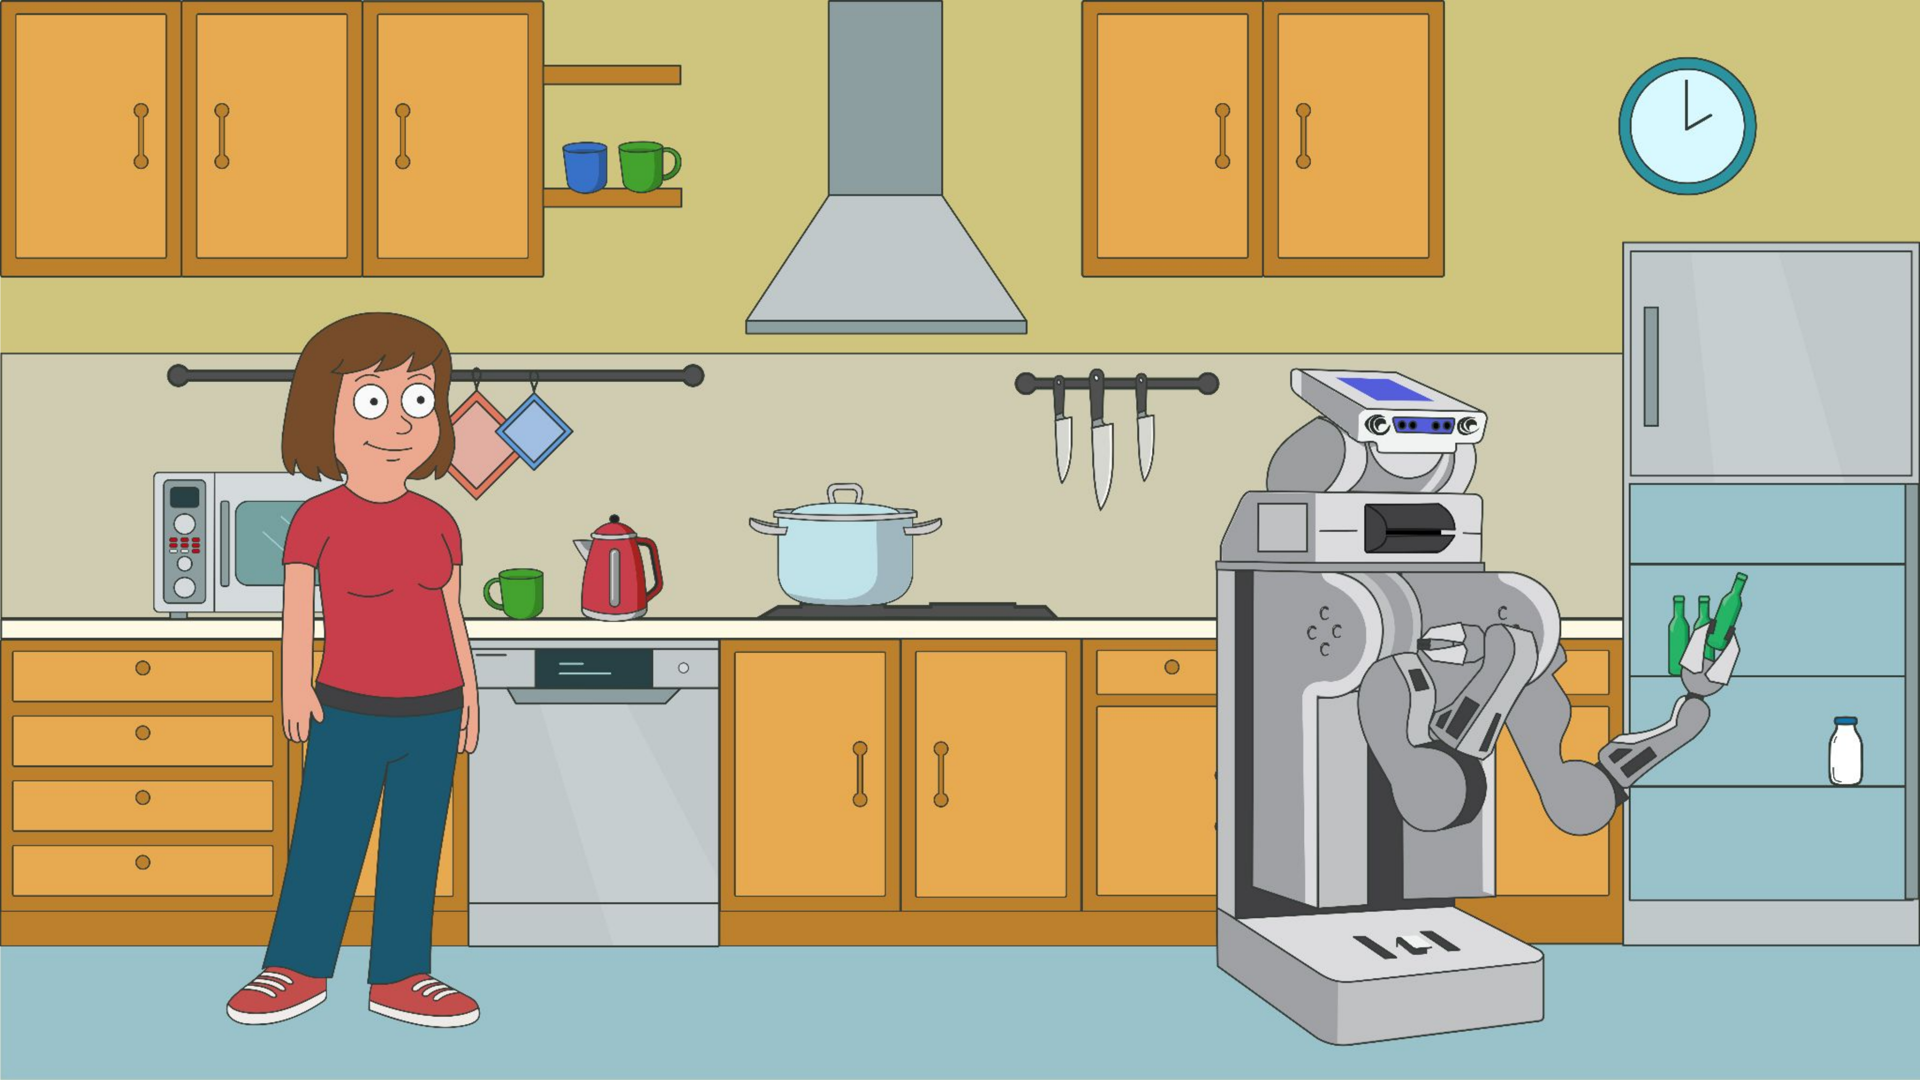

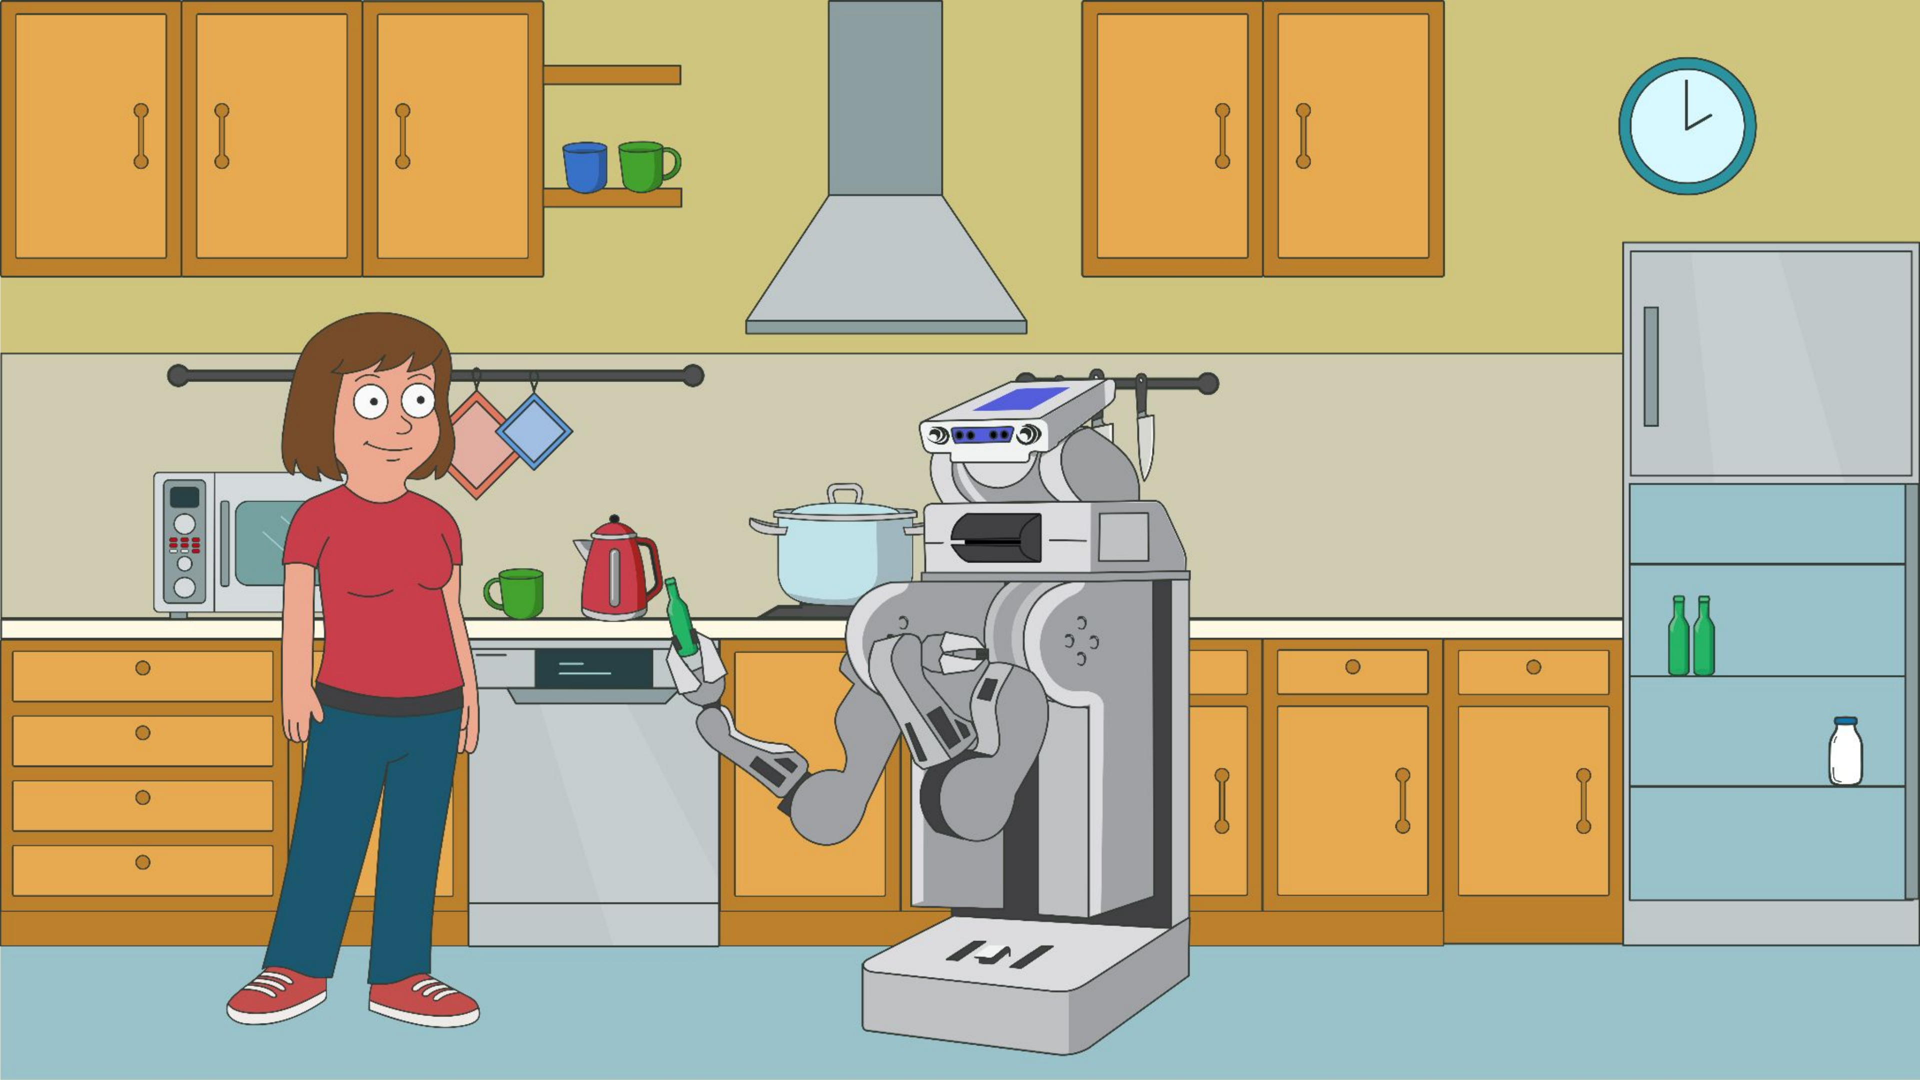

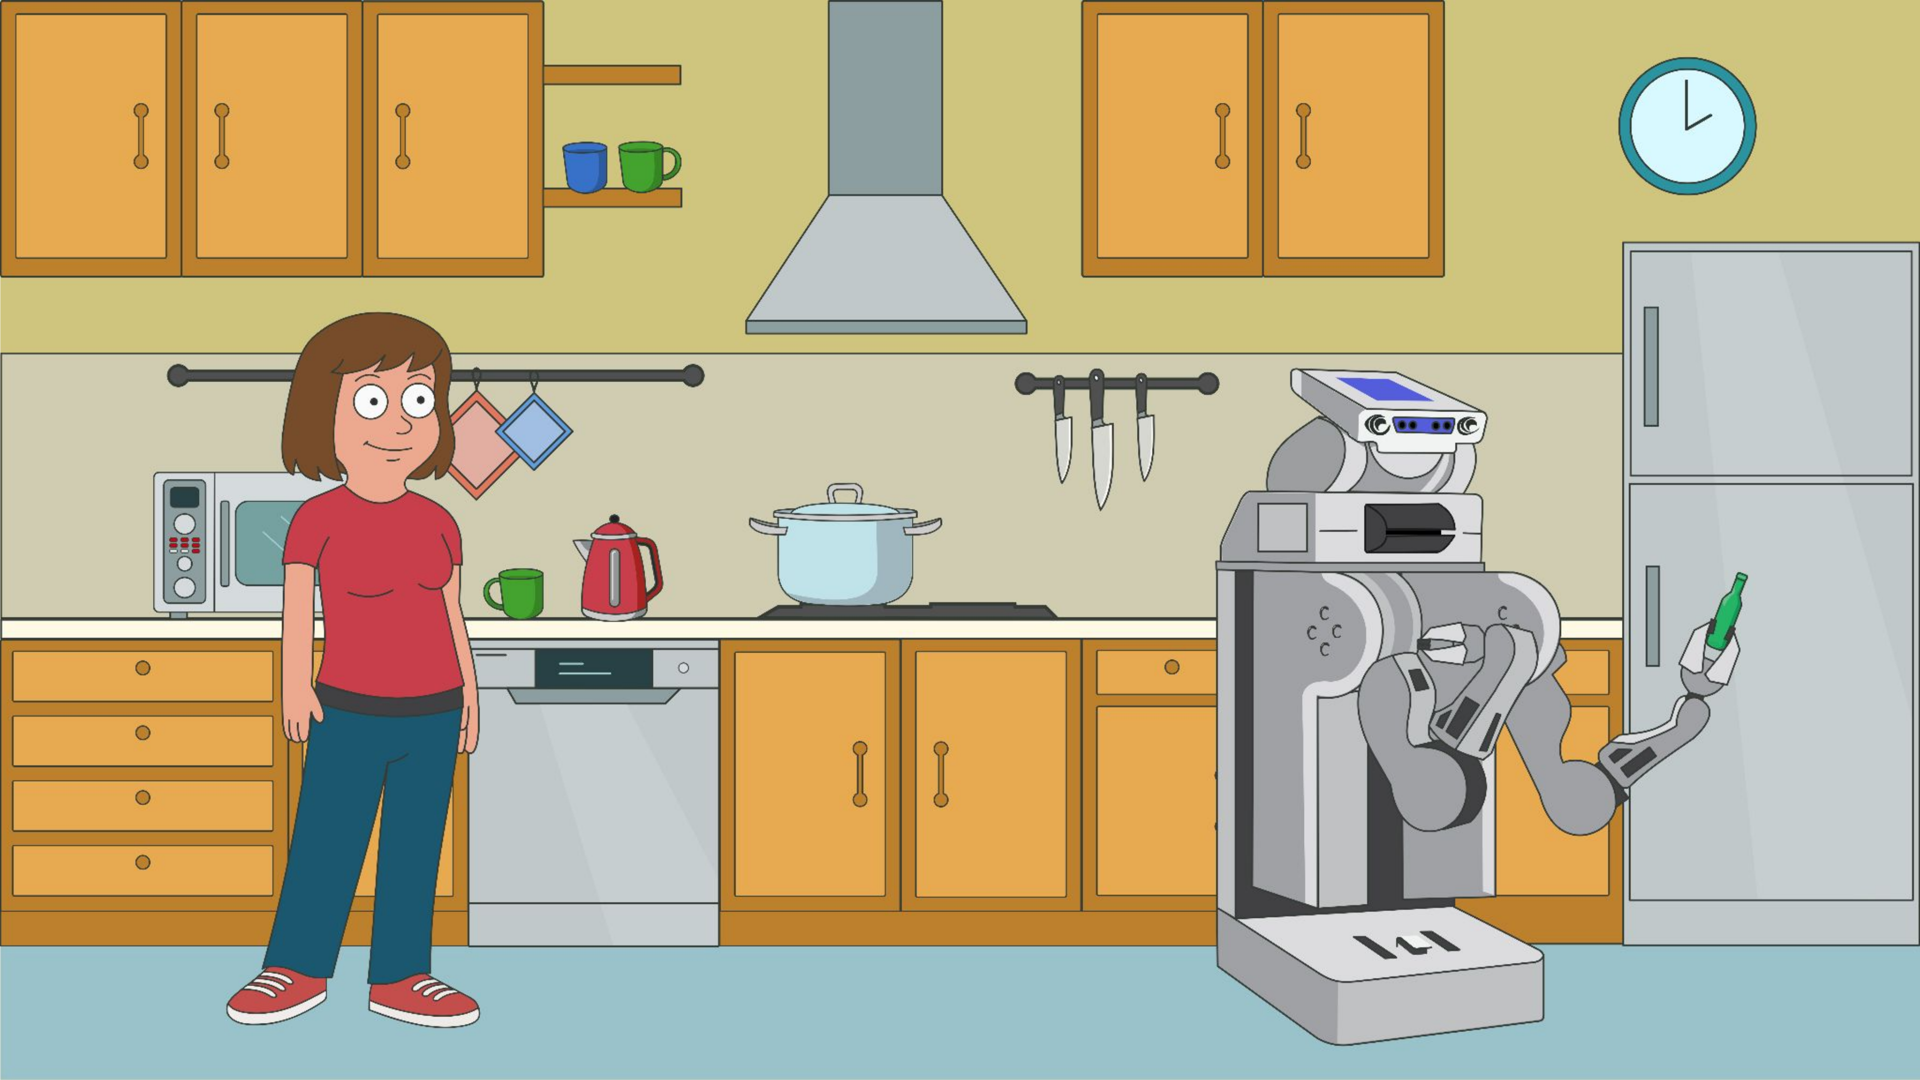

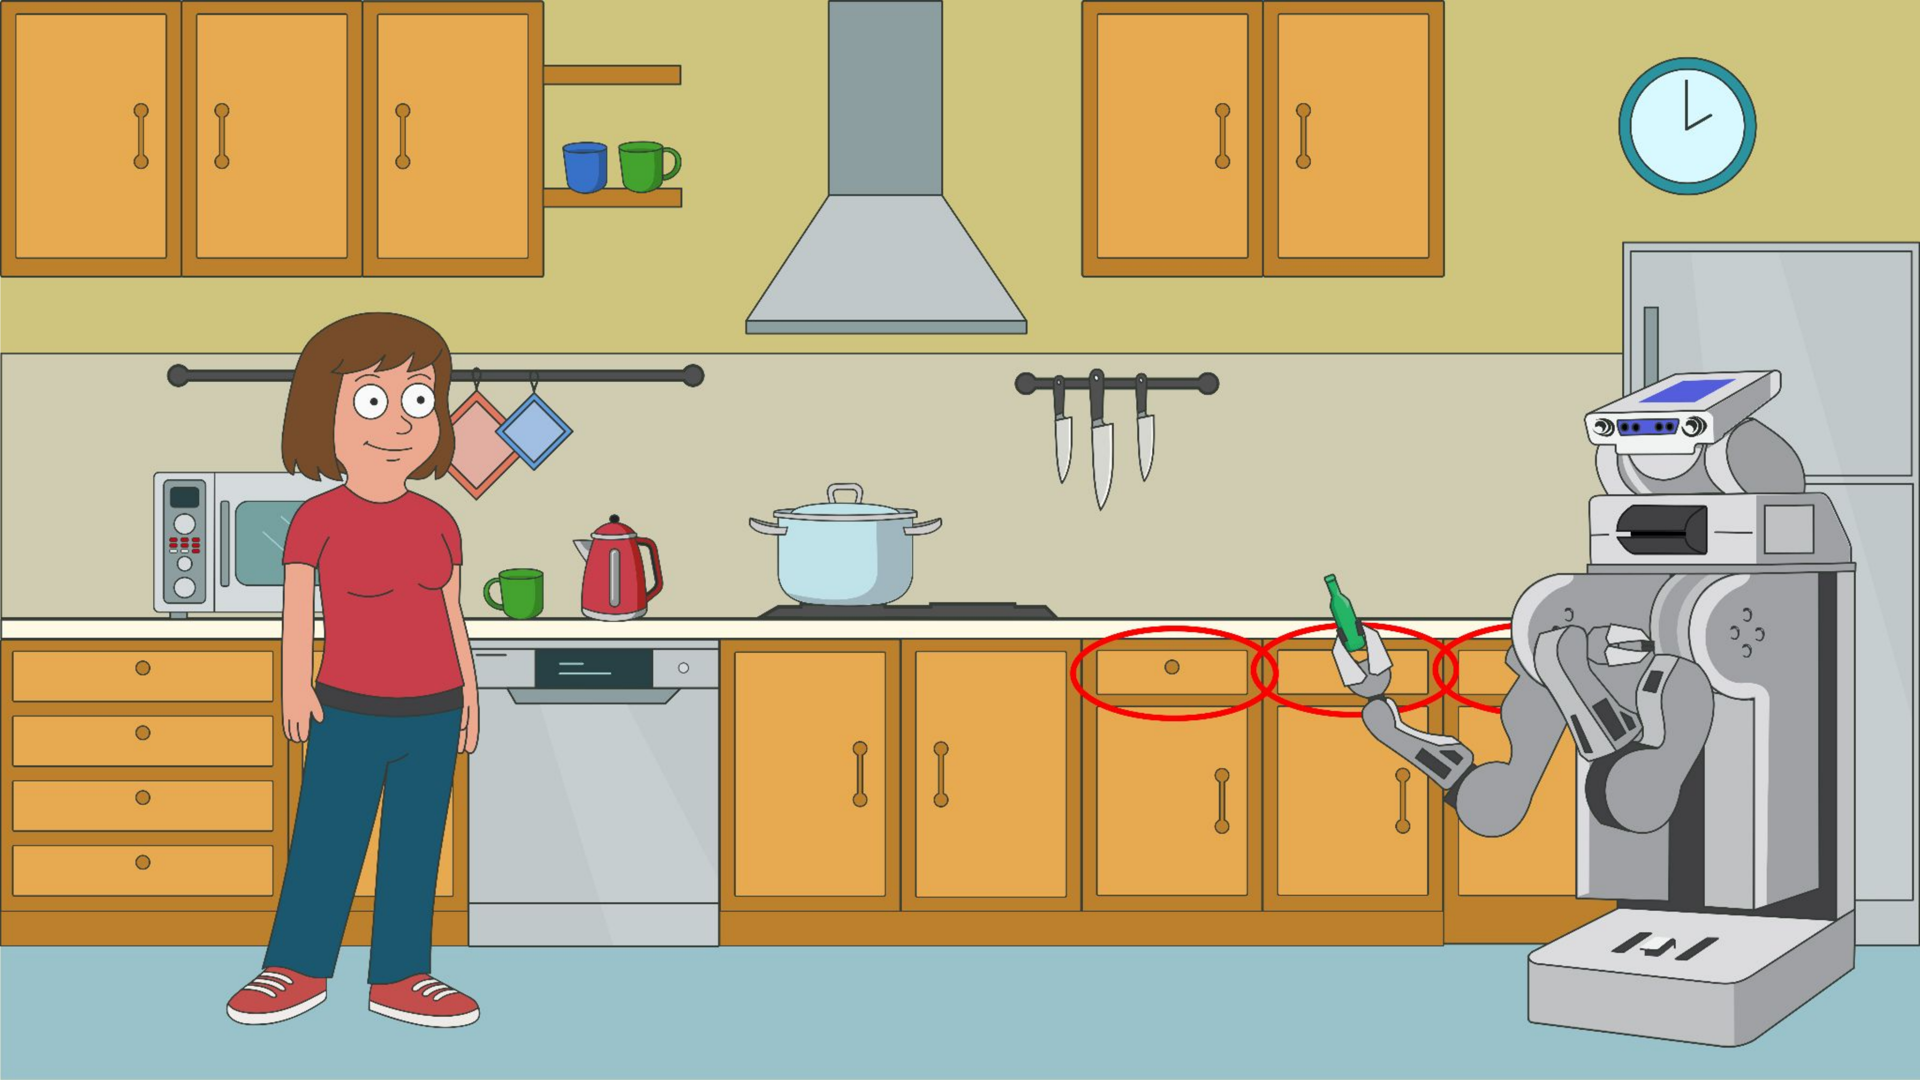

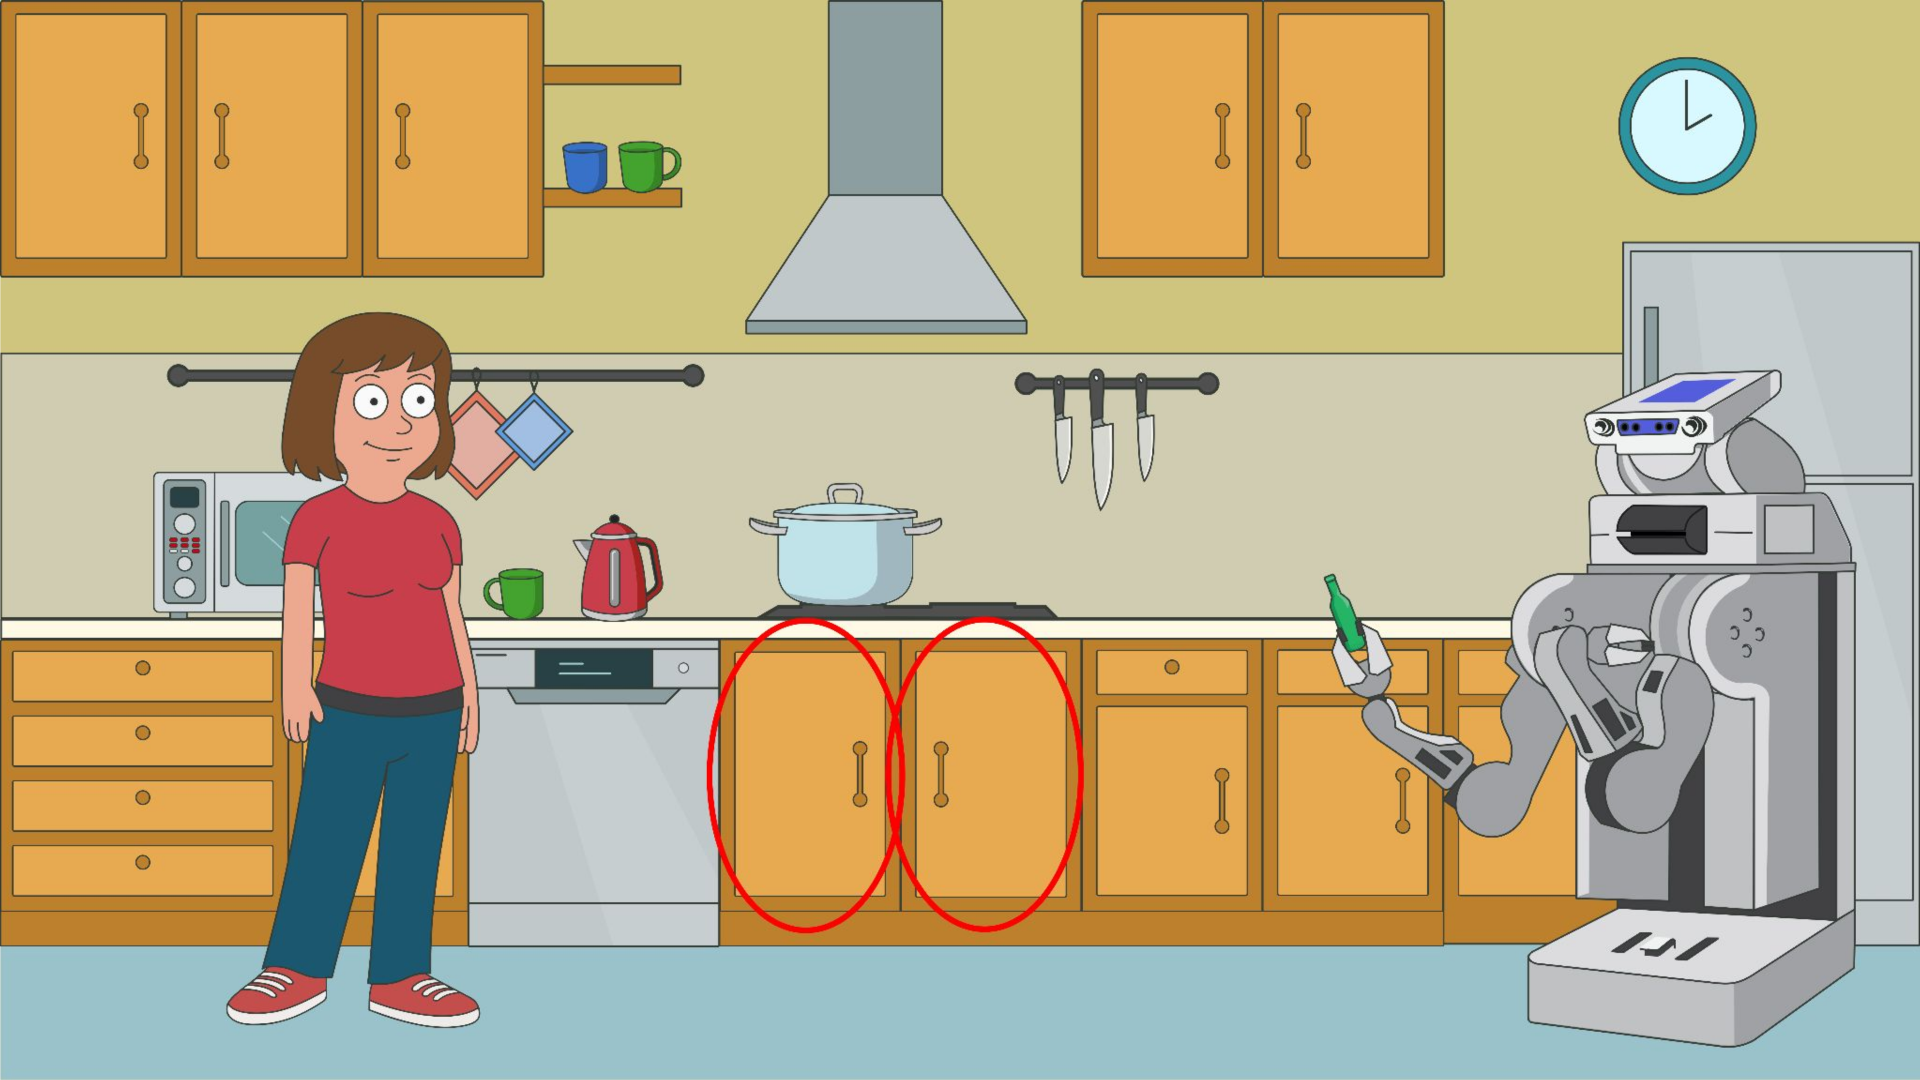

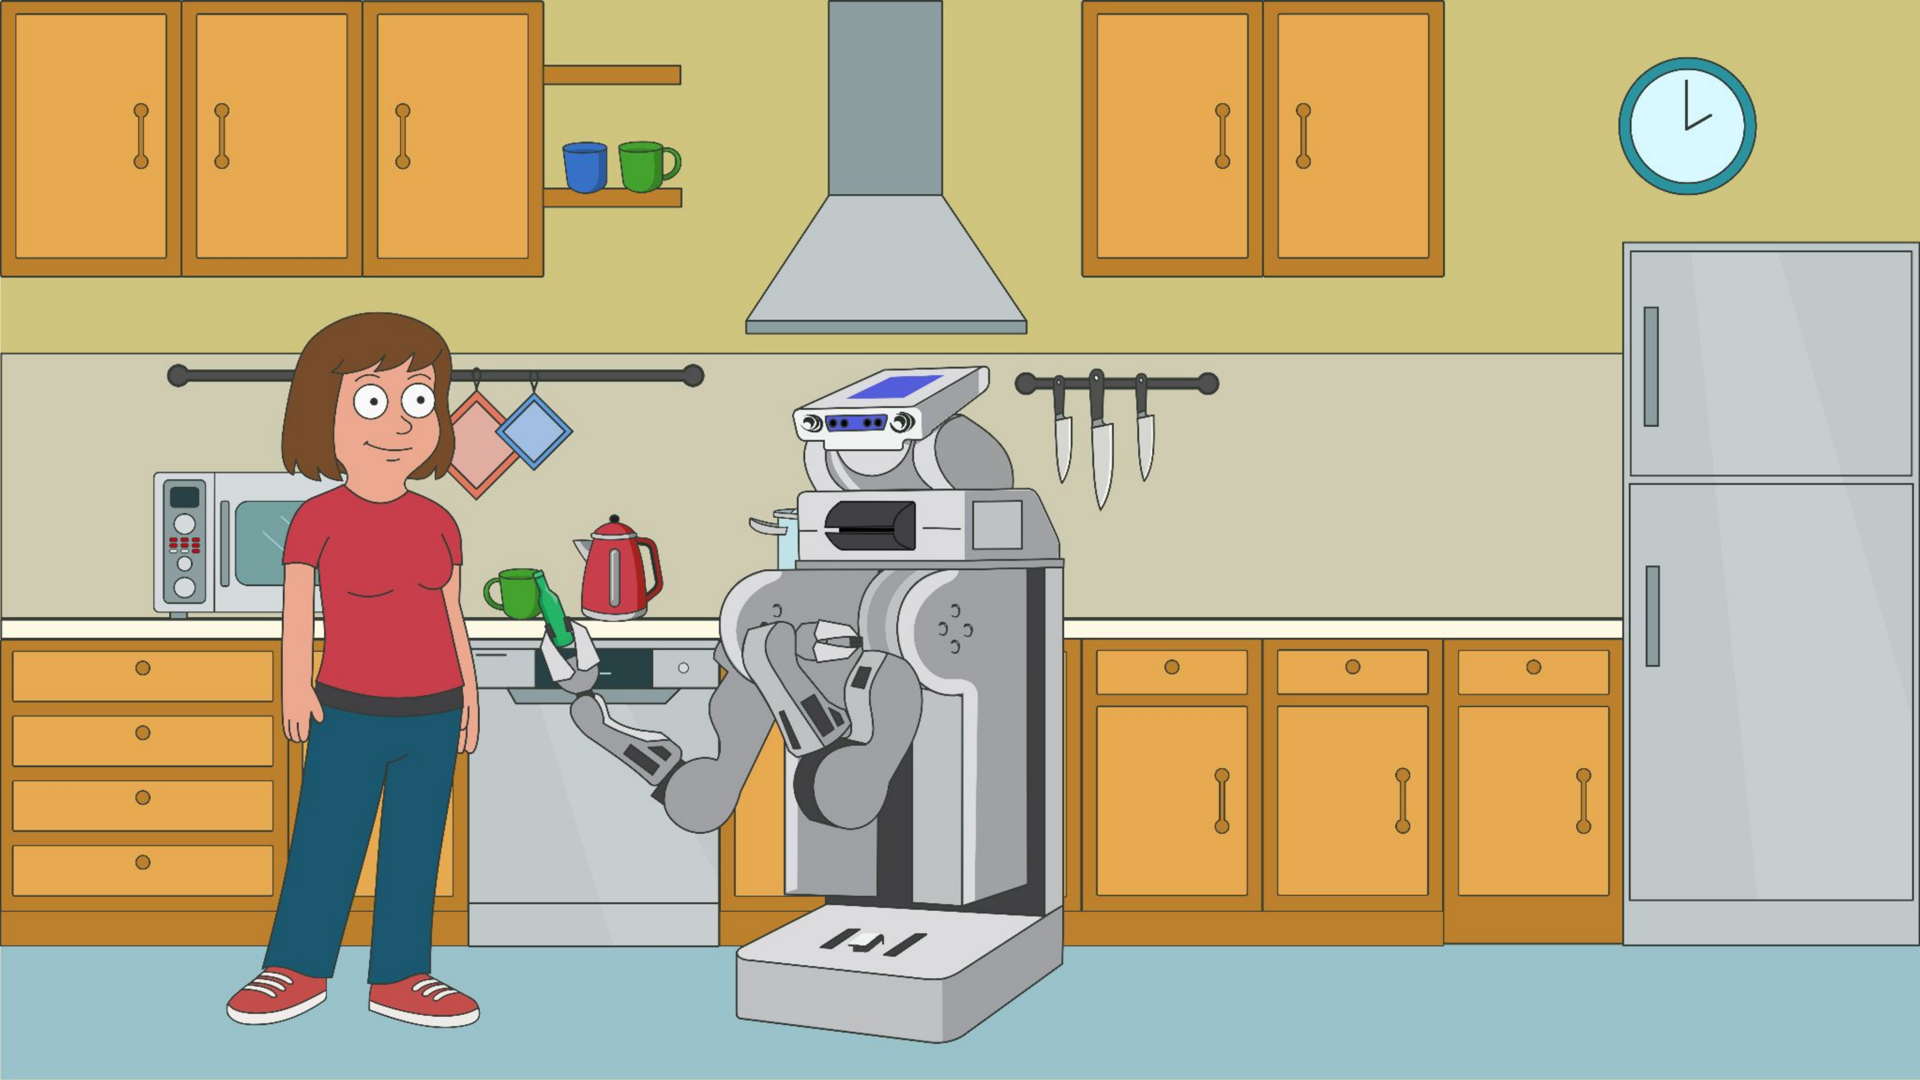

Motivation

A need of common sense knowledge

✓ LLM | Knowledge graph ✗

A need for safety and explanation

