

DSC291: Machine Learning with Few Labels

Enhancing Large Language Models: Overview

Zhiting Hu

Lecture 10, April 22, 2024

UC San Diego

HALICIOĞLU DATA SCIENCE INSTITUTE

Outline: Enhancing the Backend Beyond LMs


- Richer learning mechanisms
 - Learning with Embodied Experiences
 - Social Learning
- Multi-modal capabilities
- **Latent-space reasoning**
- Agent models with external augmentations (e.g., tools)

Latent-space Reasoning

- What's the best space for carrying out reasoning?
 - Natural language space?
 - Raw sensory space (e.g., video)?
 - **Learned** latent space?
 - Single-level / multi-level latent space?
- Consider a long-term planning problem, e.g., economic planning for U.S. in 2024
 - Extremely complex, long-horizon reasoning
 - Inefficient/infeasible with LLM token-by-token reasoning or Video Model frame-by-frame reasoning
- Multi-level latent spaces are needed for multi-granularity reasoning

Latent-space Reasoning

- But how to learn a good latent space in the first place?



Outline: Enhancing the Backend Beyond LMs

- Richer learning mechanisms
 - Learning with Embodied Experiences
 - Social Learning
- Multi-modal capabilities
- Latent-space reasoning
- **Agent models with external augmentations (e.g., tools)**

Agent models with external augmentations


- External augmentations for added capabilities:
 - **Tools:** telescope, vehicles, ...
 - **Data about a skill:** demonstration videos of climbing a snowy mountain
 - **Knowledge bases:** domain knowledge

LLMs need external tools for real-world tasks



The original price of MacBook Air is \$1580. Can you help me purchase it when it gets 10% off?

Sorry, but this is beyond my capabilities as a language model...




LLMs need external tools for real-world tasks

Lacking the abilities for



The original price of MacBook Air is \$1580. Can you help me purchase it when it gets 10% off?

Sorry, but this is beyond my capabilities as a language model...



LLMs need external tools for real-world tasks


Lacking the abilities for

Accurate math calculation



The original price of MacBook Air is **\$1580**. Can you help me purchase it when it gets **10%** off?

Sorry, but this is beyond my capabilities as a language model...



LLMs need external tools for real-world tasks


Lacking the abilities for

- Accurate math calculation



The original price of MacBook Air is \$1580. Can you help me purchase it when it gets 10% off?

Sorry, but this is beyond my capabilities as a language model...



LLMs need external tools for real-world tasks

Lacking the abilities for


- Accurate math calculation

Up-to-date knowledge



The original price of MacBook Air is \$1580. Can you help me purchase it **when it gets 10% off?**

Sorry, but this is beyond my capabilities as a language model...



LLMs need external tools for real-world tasks


Lacking the abilities for

- Accurate math calculation
- Accessing up-to-date knowledge



The original price of MacBook Air is \$1580. Can you help me purchase it when it gets 10% off?

Sorry, but this is beyond my capabilities as a language model...



LLMs need external tools for real-world tasks

Real-world actions


Lacking the abilities for

- Accurate math calculation
- Accessing up-to-date knowledge



The original price of MacBook Air is \$1580. Can you help me **purchase it** when it gets 10% off?

Sorry, but this is beyond my capabilities as a language model...



LLMs need external tools for real-world tasks


Lacking the abilities for

- Accurate math calculation
- Accessing up-to-date knowledge
- Taking real-world actions






The original price of MacBook Air is \$1580. Can you help me purchase it when it gets 10% off?

Sorry, but this is beyond my capabilities as a language model...



LLMs need external tools for real-world tasks


Augmenting language models with tools will help unlock those abilities!

- Accurate math calculation  Calculator
- Accessing up-to-date knowledge  Database
- Taking real-world actions  API/Robot
- ...






The original price of MacBook Air is \$1580. Can you help me purchase it when it gets 10% off?

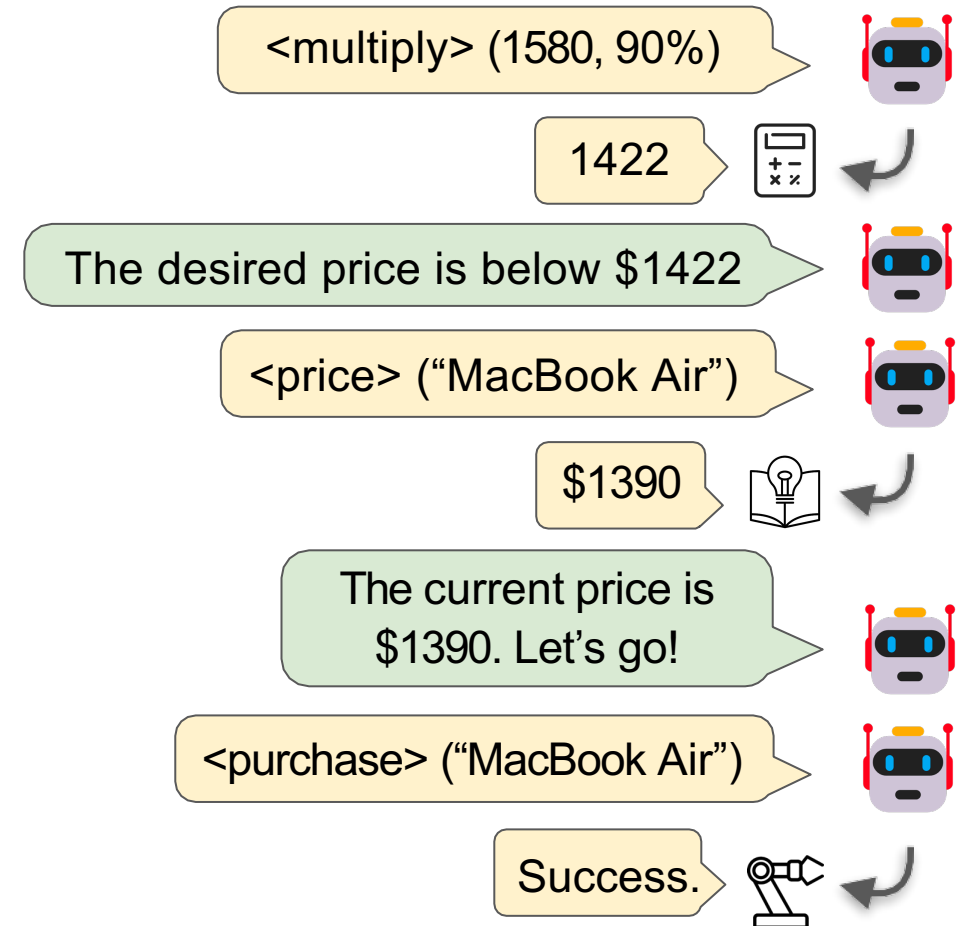
Sorry, but this is beyond my capabilities as a language model...



