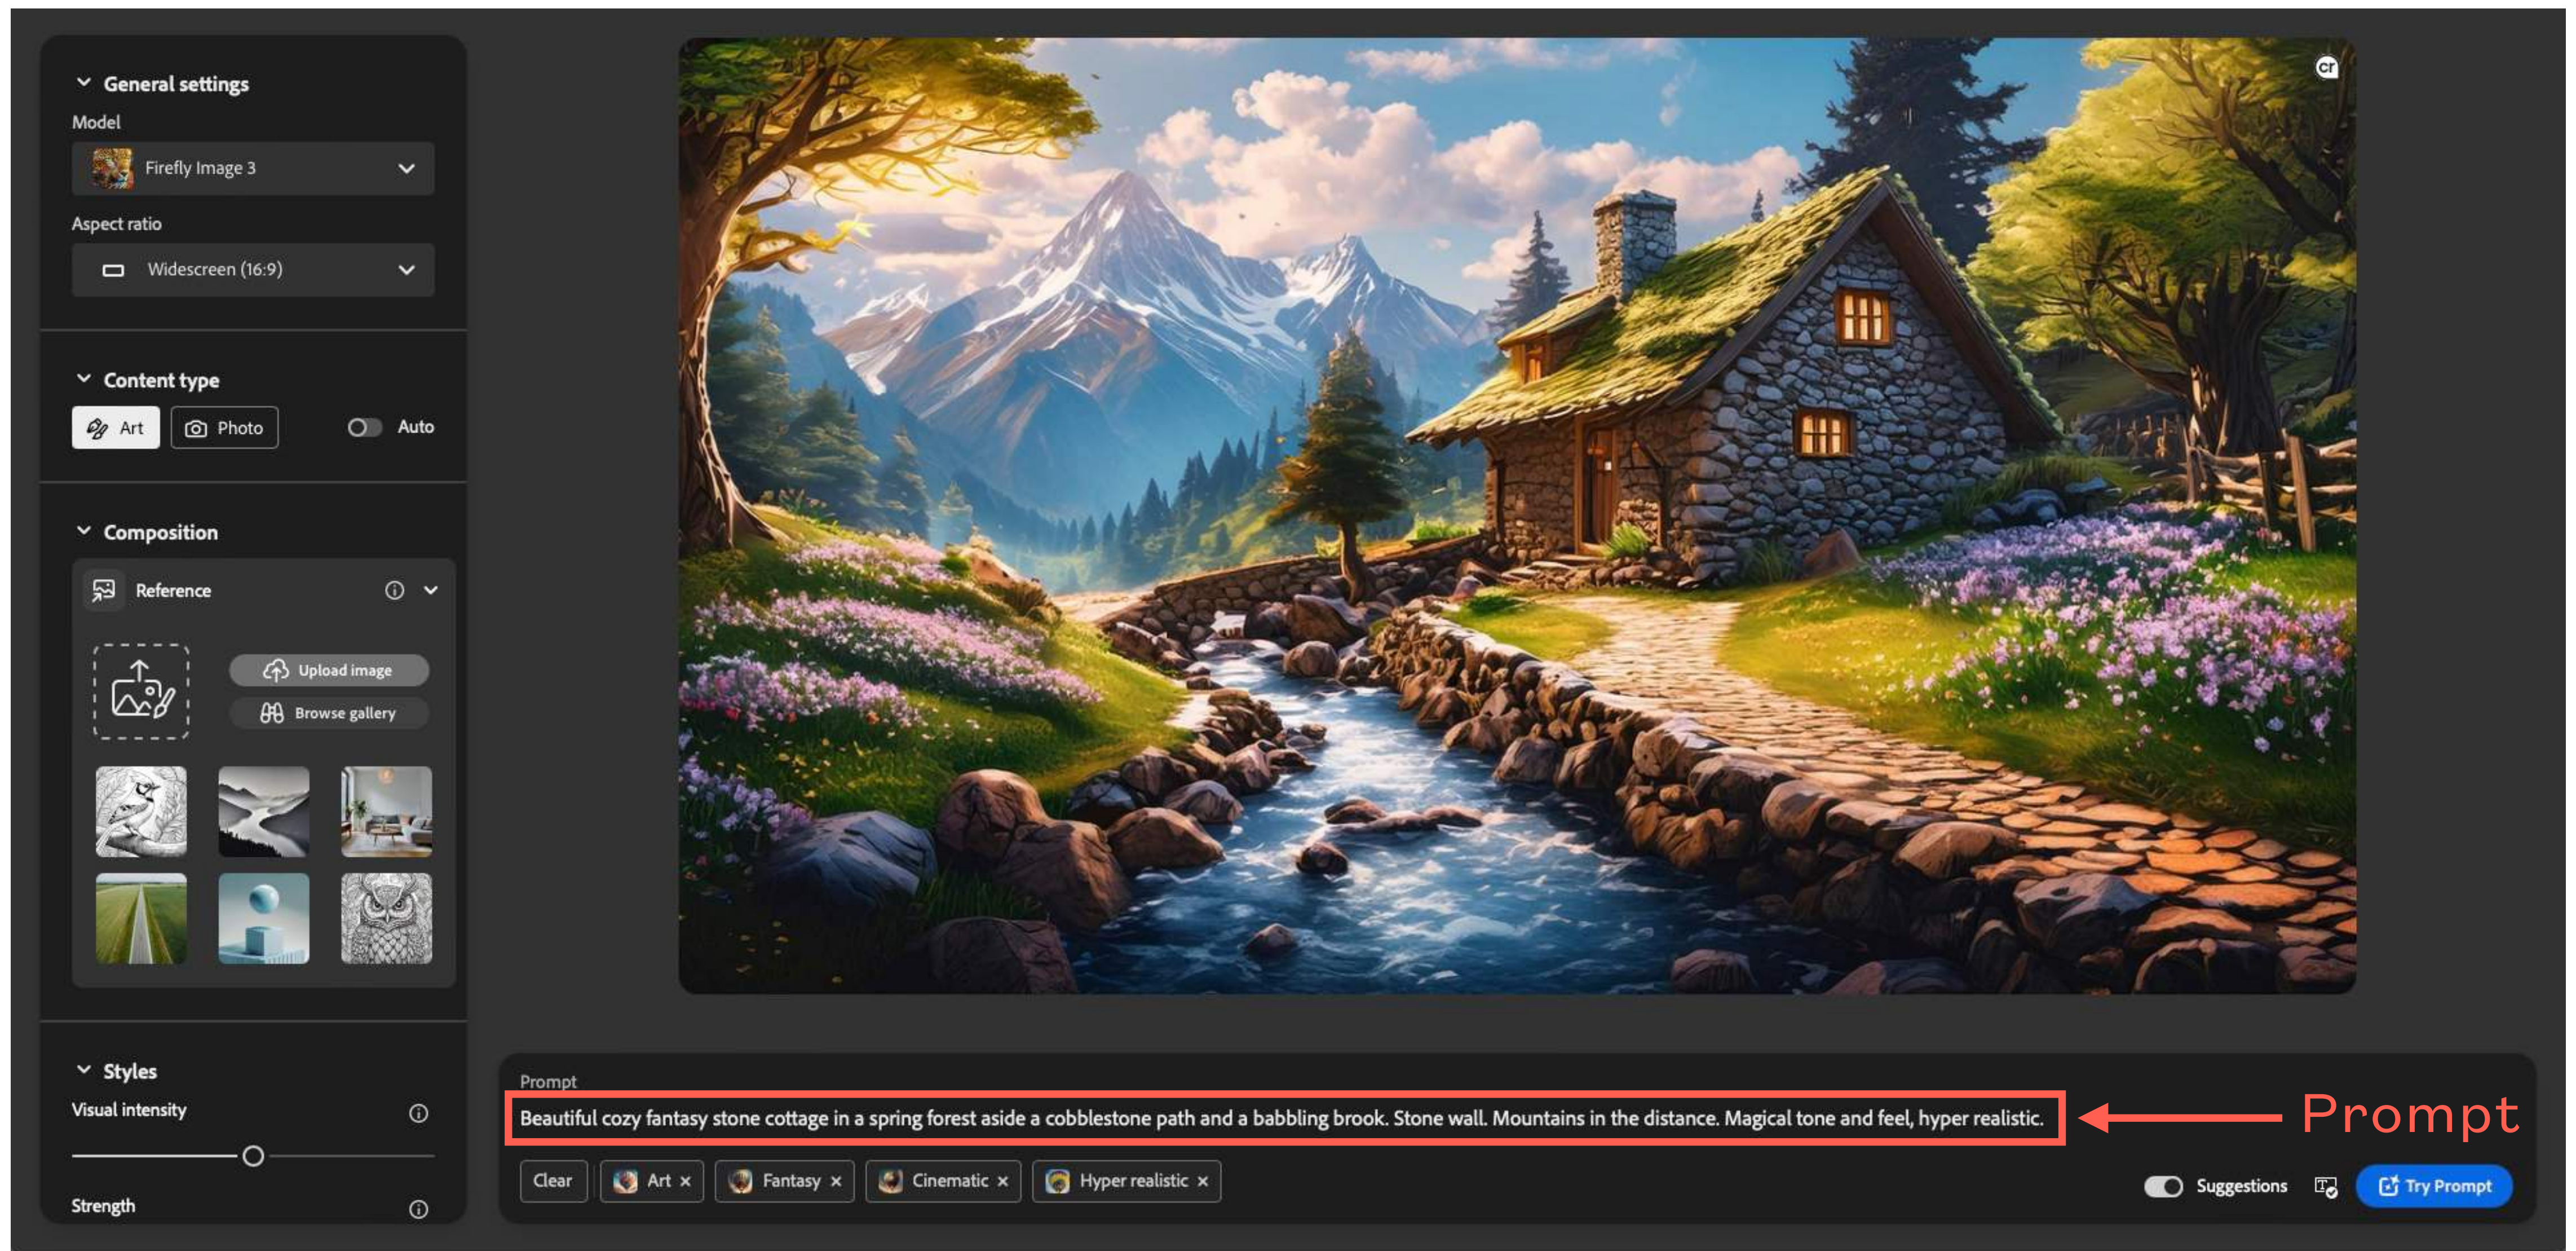




A Brief Introduction to

Prompt Engineering

BY PETE - PITTAWAT TAVEEKITWORACHAI



Agenda

- LLM
- Prompt and Prompt Engineering
- Future Outlook

(Large) Language Model

What is Language Model (LM)?

- A model that takes a words sequence as an input and predict a **next word** as an output (**text completion**)

- Example

Computer consists of two main parts, hardware and _____

What is Language Model (LM)?