✗ LLM | Knowledge graph ✓

Motivation

Robotic and HRI requirements

Limited internet connexion

Onboard computation

Reactivity

Need for privacy

Long term interaction

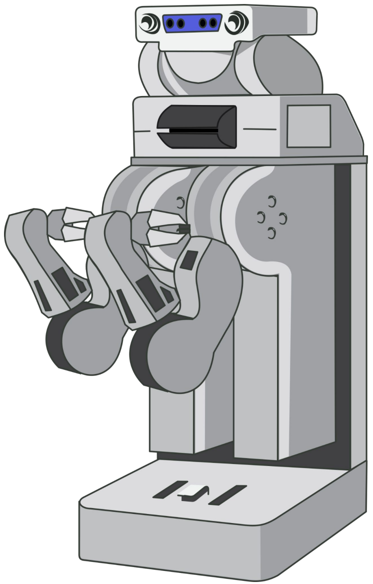

Motivation

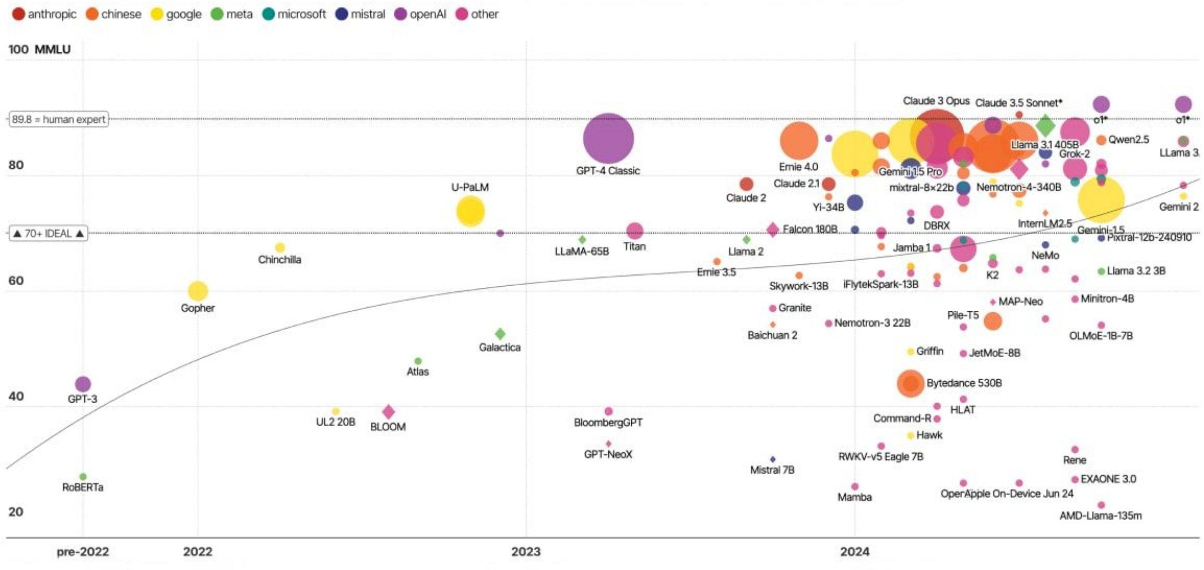

- Model size
- Input data
- Computing power

Reasoning with LLM

"**Regular**" LLM

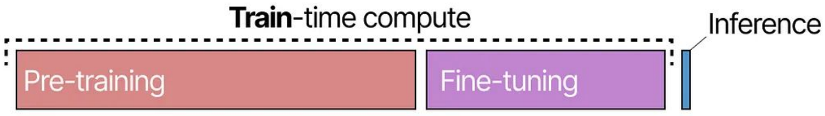