LLMs need external tools for real-world tasks

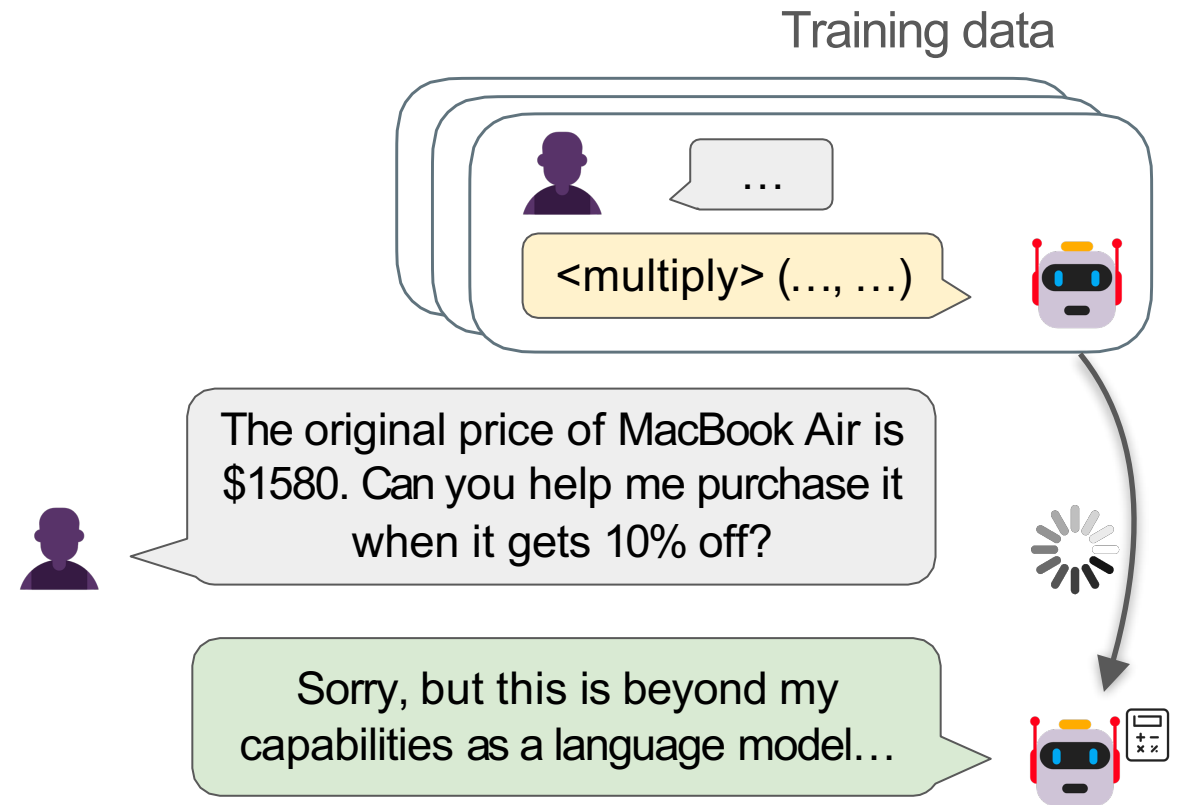
Augmenting language models with tools will help unlock those abilities!

- Accurate math calculation  Calculator
- Accessing up-to-date knowledge  Database
- Taking real-world actions  API/Robot
- ...



Teaching LLMs to Use Tools - Method #1: Fine-tuning

Train the LLM with the demonstrations of tool calling



Talm: Tool augmented language models [Parisi et al., 2022]

Toolformer: Language models can teach themselves to use tools [Schick et al., 2023]

Teaching LLMs to Use Tools - Method #1: Fine-tuning

Train the LLM with the demonstrations of tool calling

Limitations:

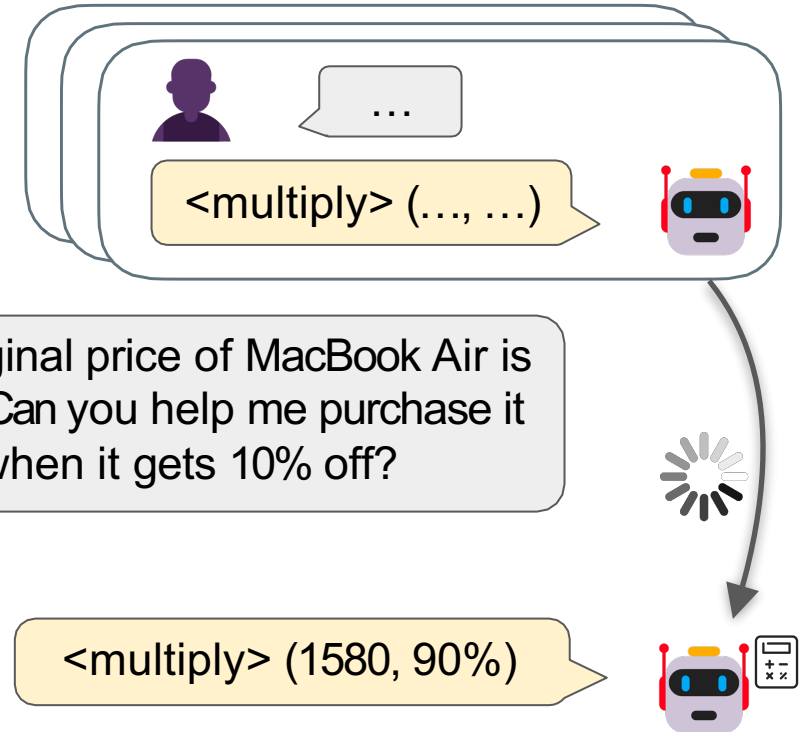
- **Not Frozen LLMs:** Fine-tuning an LLM is expensive 💰
- **Not Plug-and-play:** Once we want to add, delete or update a tool, the LLM needs to be **re-trained** 🔄



The original price of MacBook Air is \$1580. Can you help me purchase it when it gets 10% off?

<multiply> (1580, 90%)

Training data

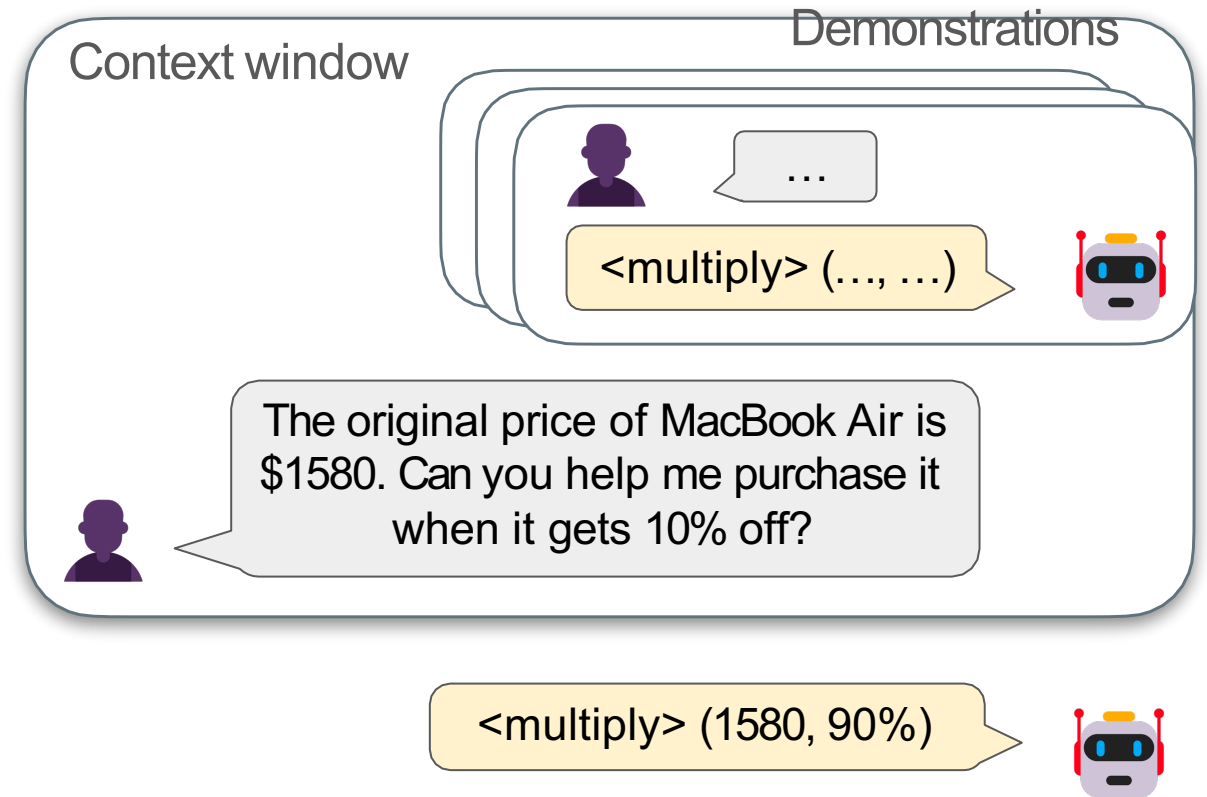


Talm: Tool augmented language models [Parisi et al., 2022]

Toolformer: Language models can teach themselves to use tools [Schick et al., 2023]