- A model that takes a words sequence as an input and predict a **next word** as an output (**text completion**)

- Example

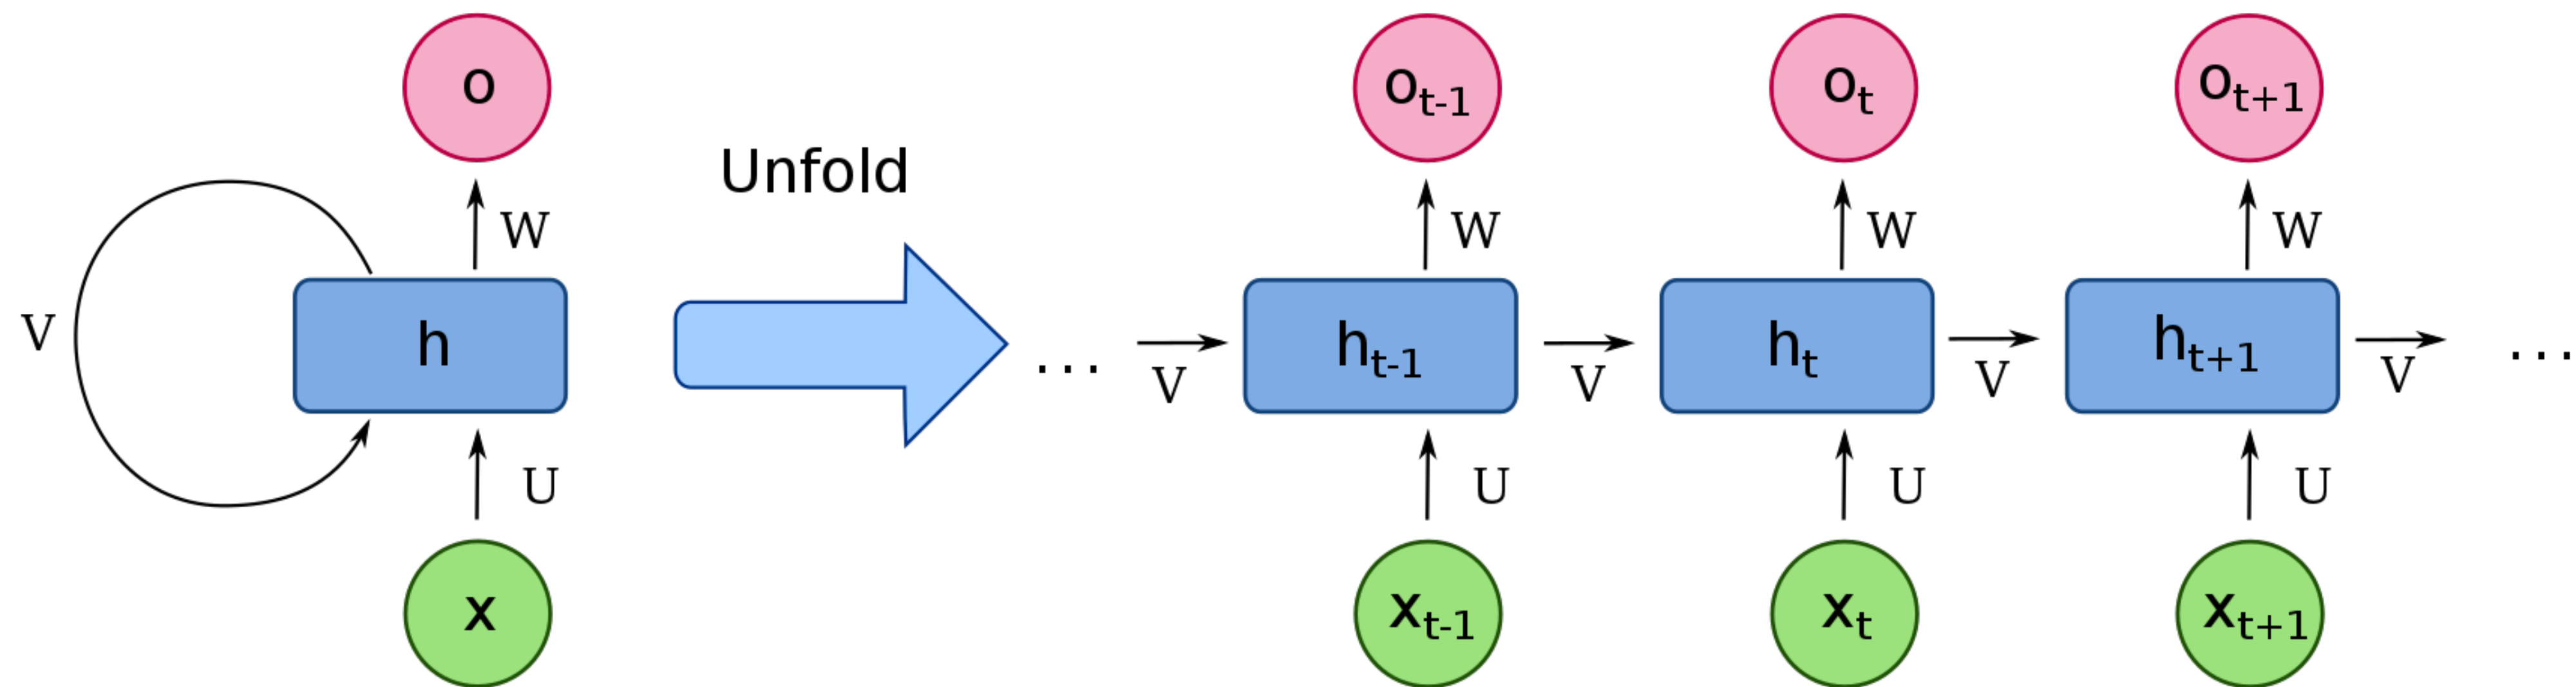
Computer consists of two main parts, hardware and **software**.

Why LM?

- Pre-text task (self-supervised)
- Fine-tune it for downstream tasks, e.g., sentiment analysis, toxicity classification, summarization

How to Train Your Dragon LM?

- In the past, we use **Recurrent Neural Network (RNN)**



How to Train Your Dragon LM?

- In the past, we use **Recurrent Neural Network (RNN)**
- One huge drawback: It is **sequential!**