Increasing model size, data and compute quantity

"**Reasoning**" LLM

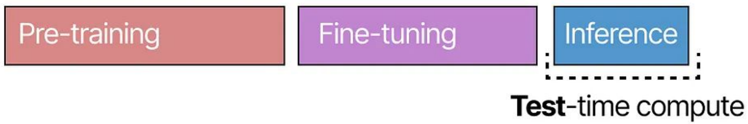

"Spending" more time (tokens) on inference time

Reasoning with LLM

“**Regular**” LLM

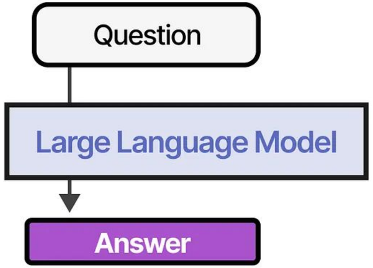

“**Reasoning**” LLM

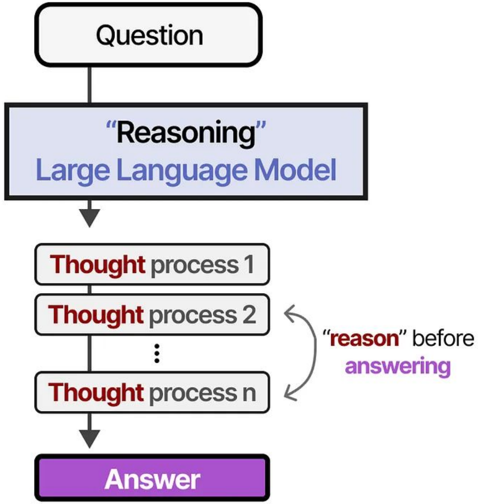

Reasoning with LLM

More time for inference

"Regular" LLM