Teaching LLMs to Use Tools - Method #2: Demonstrations

Prompting LLMs with demonstrations of tool calling



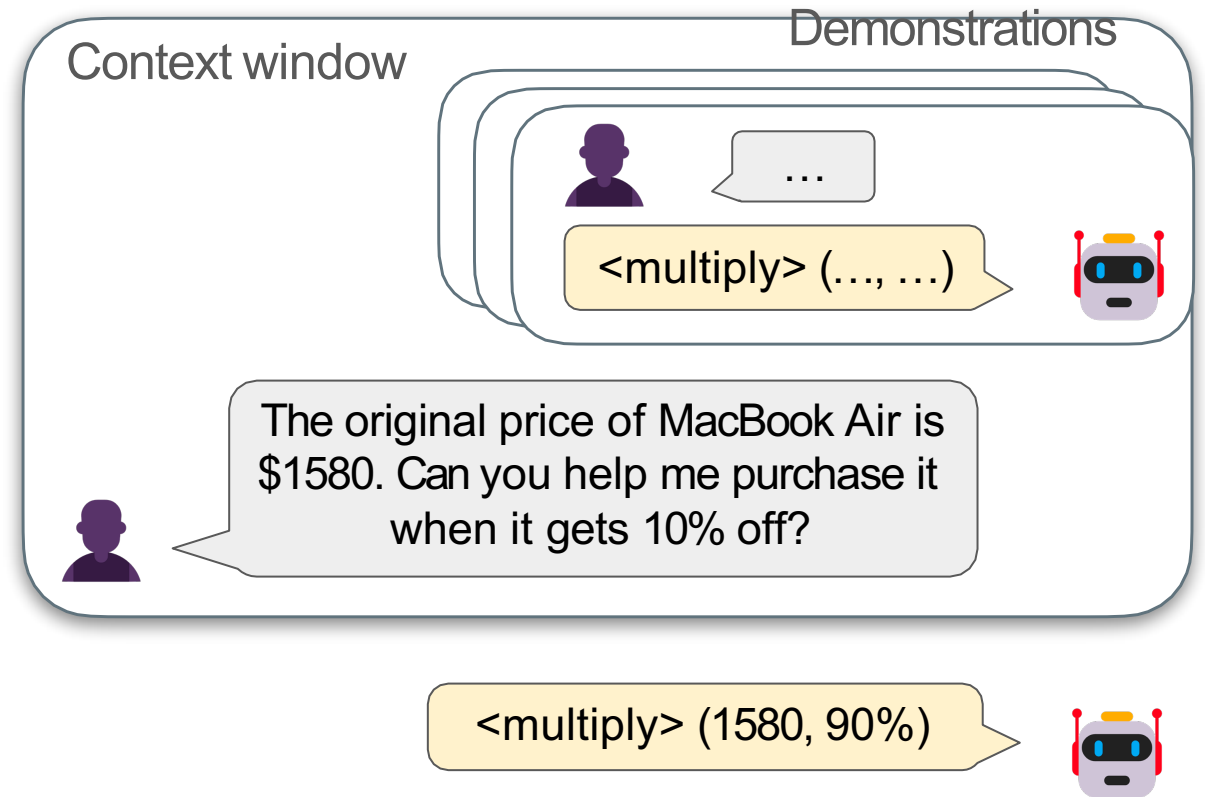
ReAct: Synergizing Reasoning and Acting in Language Models [Yao et al., 2023]
Gorilla: Large language model connected with massive apis [Patil et al., 2023]

Teaching LLMs to Use Tools - Method #2: Demonstrations

Prompting LLMs with demonstrations of tool calling

Limitations:

- **Shallow Understanding:** Can only learn from surface text instead of large-scale data 🤔
- **Limited tools:** struggles with a large tool set 🧰

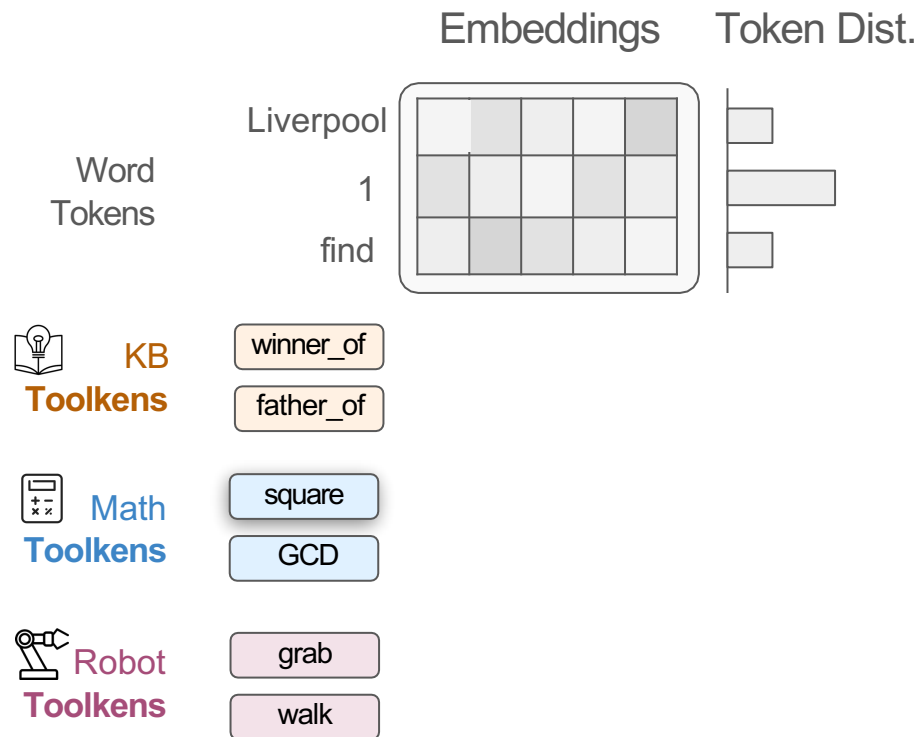


ReAct: Synergizing Reasoning and Acting in Language Models [Yao et al., 2023]
Gorilla: Large language model connected with massive apis [Patil et al., 2023]