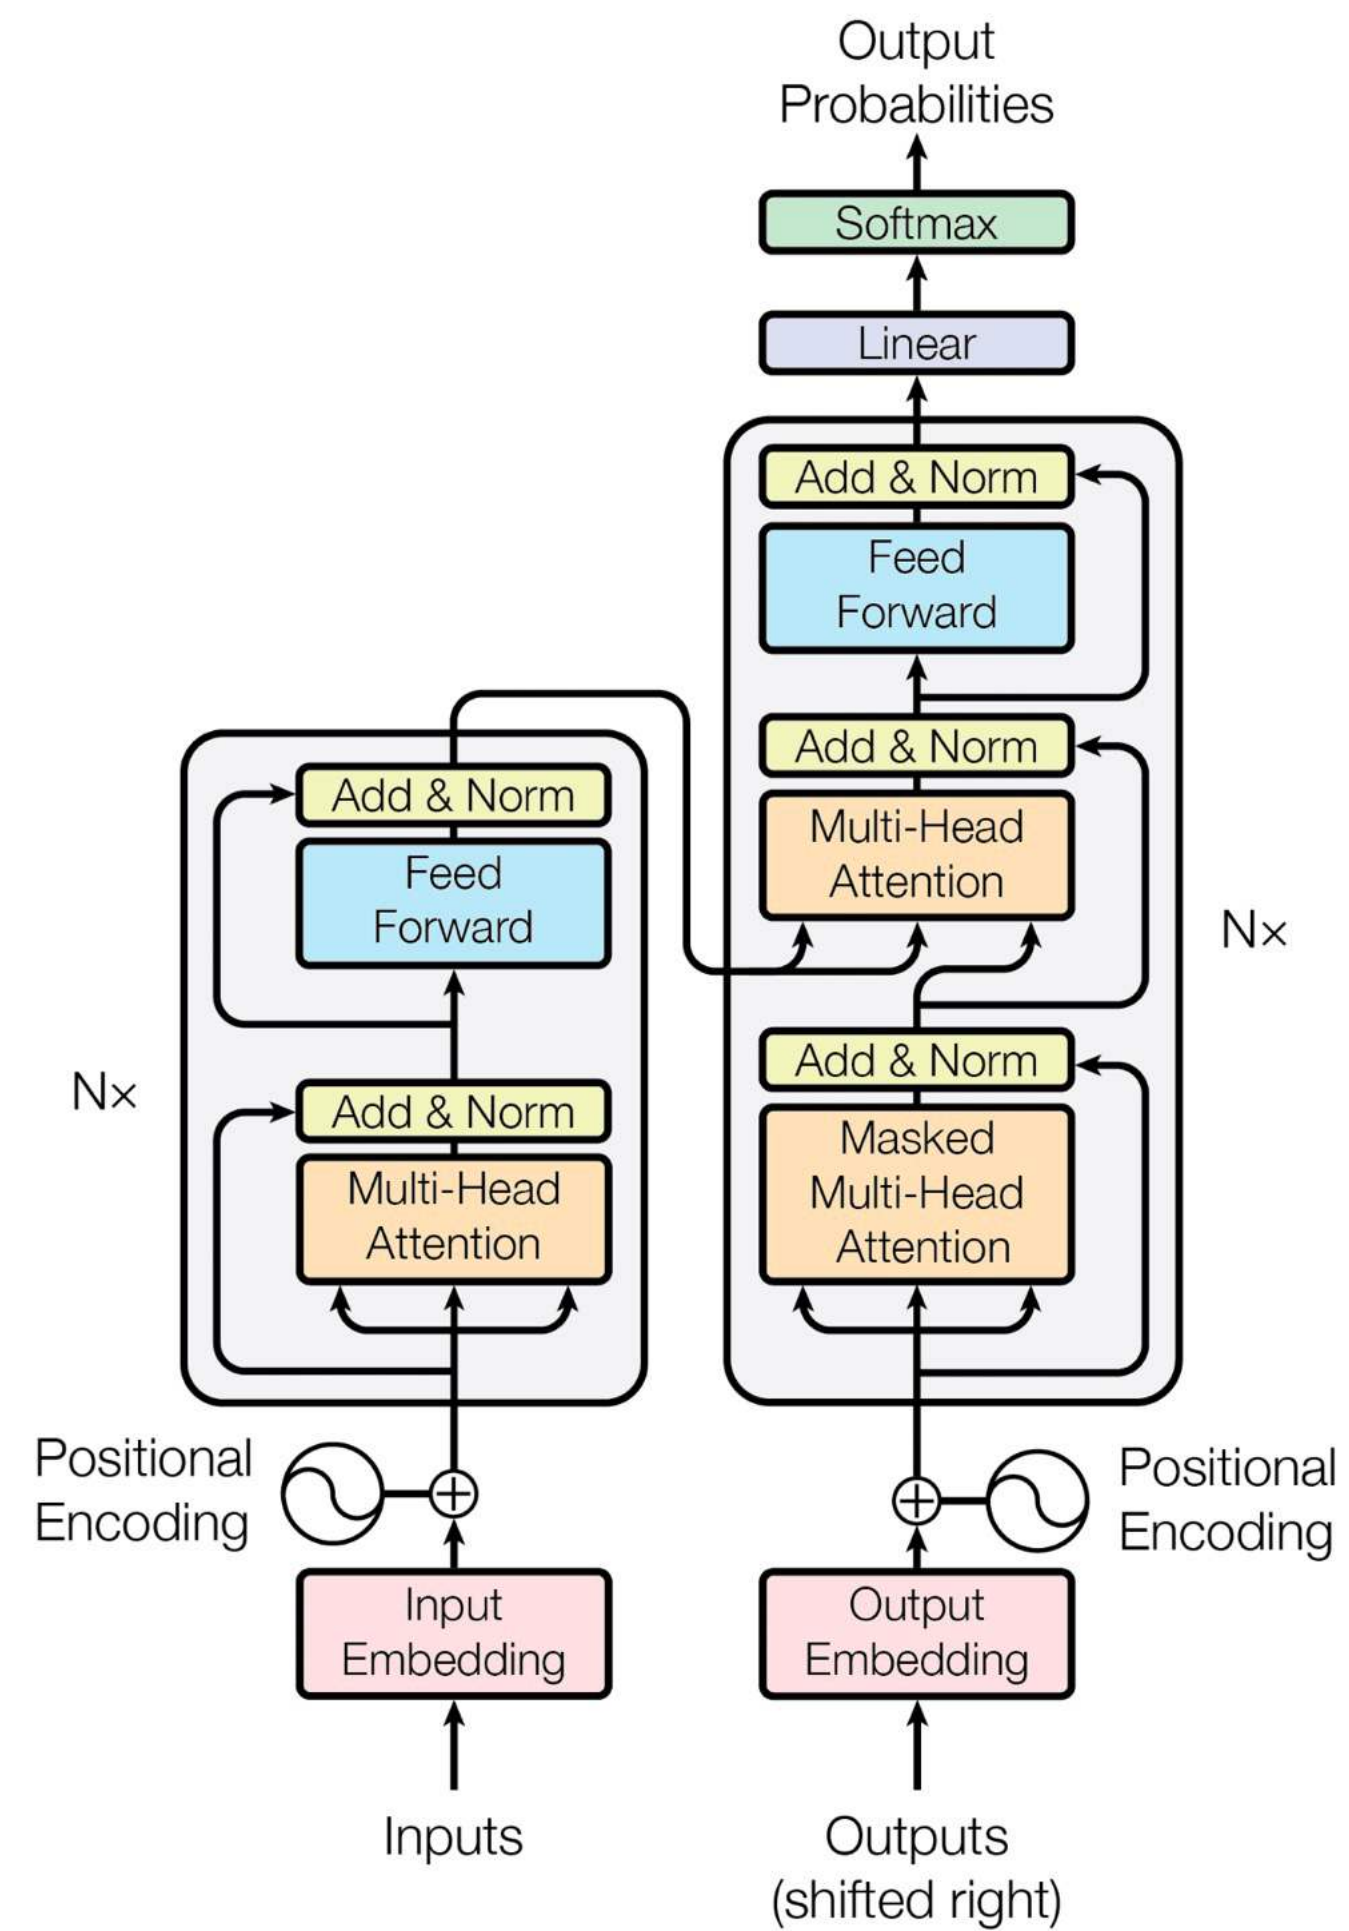
The Rise of Transformers