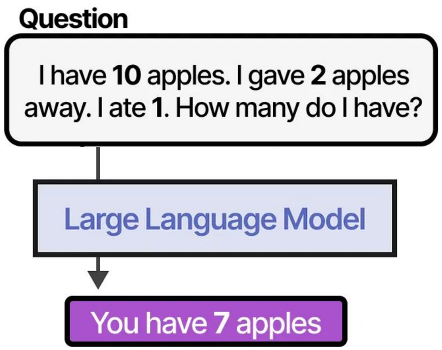

"Reasoning" LLM

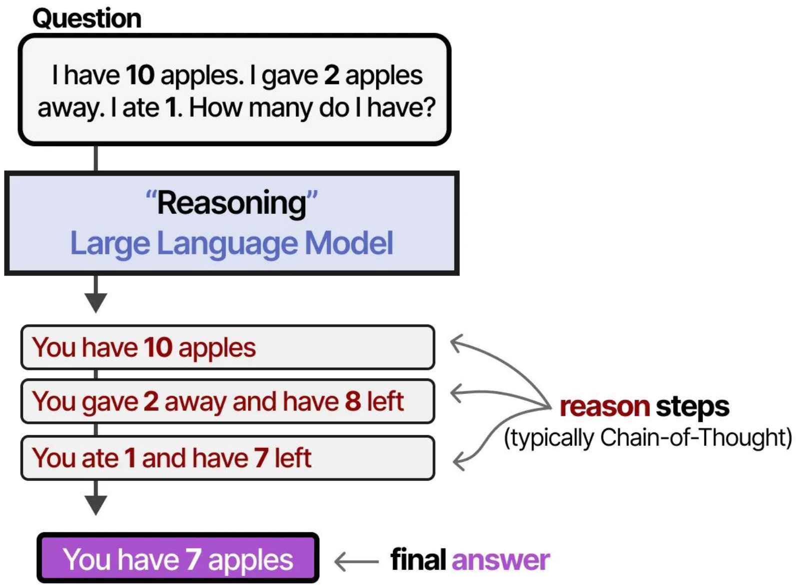

Reasoning with LLM

Input: Induce a reasoning process to the LLM

Output: Generate several inferences and judge the best

Reasoning with LLM: Search against Verifiers

1. Sampling various generations
2. Selecting the best answer (majority vote, best-of-N, beam search, etc)