Teaching LLMs to Use Tools - Method #3: Toolken

Step 1: Next token/toolken prediction

Adding **Toolkens** to the vocabulary



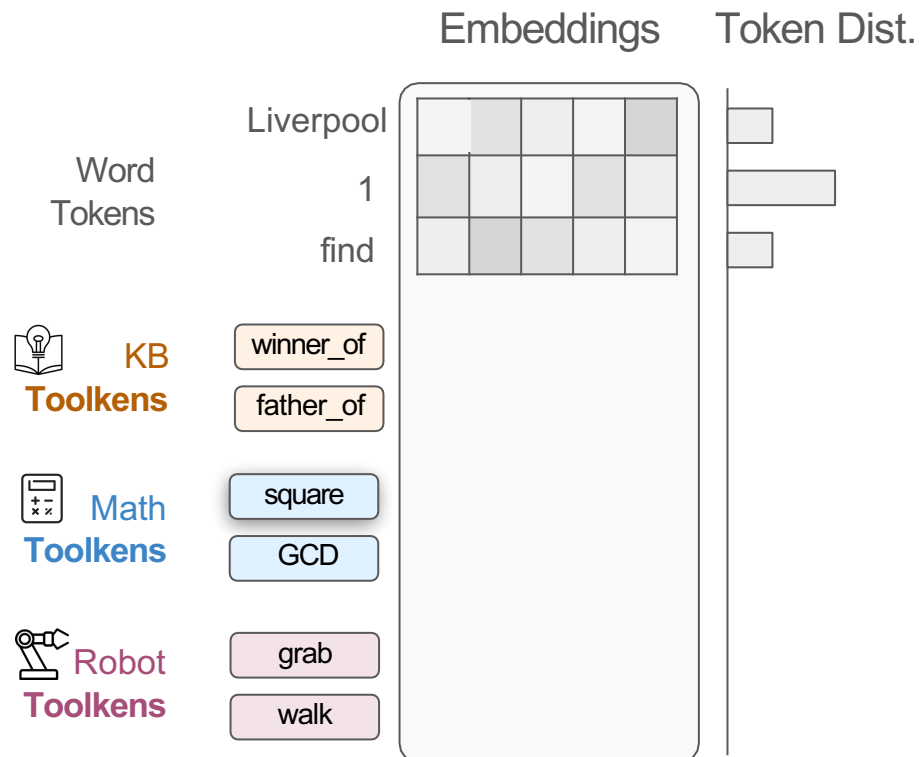
Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is 16 meters. Therefore, the area is ____

Teaching LLMs to Use Tools - Method #3: Toolken

Step 1: Next token/toolken prediction

Adding **Toolkens** to the vocabulary




Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is 16 meters. Therefore, the area is ____