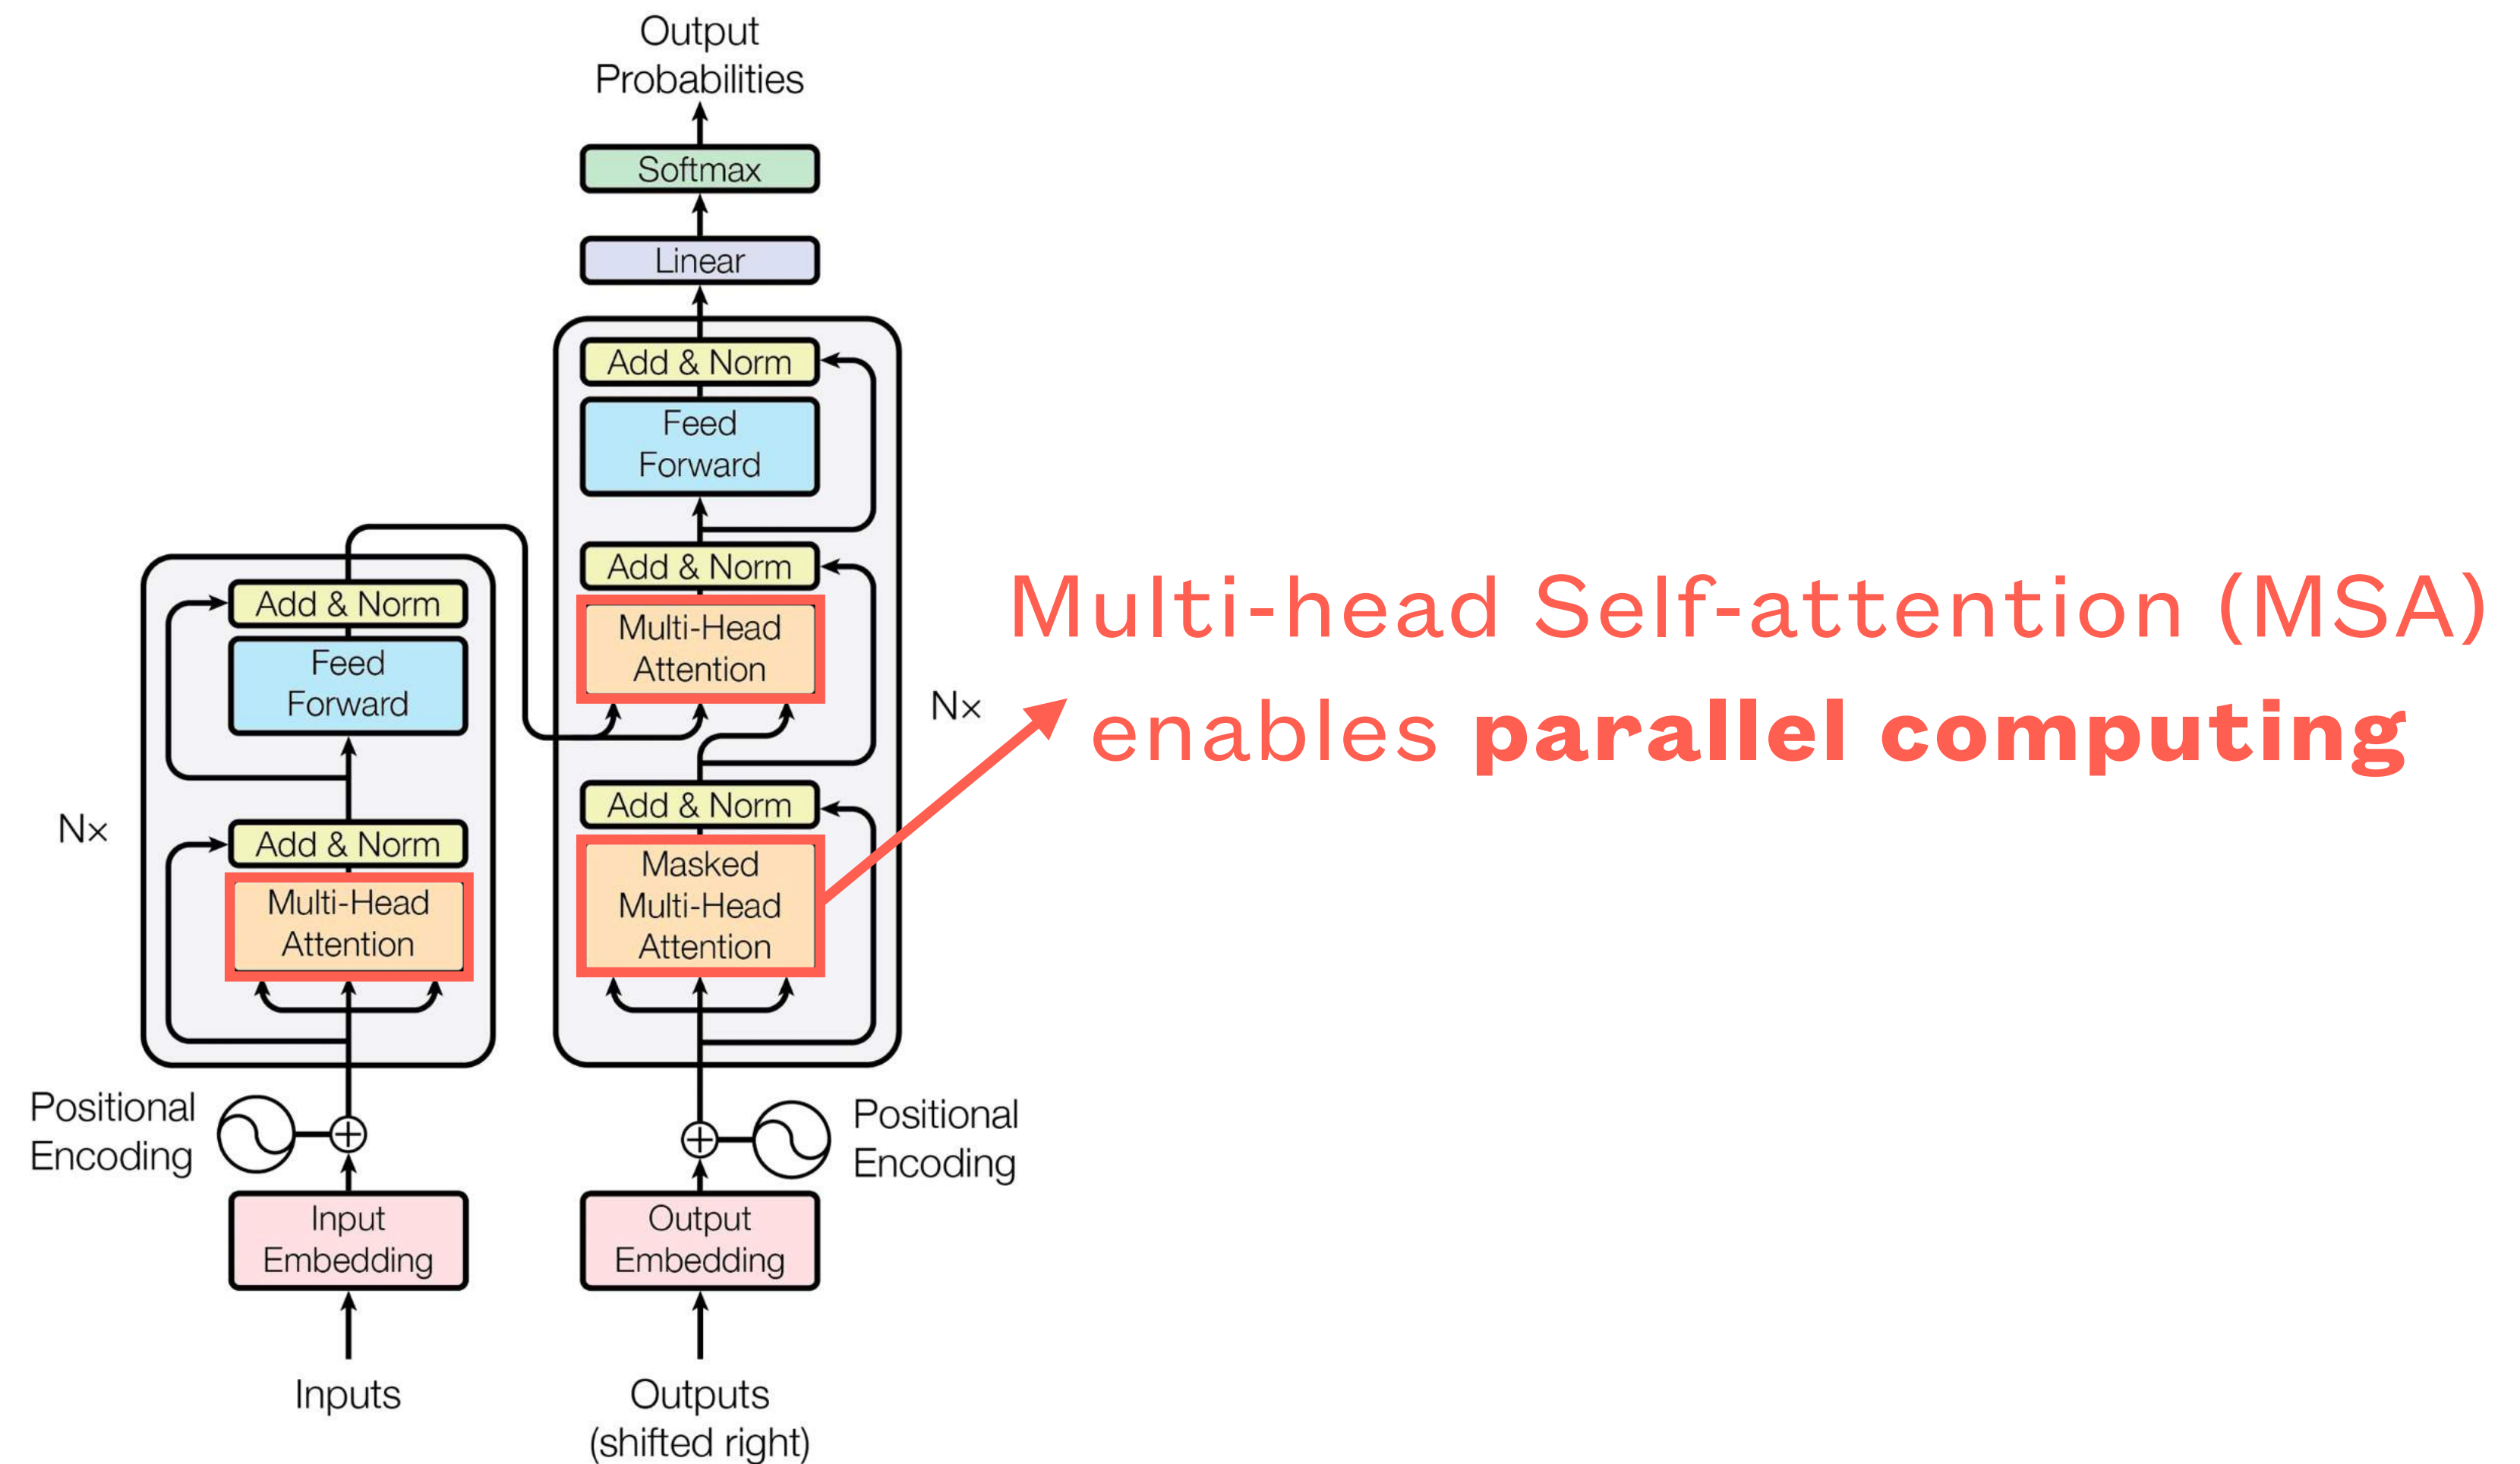
The Rise of Transformers



The Rise of