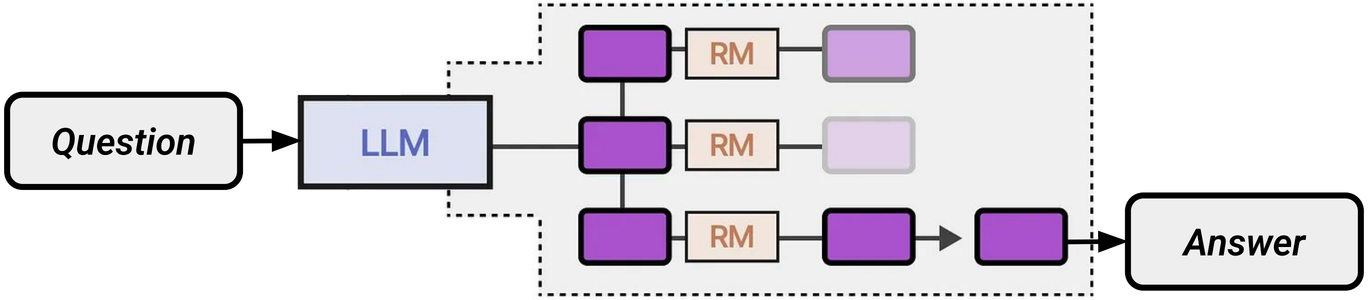

Increased correctness through **reduced statistical noise**

No need to re-train the main model

Need to train a fine-tuned LLM as **verifier** to judge

Still probabilistic output

Reasoning with LLM: Modifying Proposal Distribution

Induced with prompting techniques (in-context learning):
through N-shots, CoT[1], etc