Teaching LLMs to Use Tools - Method #3: Toolken

Step 1: Next token/toolken prediction

Adding **Toolkens** to the vocabulary



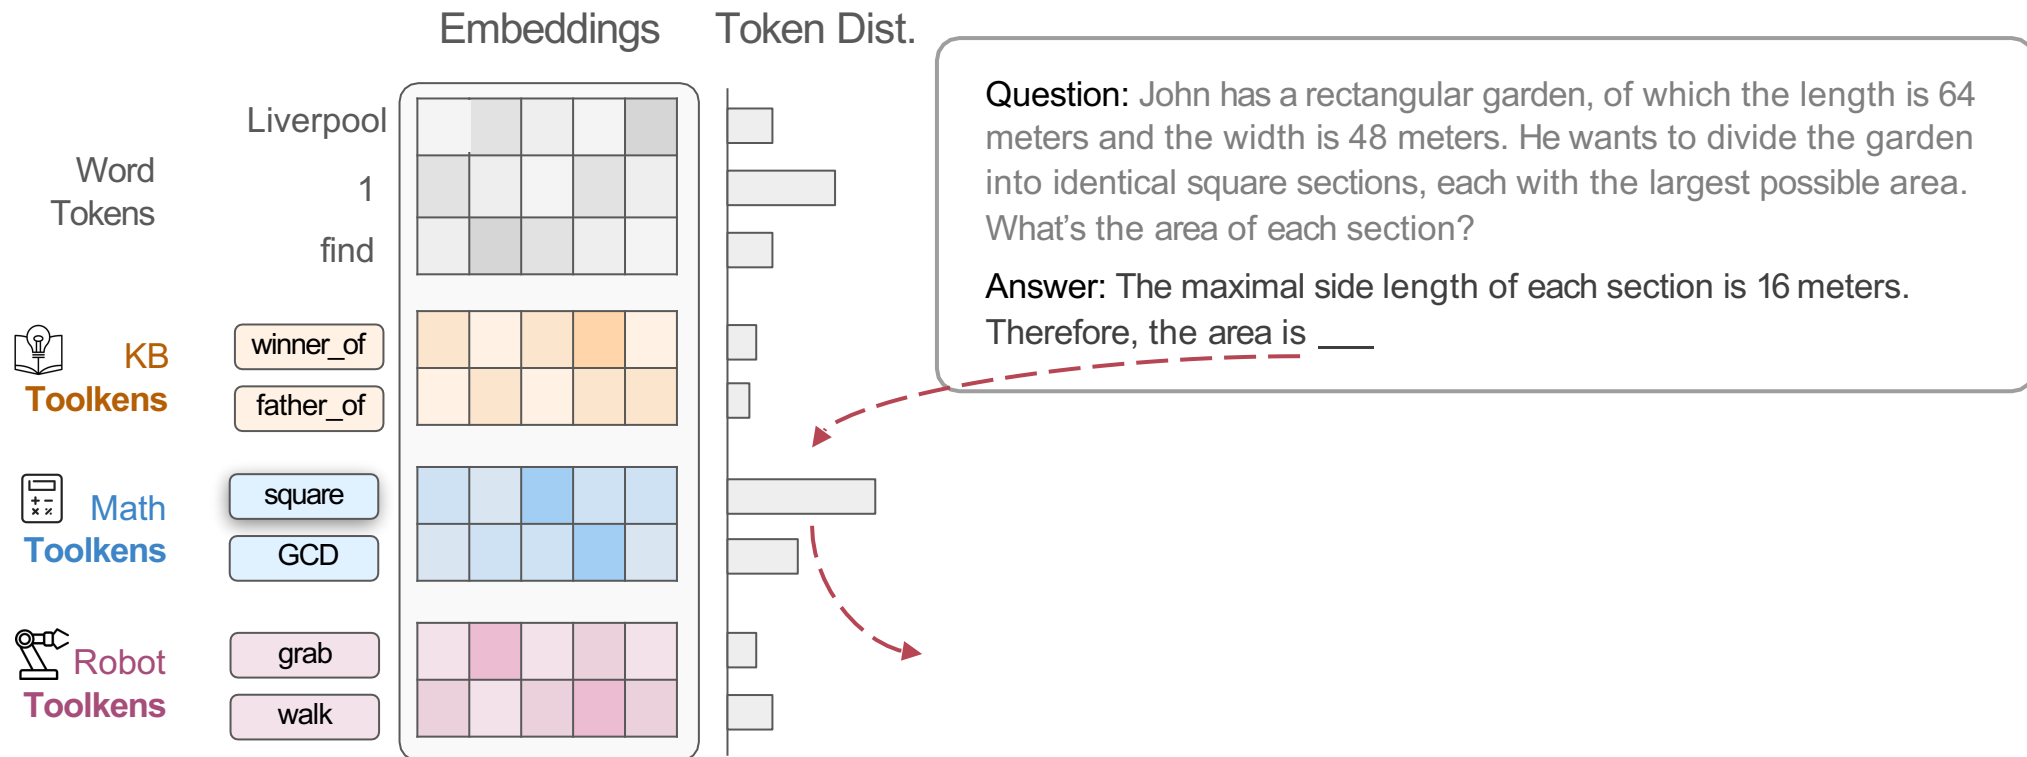
Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is 16 meters. Therefore, the area is ____