Transformers

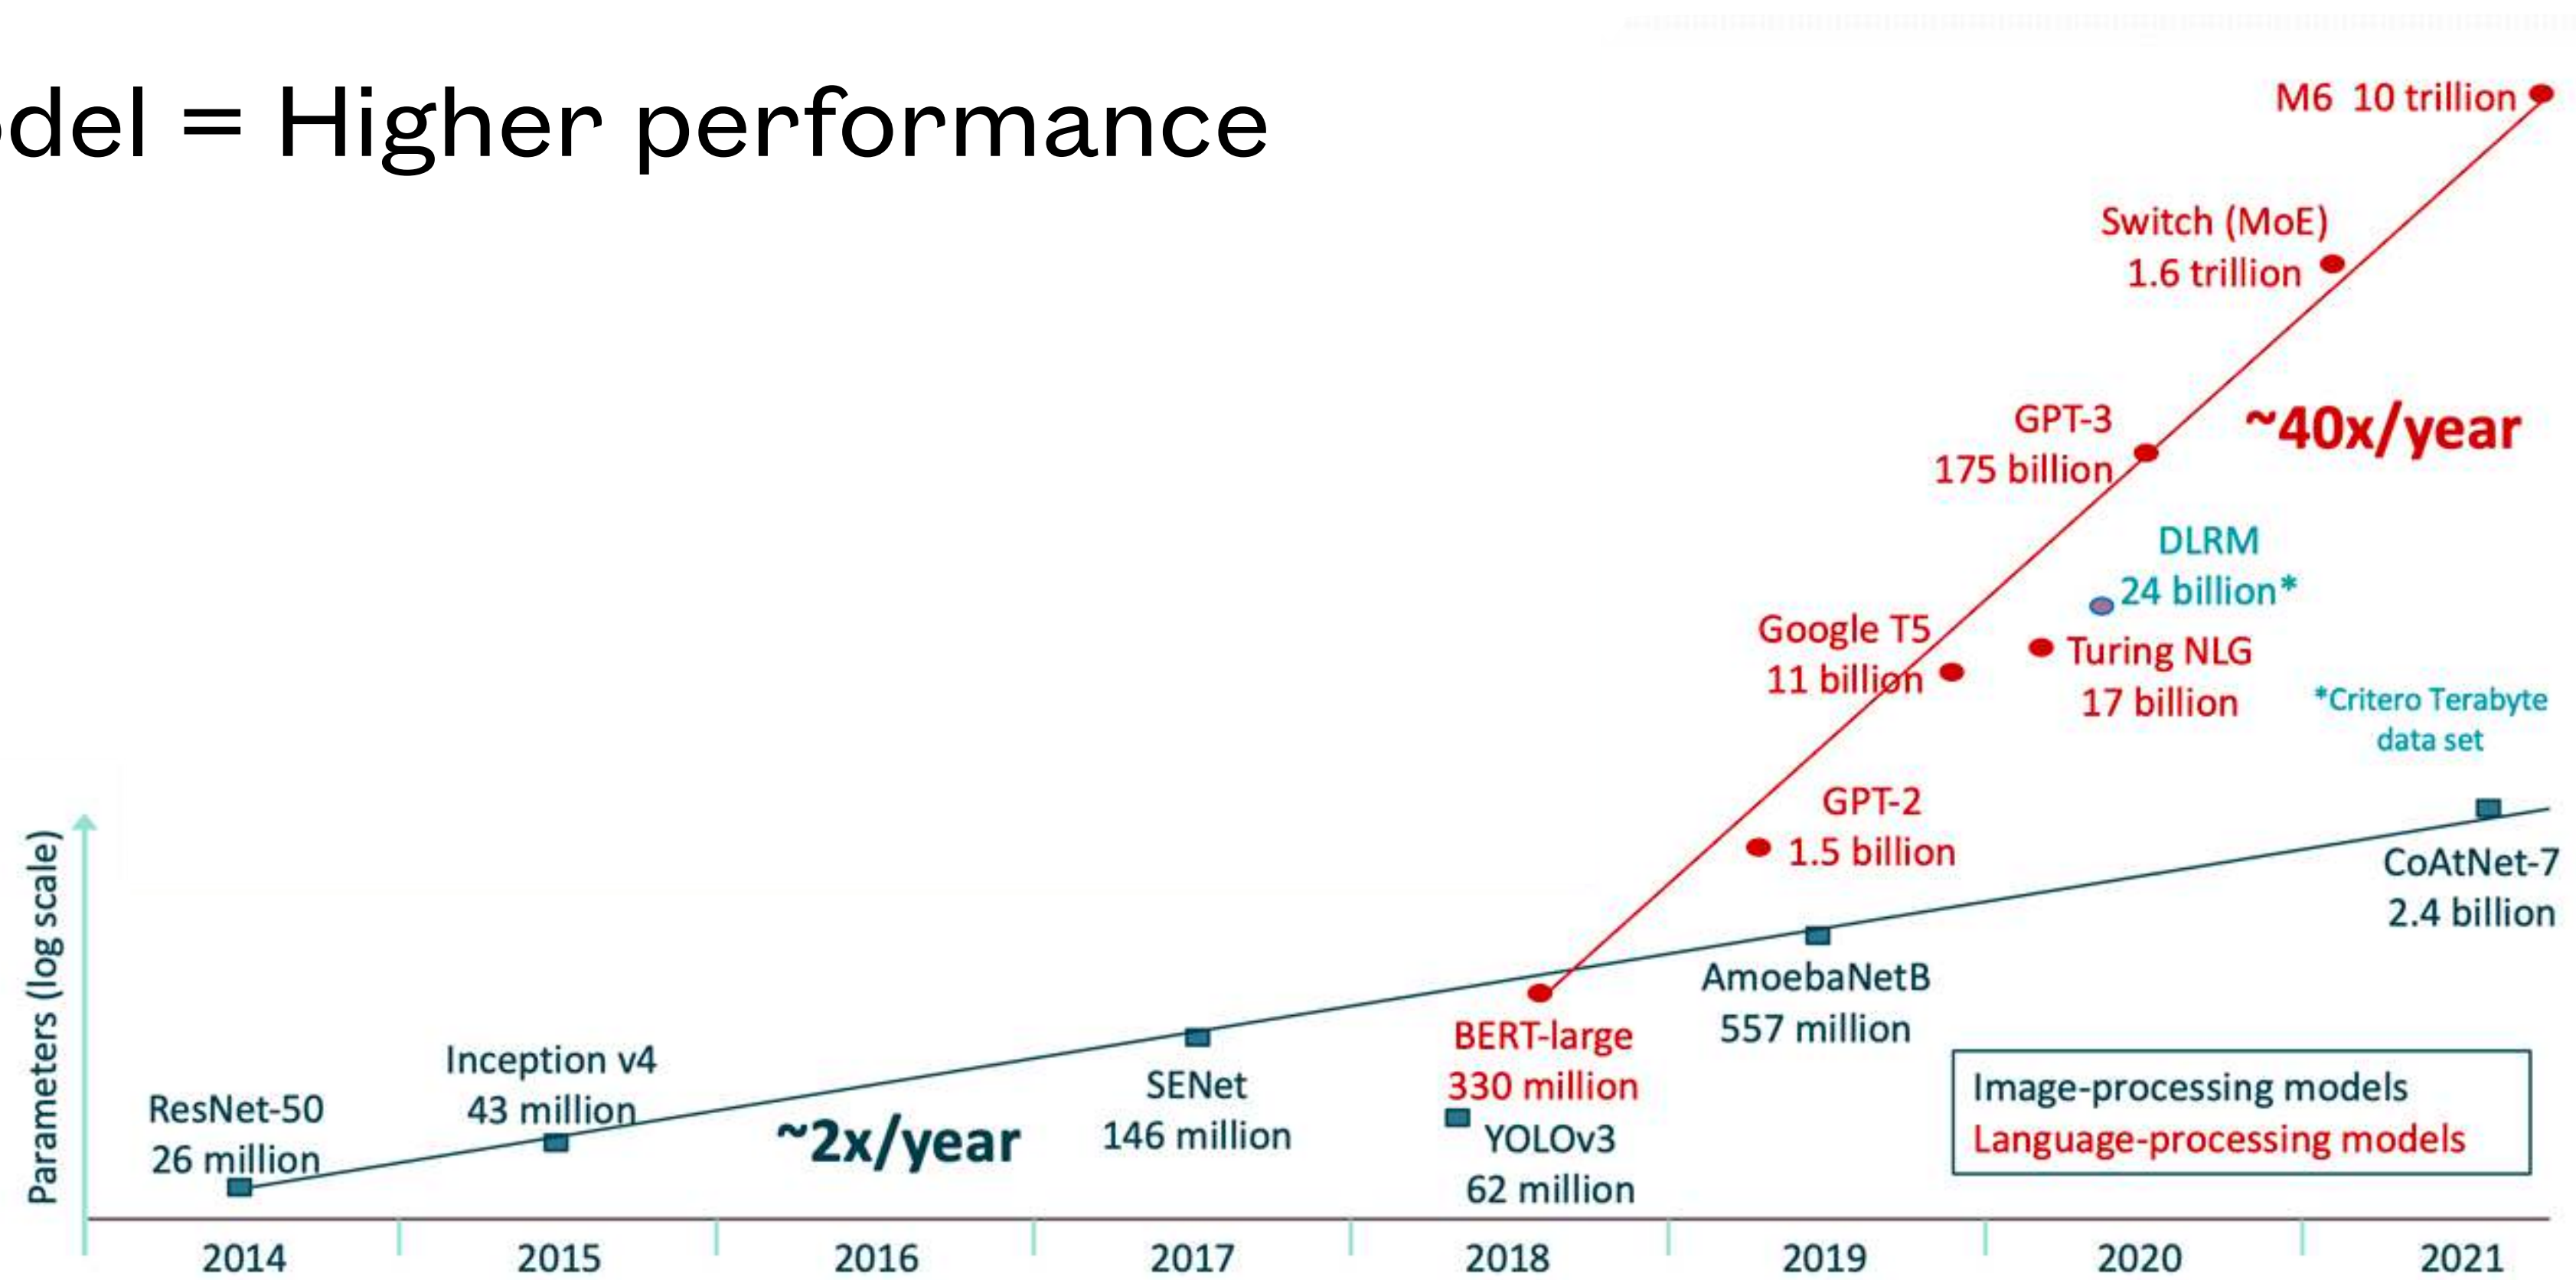


The Rise of Transformers