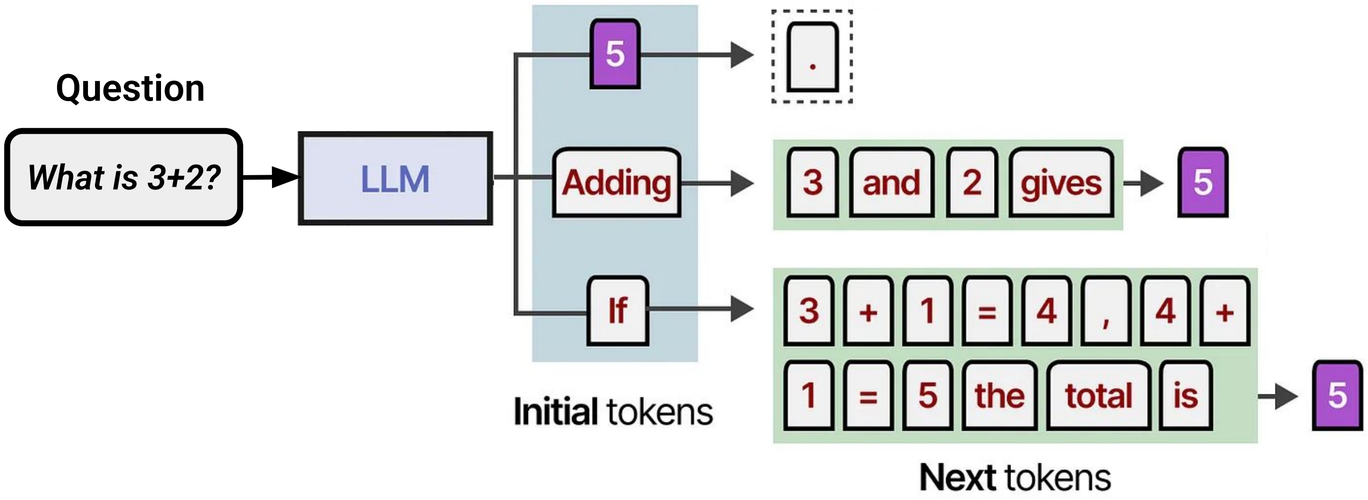

[1] Wei, J., Wang, X., et al. (2022). *Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems.*

Reasoning with LLM: Modifying Proposal Distribution

Induced with prompting techniques (in-context learning):
through N-shots, CoT[1], etc

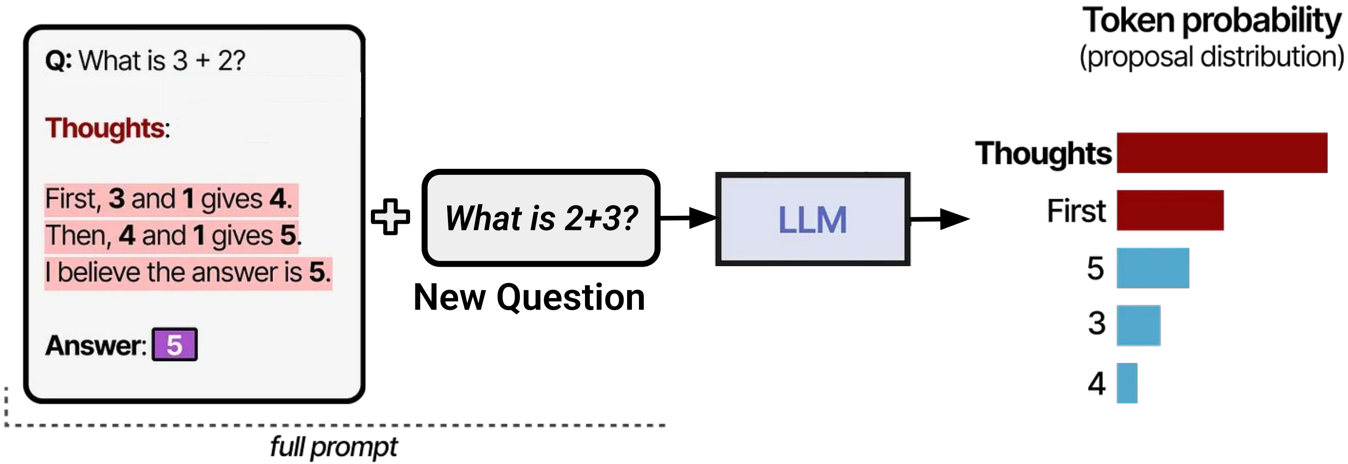

Low engineering complexity (no re-training)

Reduce brittle single-path reasoning

No correctness guarantee

Fragile, prompt-sensitive

[1] Wei, J., Wang, X., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems.

Reasoning with LLM: Modifying Proposal Distribution

Learned with re-training with (question, reasoning, answer) triplets

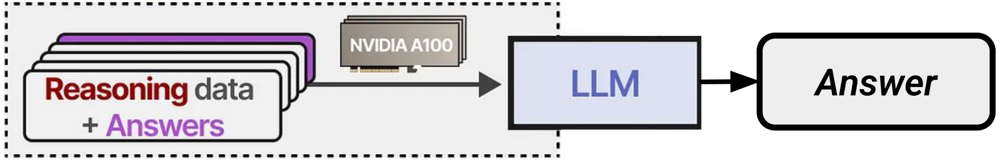

Requires Training Data

Little data available with reasoning paths

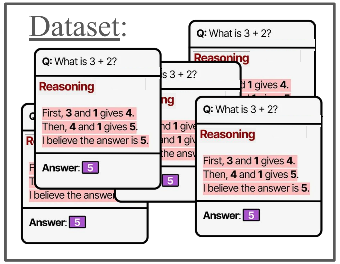

Reasoning with LLM: Modifying Proposal Distribution

Learned with re-training with (question, reasoning, answer) triplets

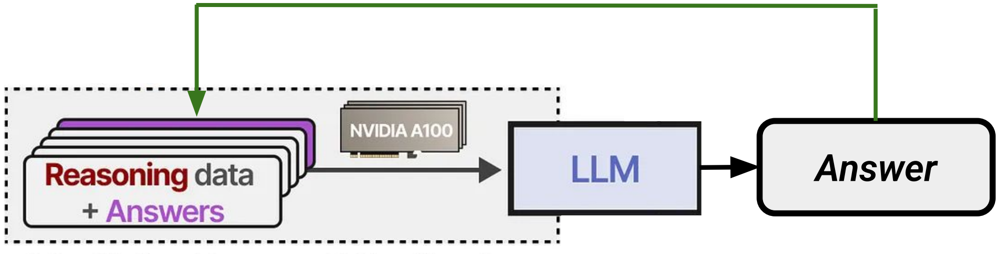

⚠️ Requires Training Data ⚠️

Little data available with reasoning paths
→ mitigated by "Self-Taught Reasoner" (STaR)[2]

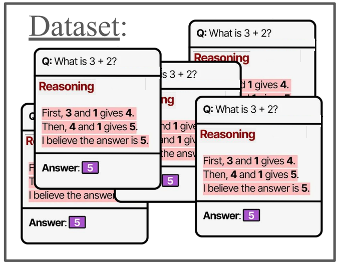

[2] Zelikman, E., Wu, Y., Mu, J., & Goodman, N. (2022). Star: Bootstrapping reasoning with reasoning. Advances in Neural Information Processing Systems.

Reasoning with LLM: Modifying Proposal Distribution

Learned with re-training with (question, reasoning, answer) triplets

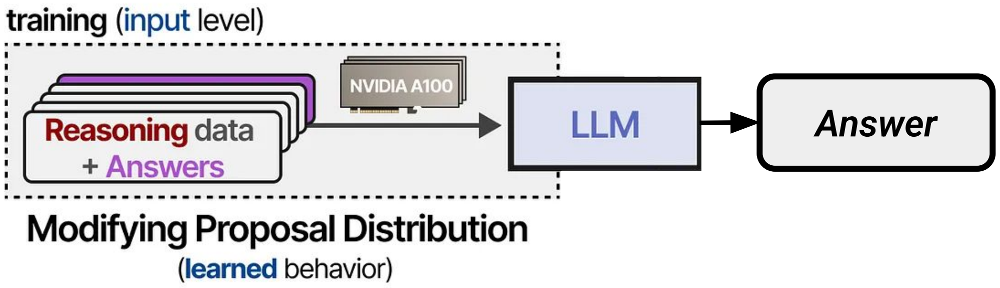

Trained systematic reasoning (explicit)

High engineering complexity (re-training)

Poor generalization and costly data

Still probabilistic output

Reasoning with LLM

Both techniques do not formally
validate the probabilistic output

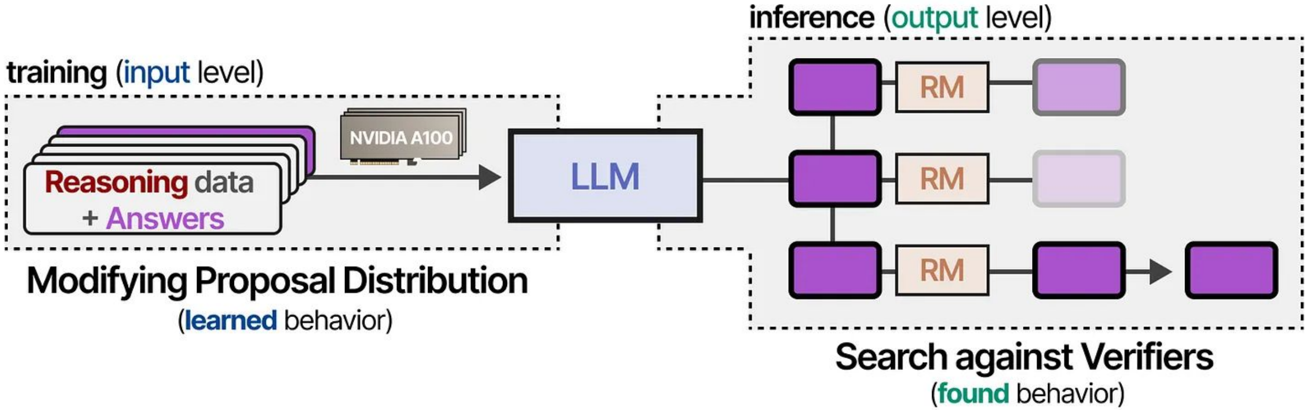

How to embed such validation in the loop?