Teaching LLMs to Use Tools - Method #3: Toolken

Step 1: Next token/toolken prediction

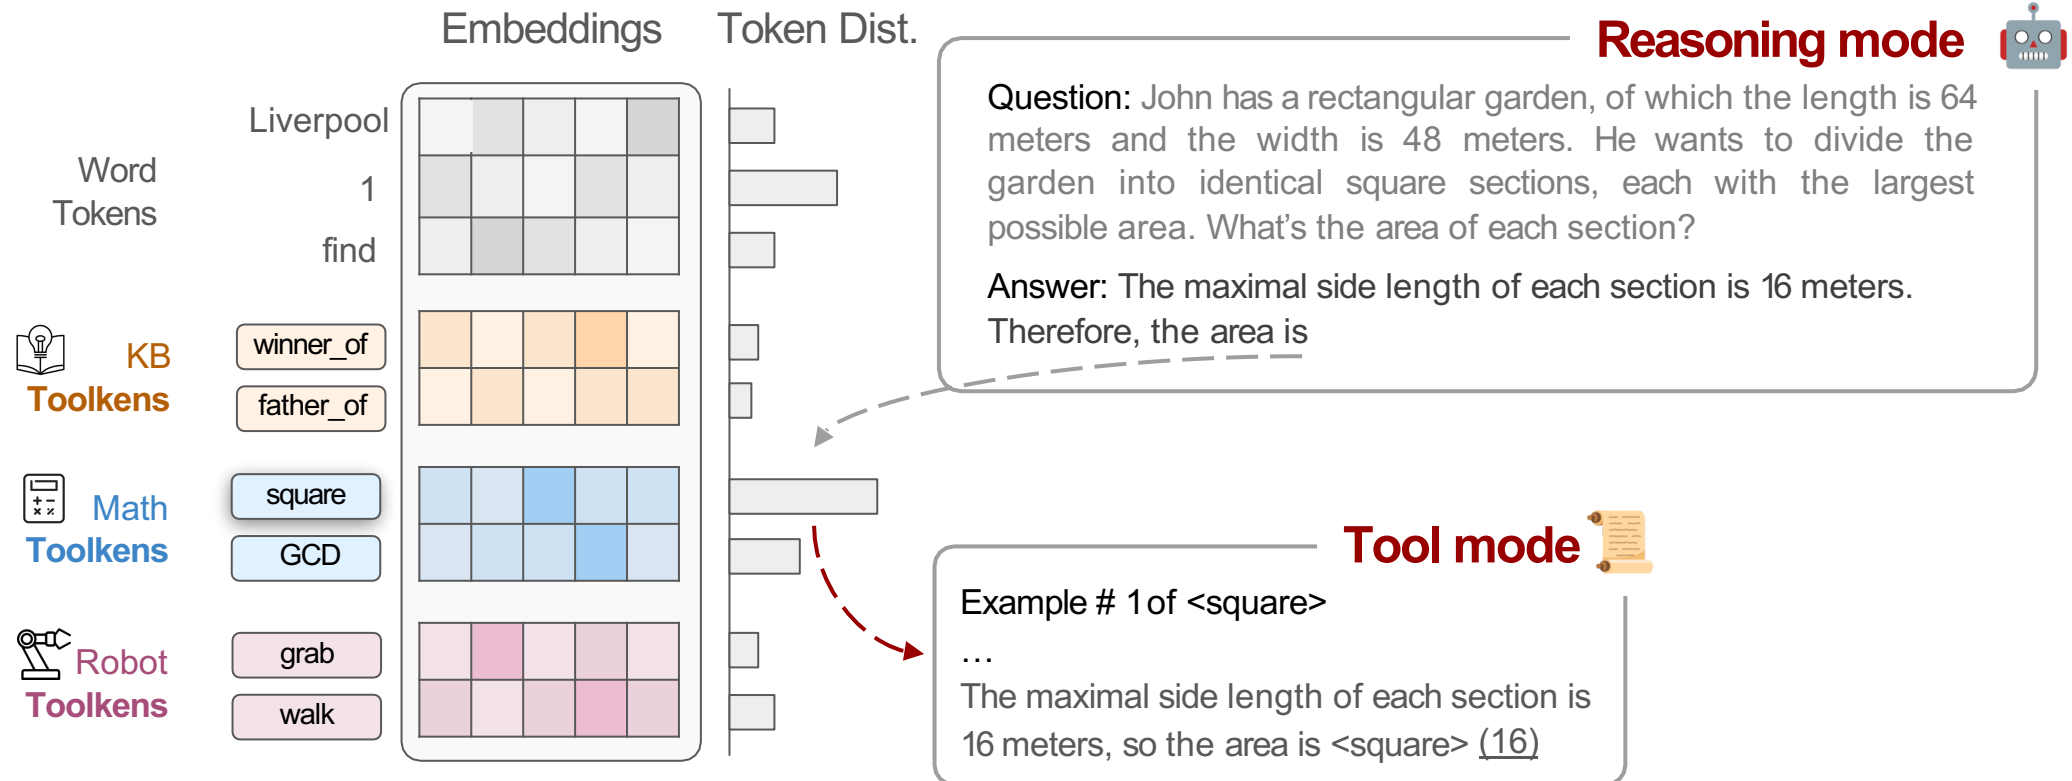
Adding **Toolkens** to the vocabulary



Teaching LLMs to Use Tools - Method #3: Toolken

Step 2: Argument prediction in a separate tool mode

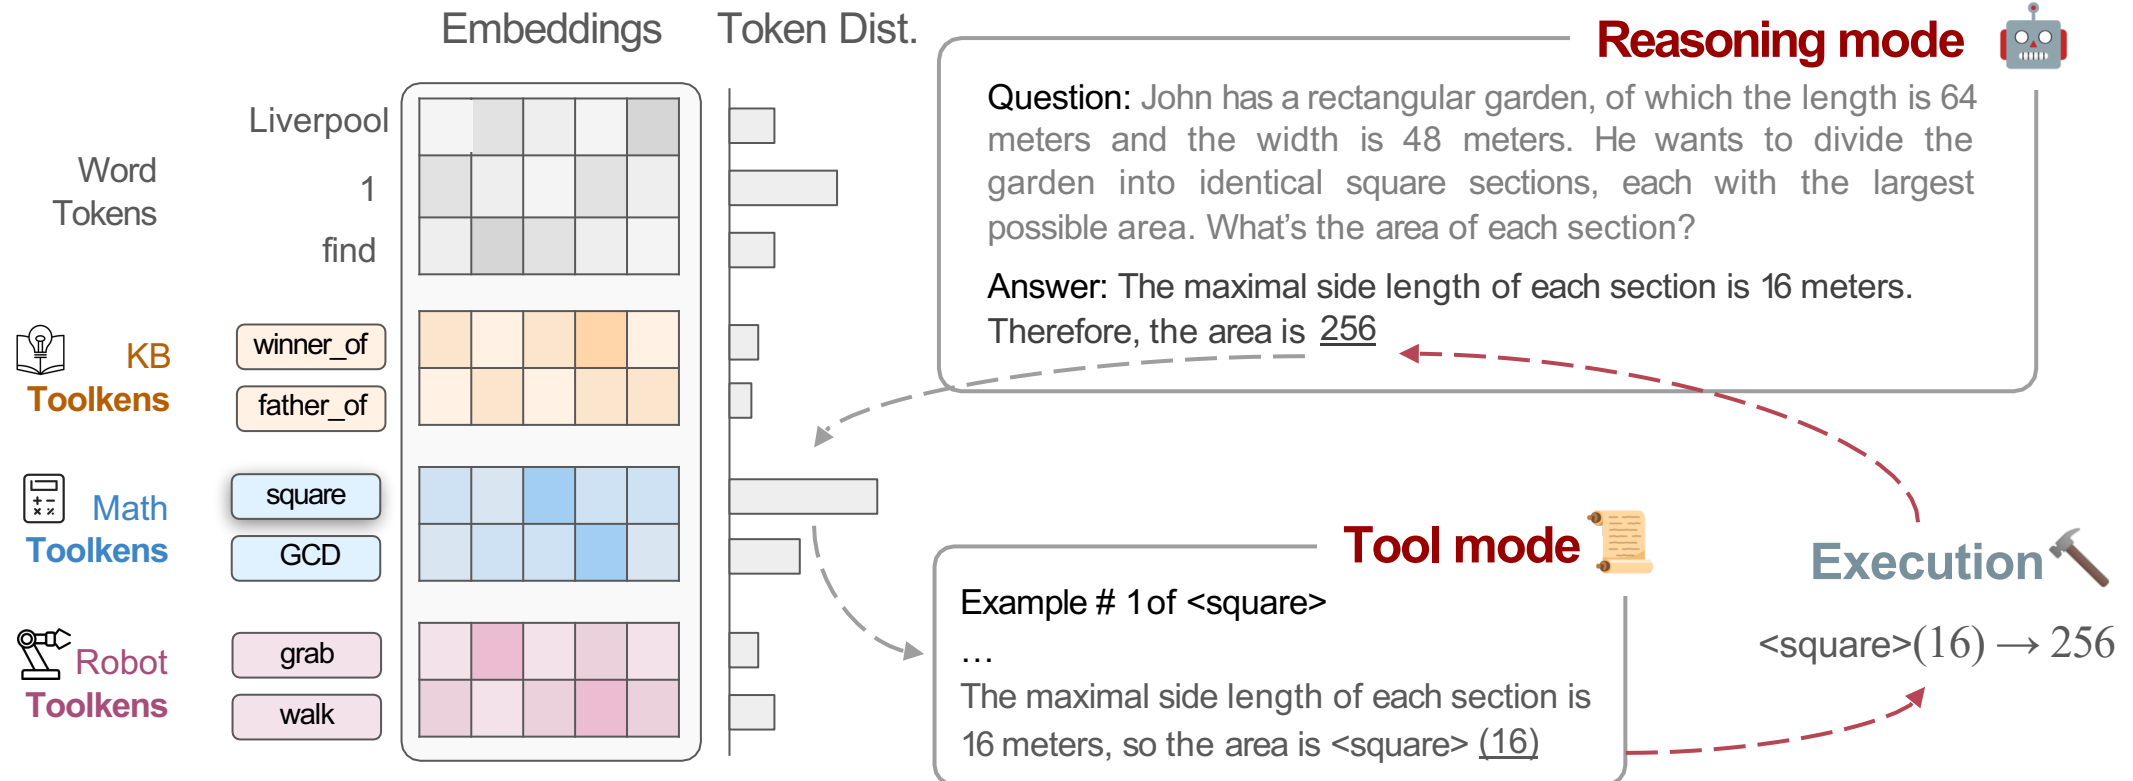
Generating arguments with **in-context learning**