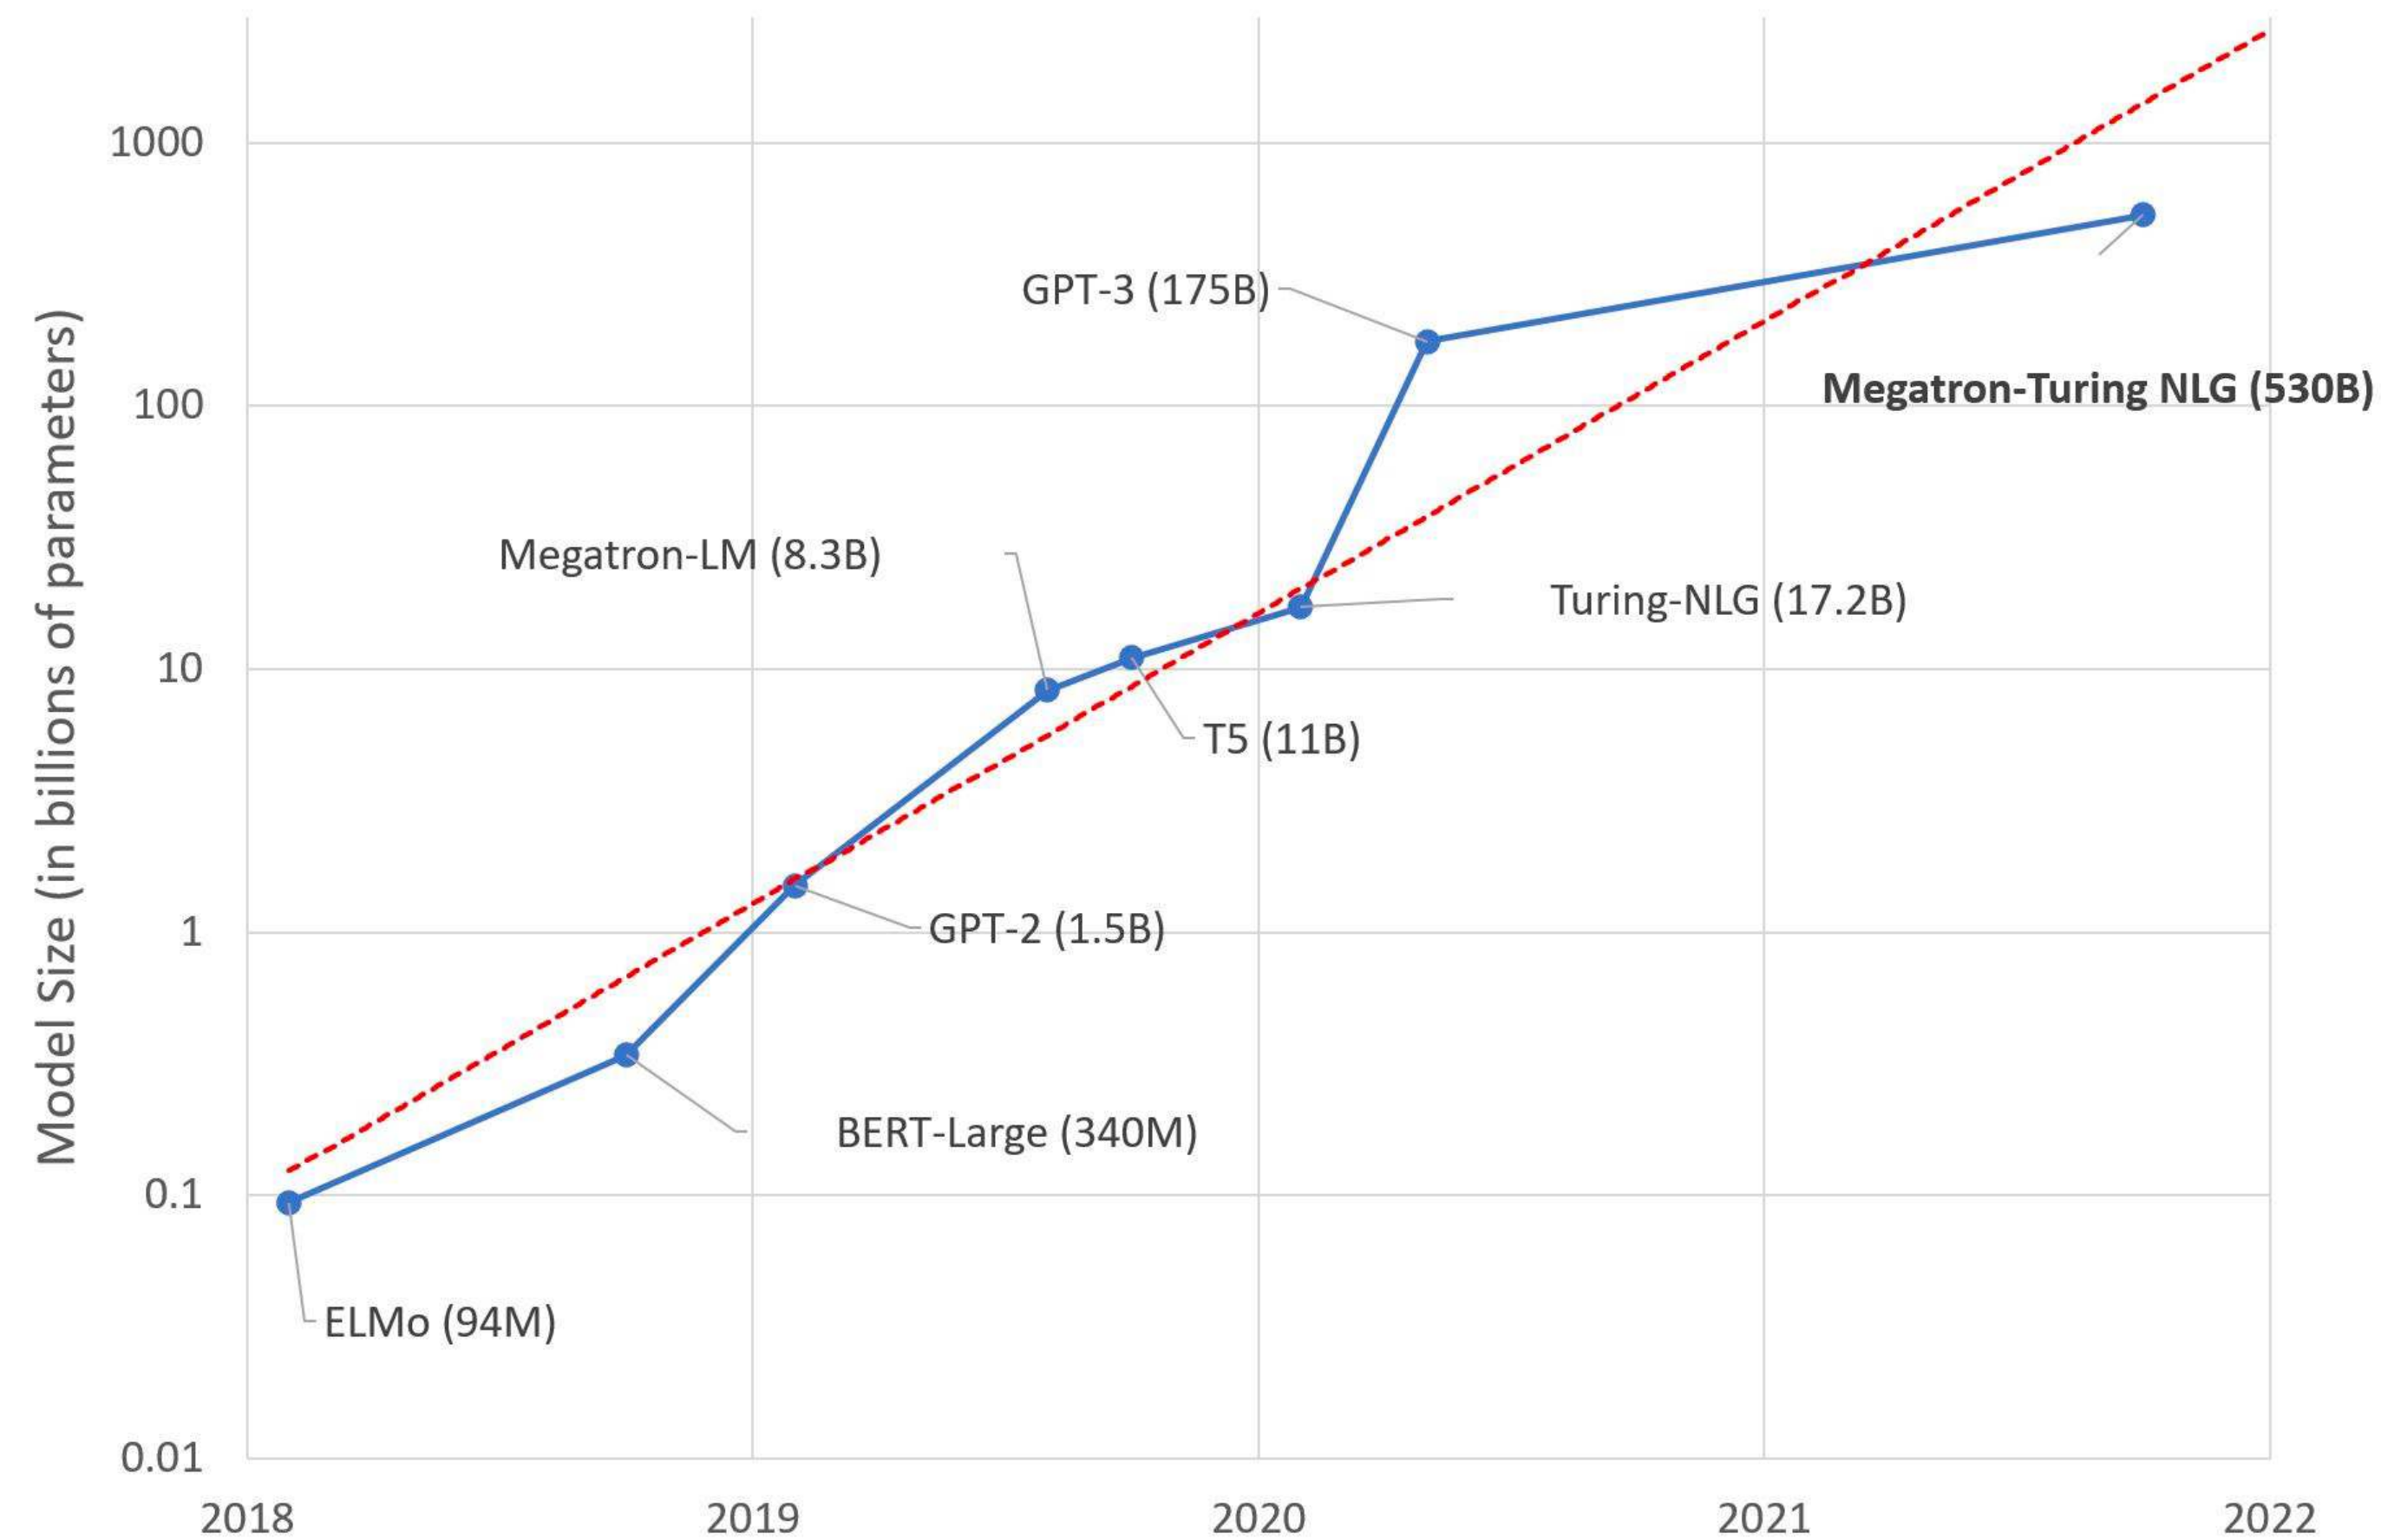


Why We Need Large LM (LLM)?