Reasoning with LLM

Robotic and HRI requirements

Limited internet connection

Onboard computation

Reactivity → depend on the chosen reasoning technique

Need for privacy

Long term interaction

Benefits

Increased explainability
Increased correctness

Limitations

Probabilistic output → No guarantee
No formal validation

Agentic LLM

Compensate for their disadvantage through external **tools**, **memory**, and **retrieval** systems

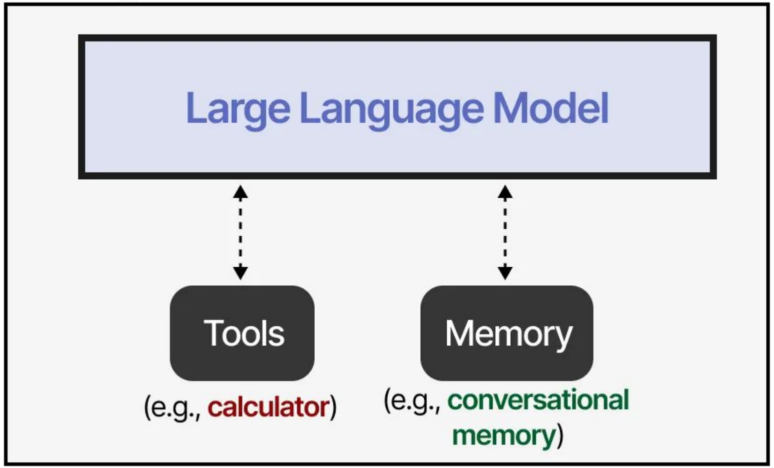

Agentic LLM: Short-Term Memory

Recent memory in Conversation History (Context window)

Up-to-date information matching the conversation

Useful for specialized LM

Limited Context Window Length

Privacy concerns from shared user history among multiple users

Relevance of "diluted" information

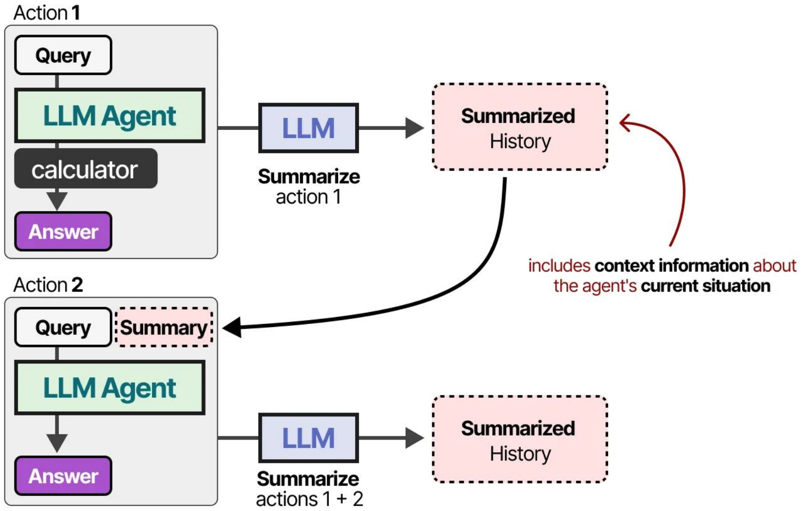

Agentic LLM: Long-Term Memory

1. Store each of the conversations in a database
2. Embed the new prompt and output the closest conversation

Up-to-date information matching the conversation

Unlimited Storage of Data

Relevance of the retrieved information decreases with quantity

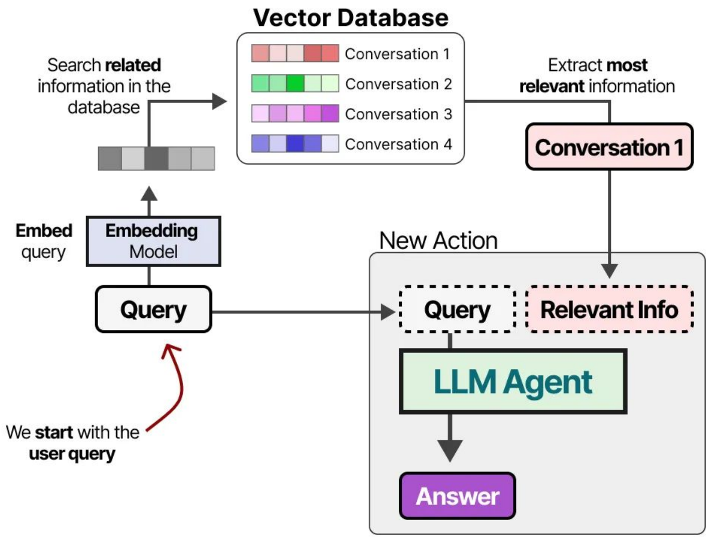

Agentic LLM: Tools

Tools allow to interact with external systems
(get data or perform action)

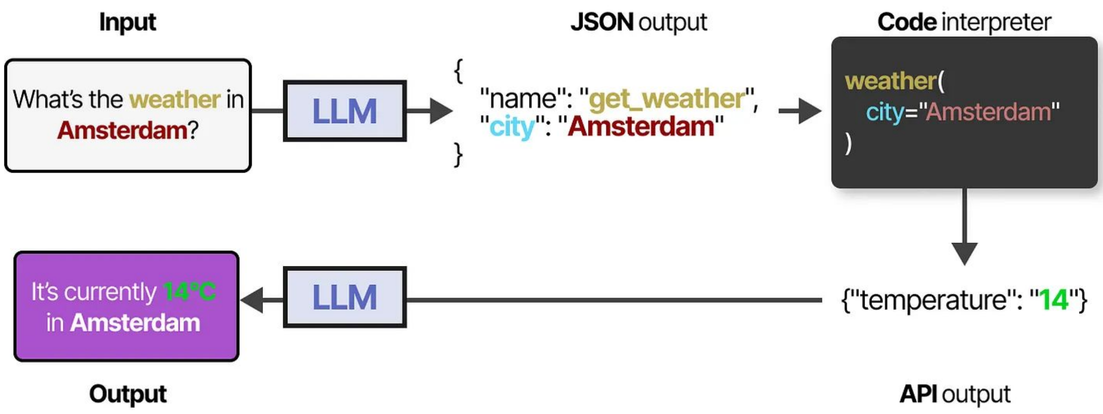

Fetch up-to-date information
from several sources

Can interact with model-based systems

Prevent LLM corruption of tool-acquired information?

Agentic LLM

From reasoning as a sequence of steps, to planning as a sequence of actions

Planning: Provide intermediary steps

Tools: Fetch or Validate information

Memory: Store/Retrieve additional knowledge

High computational cost (multiple specialized models running concurrently)

Probabilistic decision-making lacks guaranteed outcomes

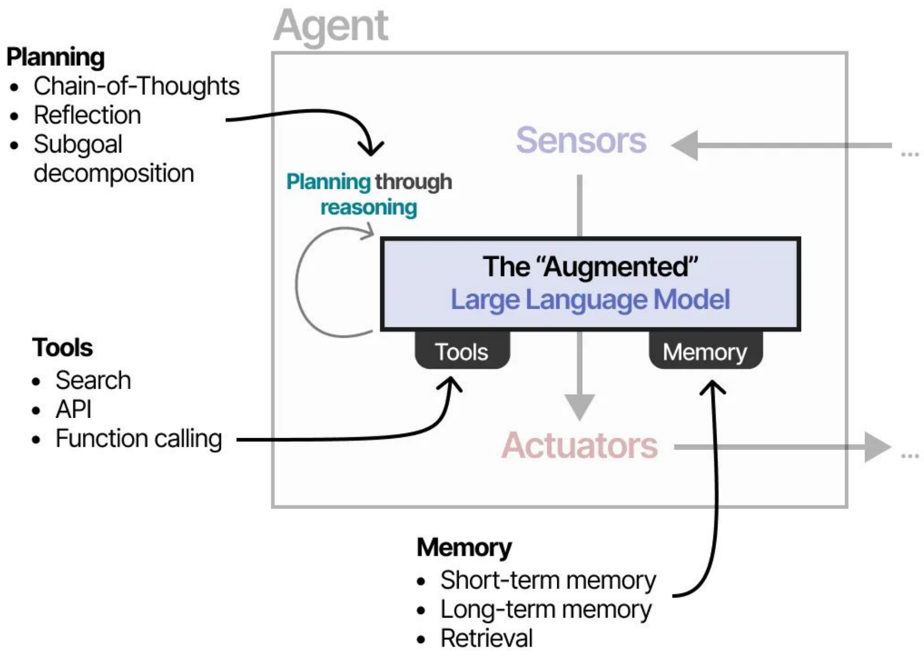

Agentic LLM: Validation?

Inference Rule[3]

(hasCapability some Grasping)(?a), (hasDisposition some Graspable)(?o), isReachableBy(?o,?a), hasEndEffector(?a,?g), **hasOpeningWidth(?g,?w1)**, **hasHoldingPartWidth(?o,?w2)**, greaterThan(?w1,?w2) → canGrasp(?a, ?o)

Human-expert

"The robot can grasp the suitcase because it has a grasping capability thanks to its gripper, which has an opening width greater than the **suitcase's holding part**'s **width** and the suitcase being reachable."