Teaching LLMs to Use Tools - Method #3: Toolken

Step 3: Execute the tool call and return the result

Finally, the tool call is **executed** and the result is **sent back** to the reasoning mode



Teaching LLMs to Use Tools - Method #3: Toolken

Example - Math Reasoning

Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer:

LLaMA-13B/33B



Math tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Math Reasoning

Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is

LLaMA-13B/33B



Math tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Math Reasoning

Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is `GCD` (64, 48)

LLaMA-13B/33B



Math tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Math Reasoning

Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is 16

LLaMA-13B/33B



Math tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Math Reasoning

Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is 16 meters. Therefore, the area is

LLaMA-13B/33B



Math tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Math Reasoning

Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is 16 meters. Therefore, the area is (16)

LLaMA-13B/33B



Math tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Math Reasoning

Question: John has a rectangular garden, of which the length is 64 meters and the width is 48 meters. He wants to divide the garden into identical square sections, each with the largest possible area. What's the area of each section?

Answer: The maximal side length of each section is 16 meters. Therefore, the area is 256 square meters

LLaMA-13B/33B



Math tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Knowledge-based QA

Question: Which team is the winner of 2005-06 FA CUP?

Answer:

LLaMA-13B/33B



KB tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Knowledge-based QA

Question: Which team is the winner of 2005-06 FA CUP?

Answer: The winner is

LLaMA-13B/33B



KB tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Knowledge-based QA

Question: Which team is the winner of 2005-06 FA CUP?

Answer: The winner is `winner_of` (2005-06 FA CUP)

LLaMA-13B/33B



KB tools

Teaching LLMs to Use Tools - Method #3: Toolken

Example - Knowledge-based QA

Question: Which team is the winner of 2005-06 FA CUP?

Answer: The winner is **Liverpool**

LLaMA-13B/33B



KB tools

Agent models with external augmentations

- External augmentations for added capabilities:
 - **Tools**: telescope, vehicles, ...
 - **Data about a skill**: demonstration videos of climbing a snowy mountain
 - **Knowledge bases**: domain knowledge
- Agent automatically chooses appropriate augmentations for a given task
 - How to represent millions of potential augmentations?
 - Learning unified embedding of tools, data, knowledge [Hao et al., 2023]
- Another dimension rarely considered so far: constraint by **budget**
 - Different augmentations will invoke different costs (financial, time, etc.)
 - Need to strike the optimal balance between task performance vs costs

Key Takeaways

- Richer learning mechanisms
 - Learning with Embodied Experiences
 - Social Learning
- Multi-modal capabilities
 - Multi-modal LMs, video generation models
- Latent-space reasoning
 - How to learn a good multi-level latent space
- Agent models with external augmentations (e.g., tools)
 - Unified embedding, budget for augmentations

Discussion

- **No Free Lunch (NFL) theorem** (suggested reading of Lecture#10):
 - No single machine learning algorithm is universally the best-performing algorithm for all problems
- Do generalist models (LLMs) violate this theorem?
- Does "the Bitter Lesson" contradict with this theorem?
 - (suggested reading of Lecture#6)

Questions?