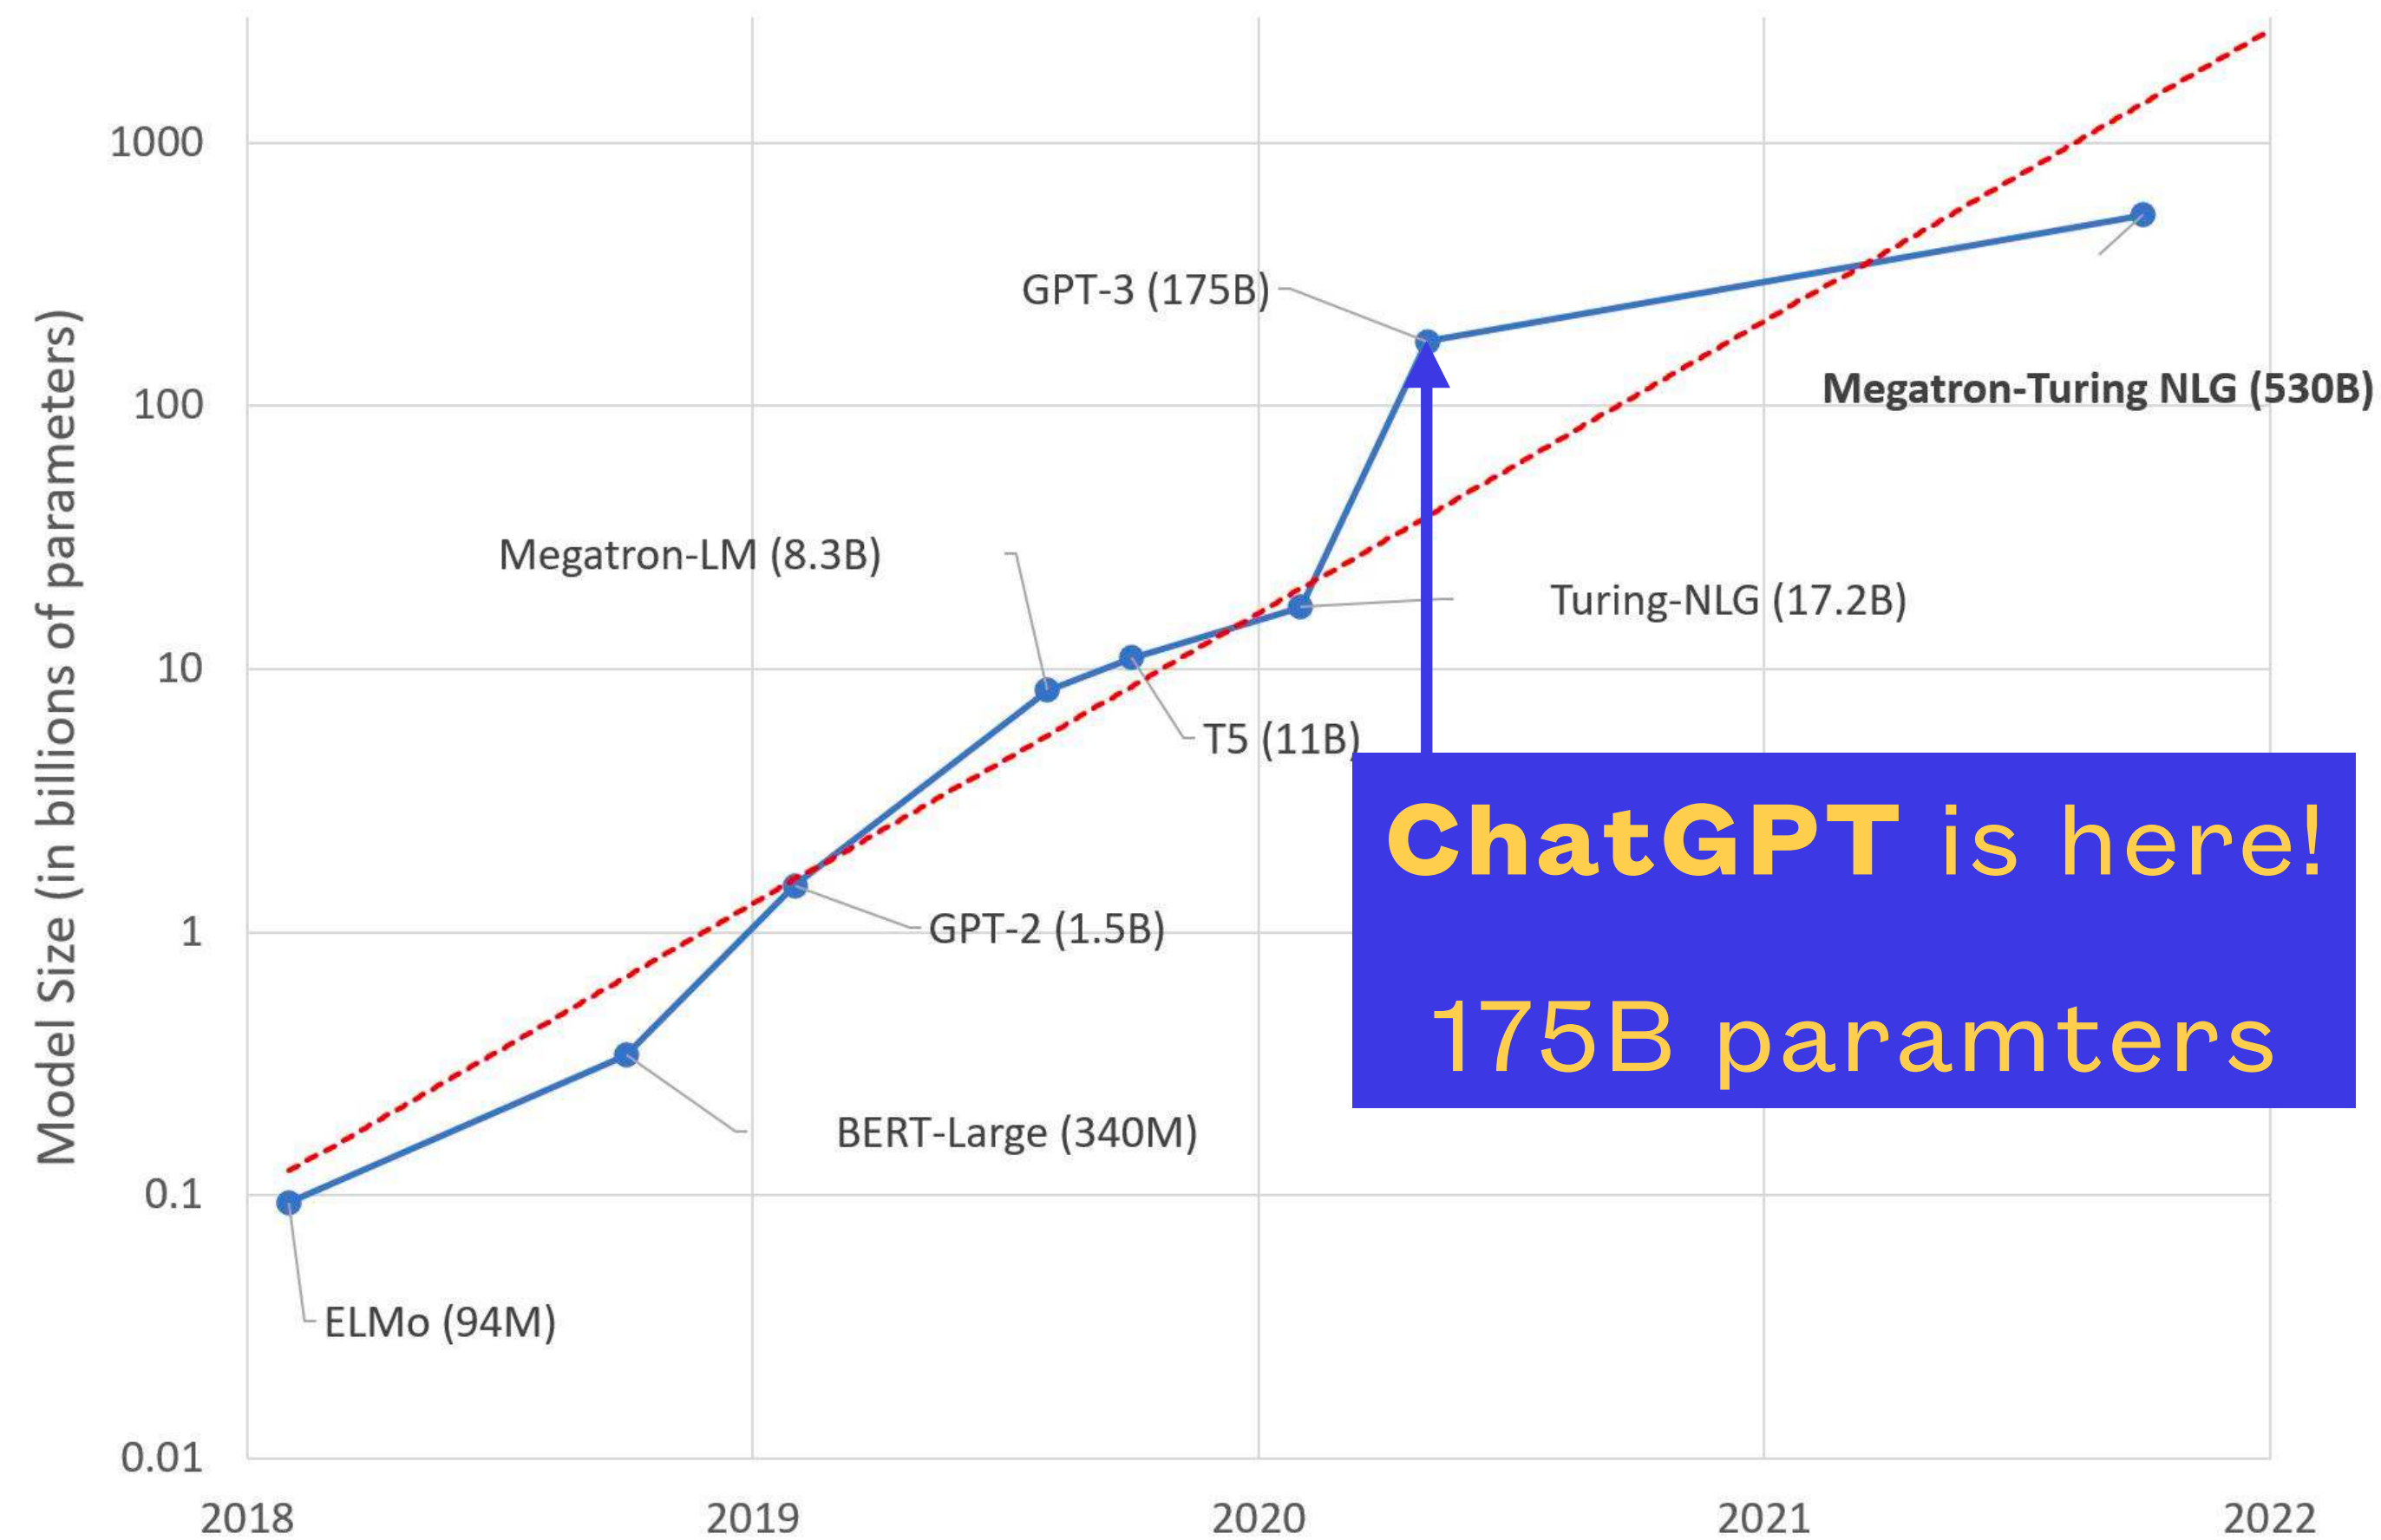
- Larger model = Higher performance



Why We Need Large LM (LLM)?