LLM-based

"The robot can grasp the object because it has a **gripper** that can grasp objects and **its holding part width** is greater than the object's width."

Validation of consistent output required

[3] *Dussard, B., Clodic, A., & Sarthou, G. (2025). Evaluating Embeddable Language Models in Verbalizing Rule-based Inferences through Justifications. In 34th Edition-IEEE RO-MAN.*

Agentic LLM: Multi agent

Each agent could be using a specific set of tools
One supervisor organizes the team

Reduces risk of generic models being inefficient for specific tasks

Parallelism for increased speed

Increased computation footprint

Agent communication validity is not guaranteed (uncertainty amplifies)

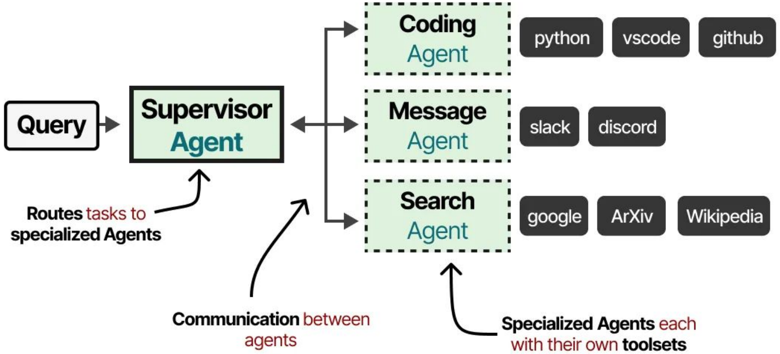

Agentic LLM

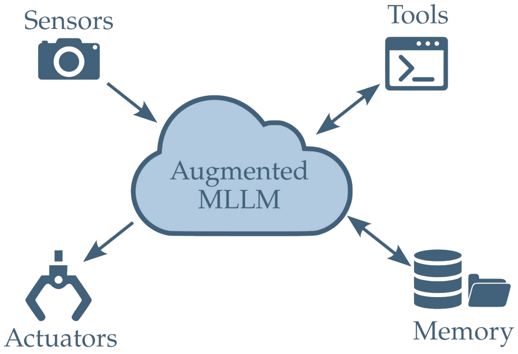

Benefits

Autonomous access to external resources
Reduced hallucinations thanks to memory + tools

Limitations

Higher footprint for Team of LLMs
LLM corruption of tool-acquired information?

Agentic LLM

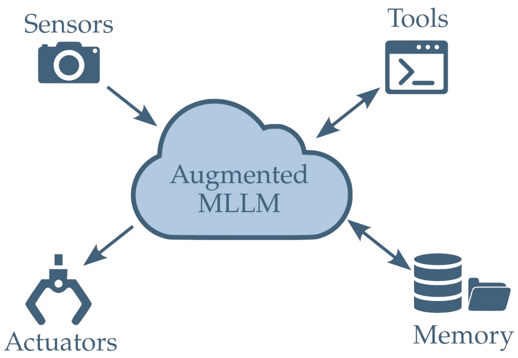

Robotic and HRI requirements

Onboard Computation
(how many LMs are running)

Reactivity
(length of reasoning chain)

Long term interaction
(requires storing personal data)

Personal Data Privacy
(jeopardized if access to Internet)

Limited Internet Connexion

Paradigm shift

LM as a cognitive tool

Generate heuristics + Help to acquire new knowledge

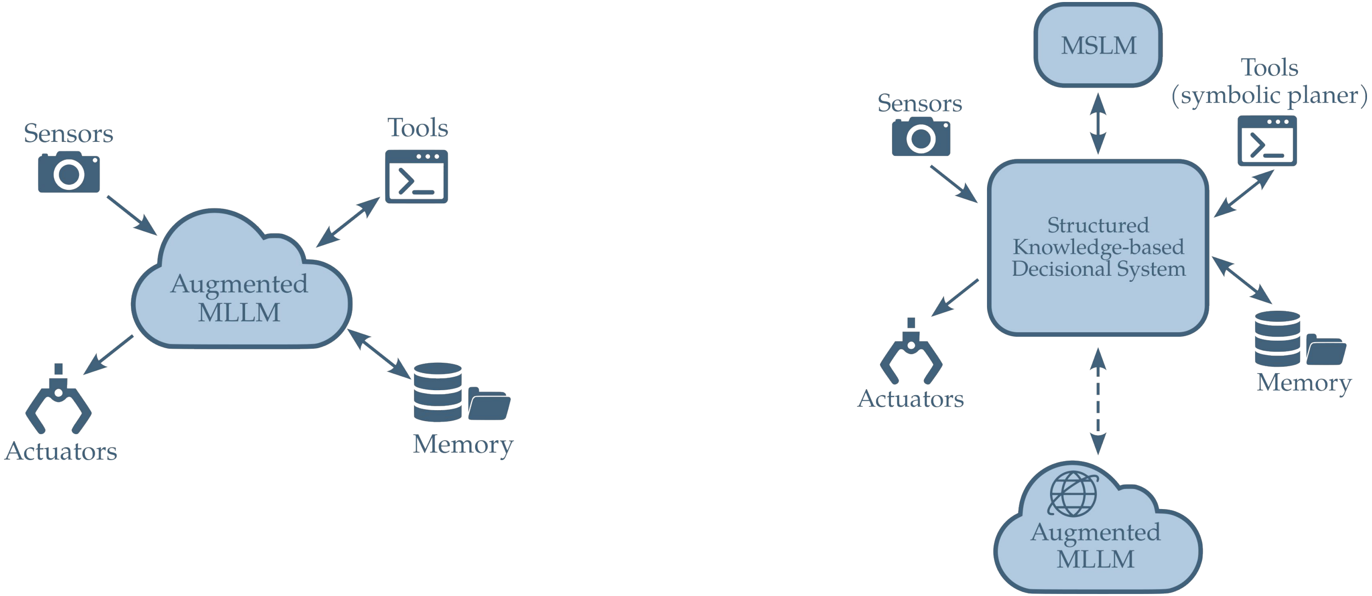

Increase robustness and explainability

Paradigm shift

1. Proactive identification of knowledge gaps.

2. Augmented querying where the query is enriched with existing knowledge and rules to ensure contextual consistency.

3. Transformation of unstructured responses into formal schemas and insertion into the knowledge graph.

4. Verification and identification of conflicts, and iterative resolution of these inconsistencies, potentially via iterations with the LLM.

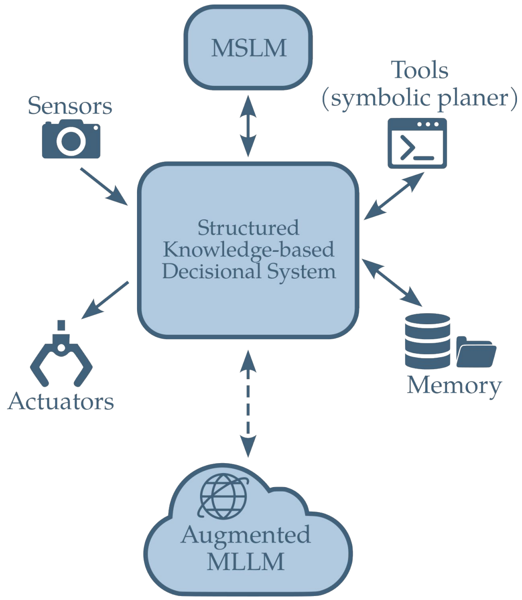

Paradigm shift

Learning structured knowledge

Reducing LM need and query

Doing the same with Human interaction

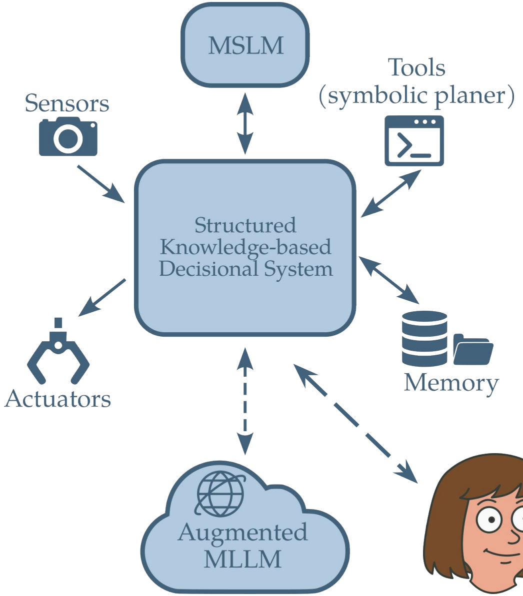

Paradigm shift

Need to think
about
long term interaction
and
respect for privacy

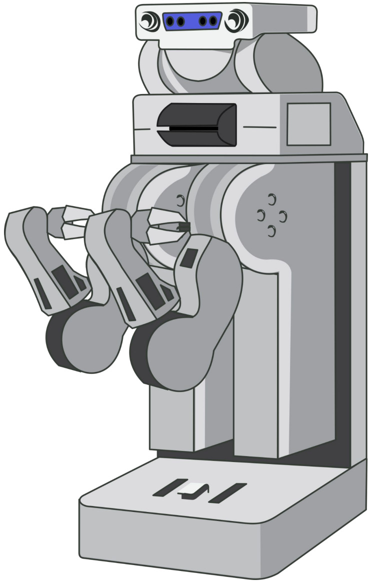

References

[1] Wei, J., Wang, X., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems.

[2] Zelikman, E., Wu, Y., Mu, J., & Goodman, N. (2022). Star: Bootstrapping reasoning with reasoning. Advances in Neural Information Processing Systems.

[3] Dussard, B., Clodic, A., & Sarthou, G. (2025). Evaluating Embeddable Language Models in Verbalizing Rule-based Inferences through Justifications. In 34th Edition-IEEE RO-MAN.