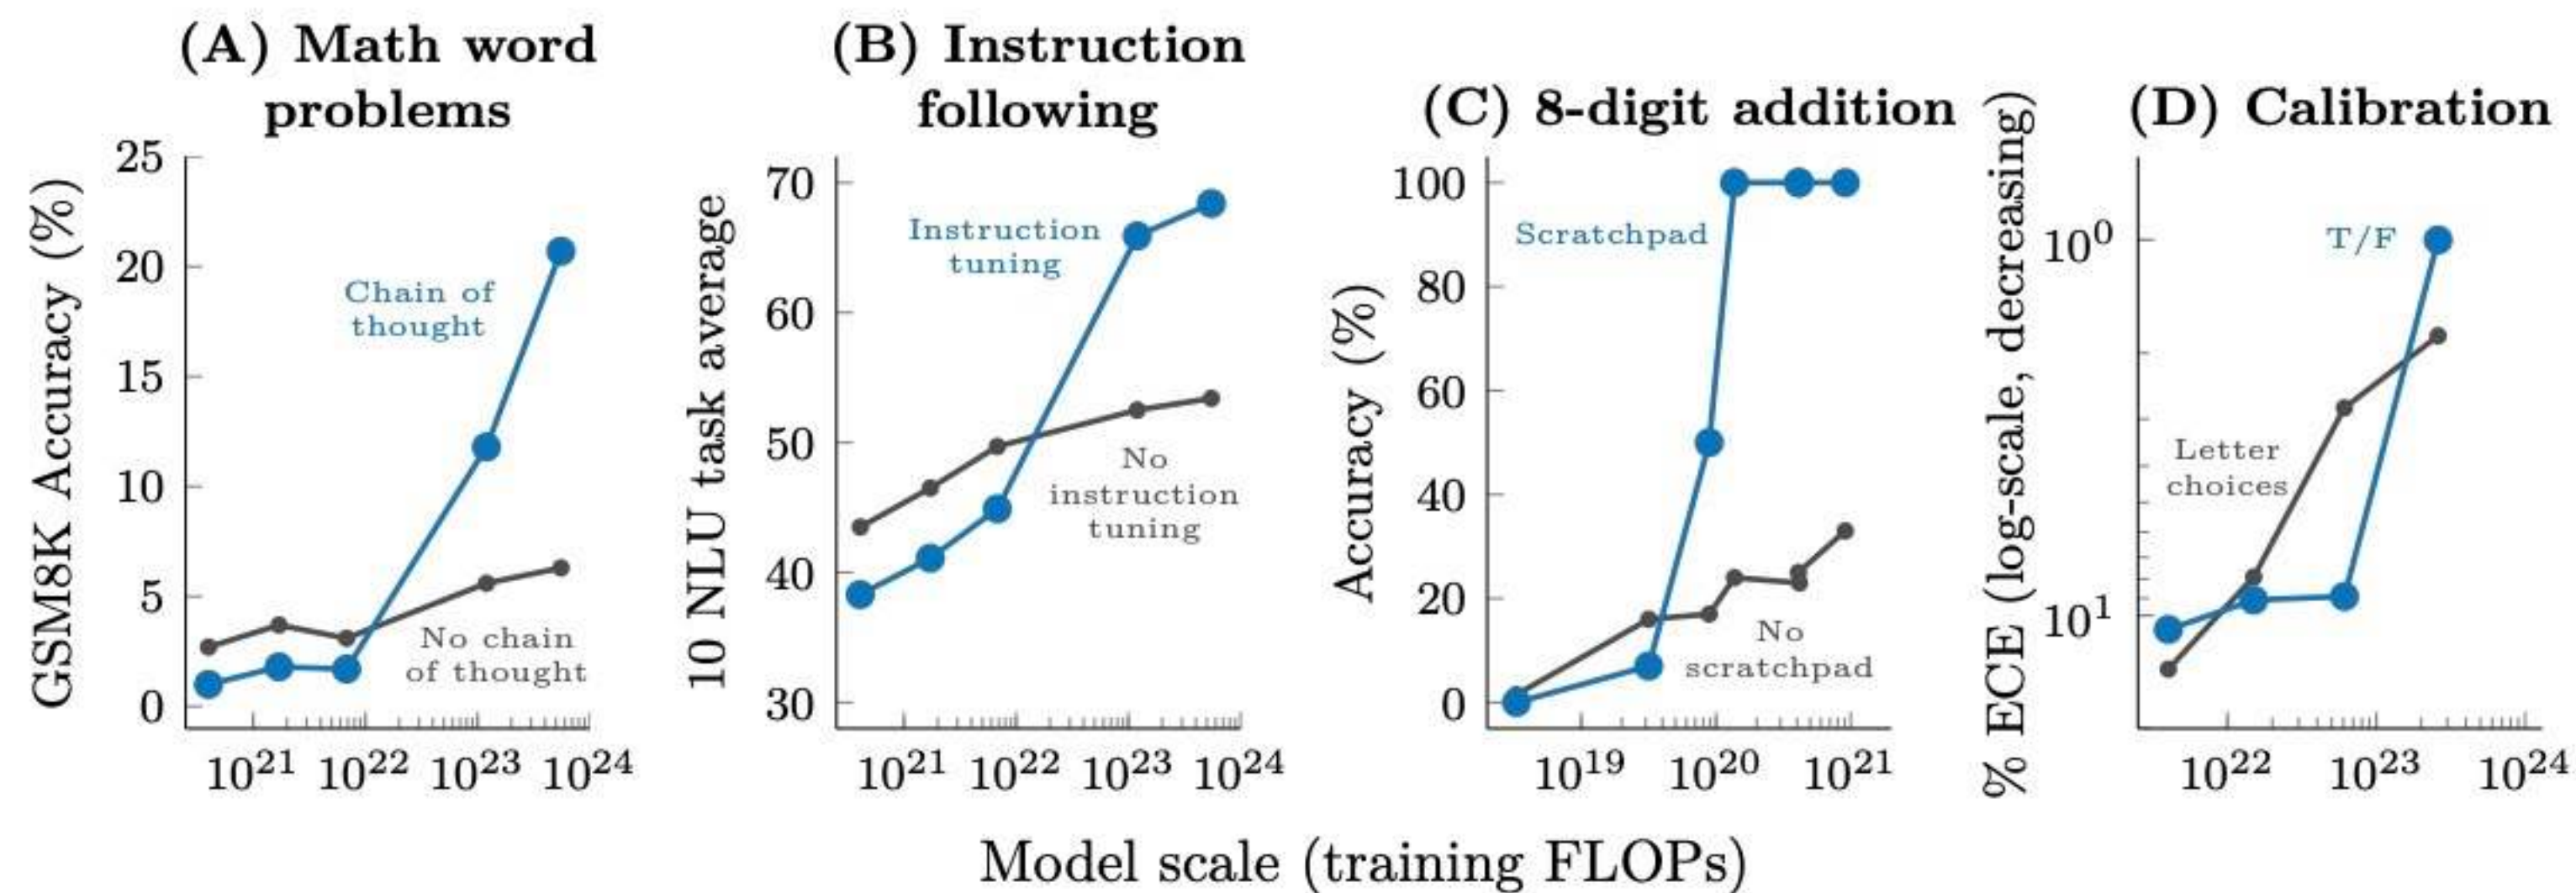


Why We Need Large LM (LLM)?



Emergent Abilities: LLM Exclusivity

- When the model get bigger to a certain point, it gains **emergent abilities**
- Can perform **untrained** tasks with high performance

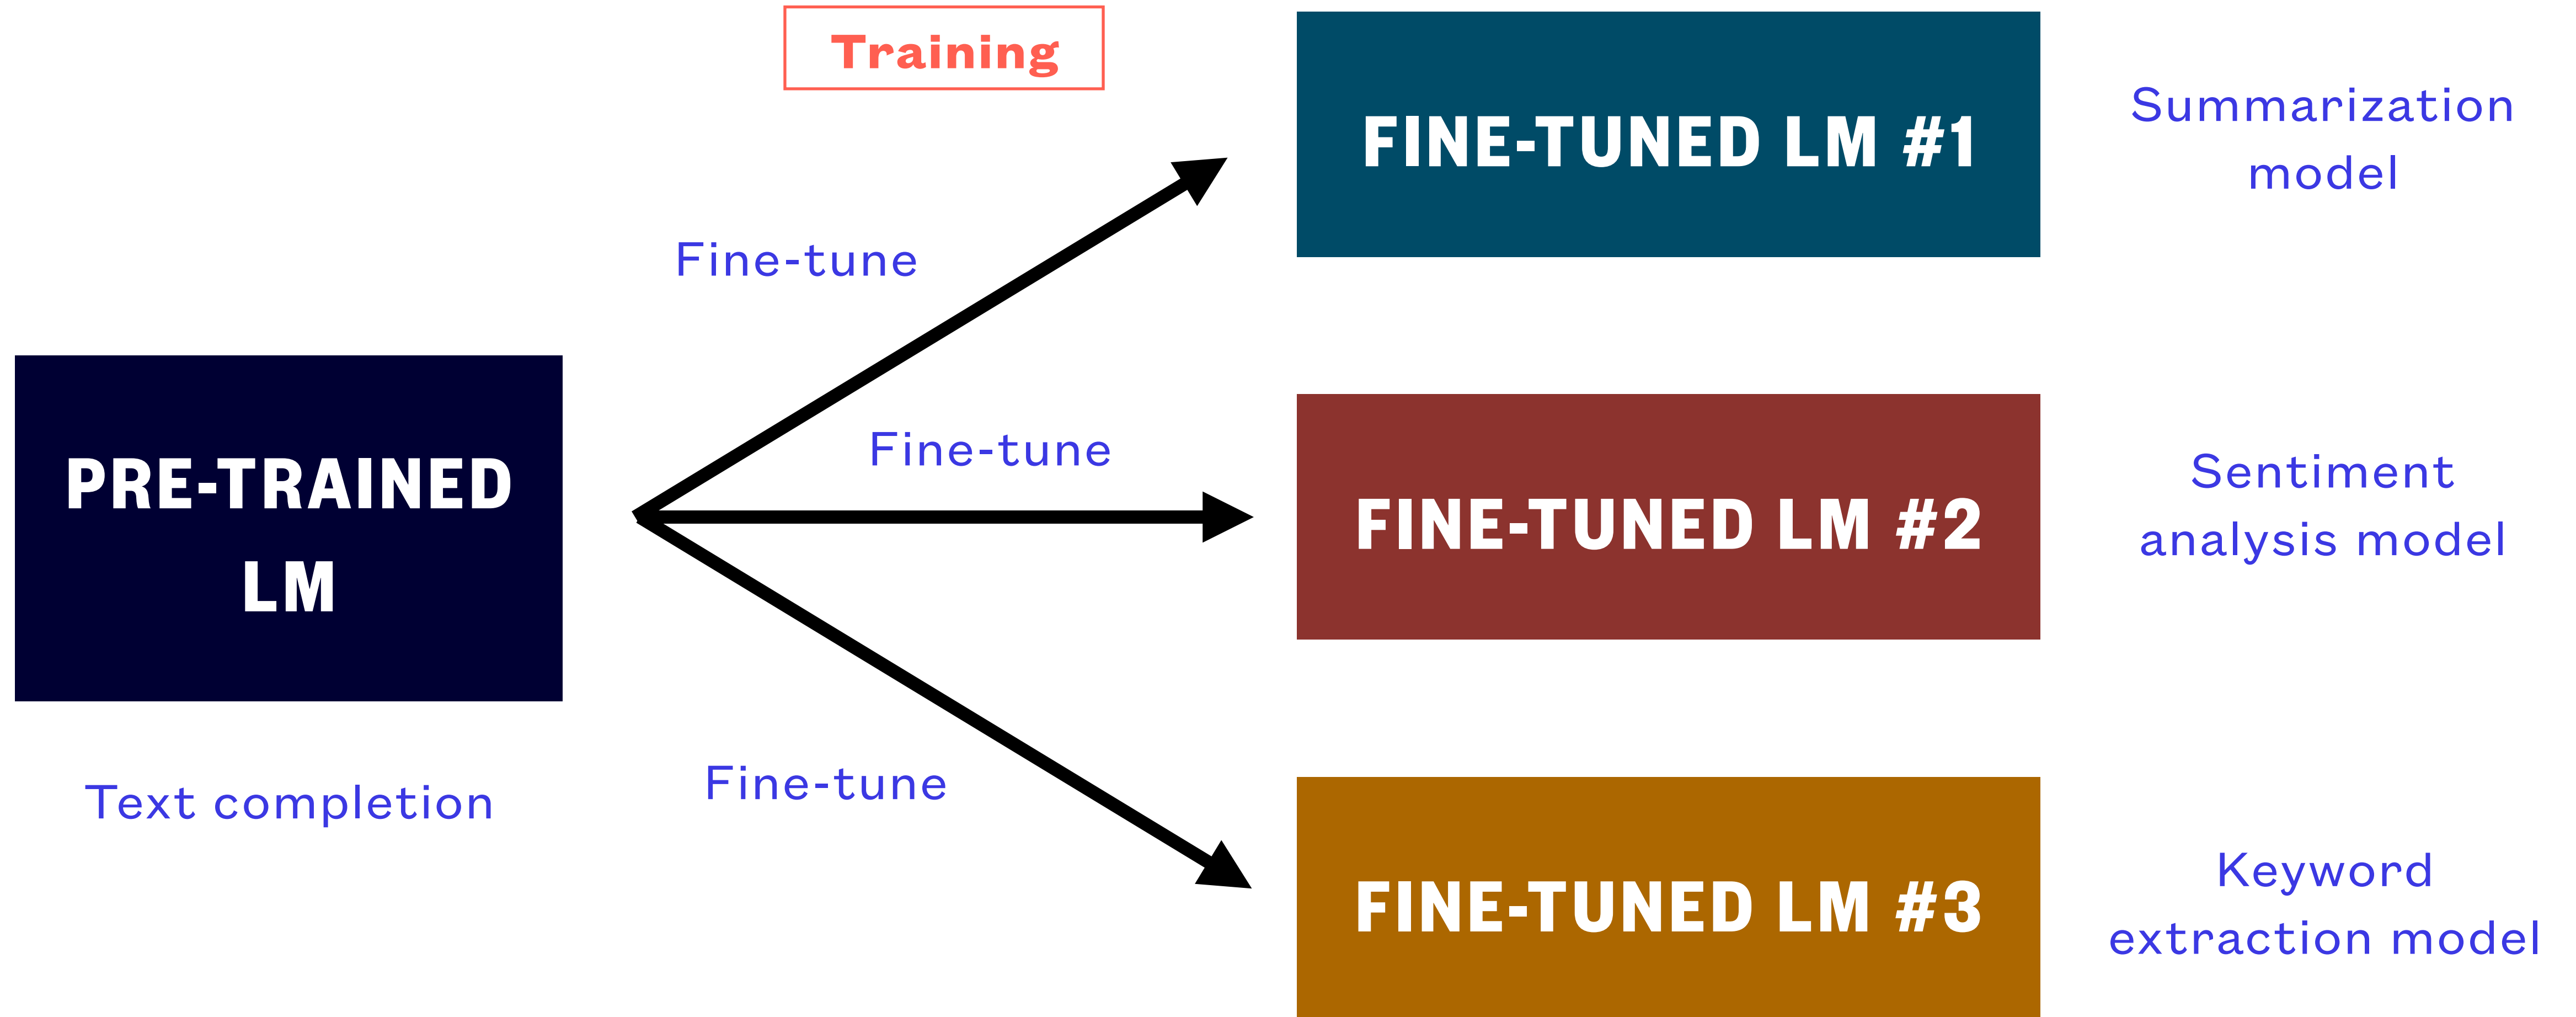


LM vs LLM

**PRE-TRAINED
LM**

Text completion

LM vs LLM



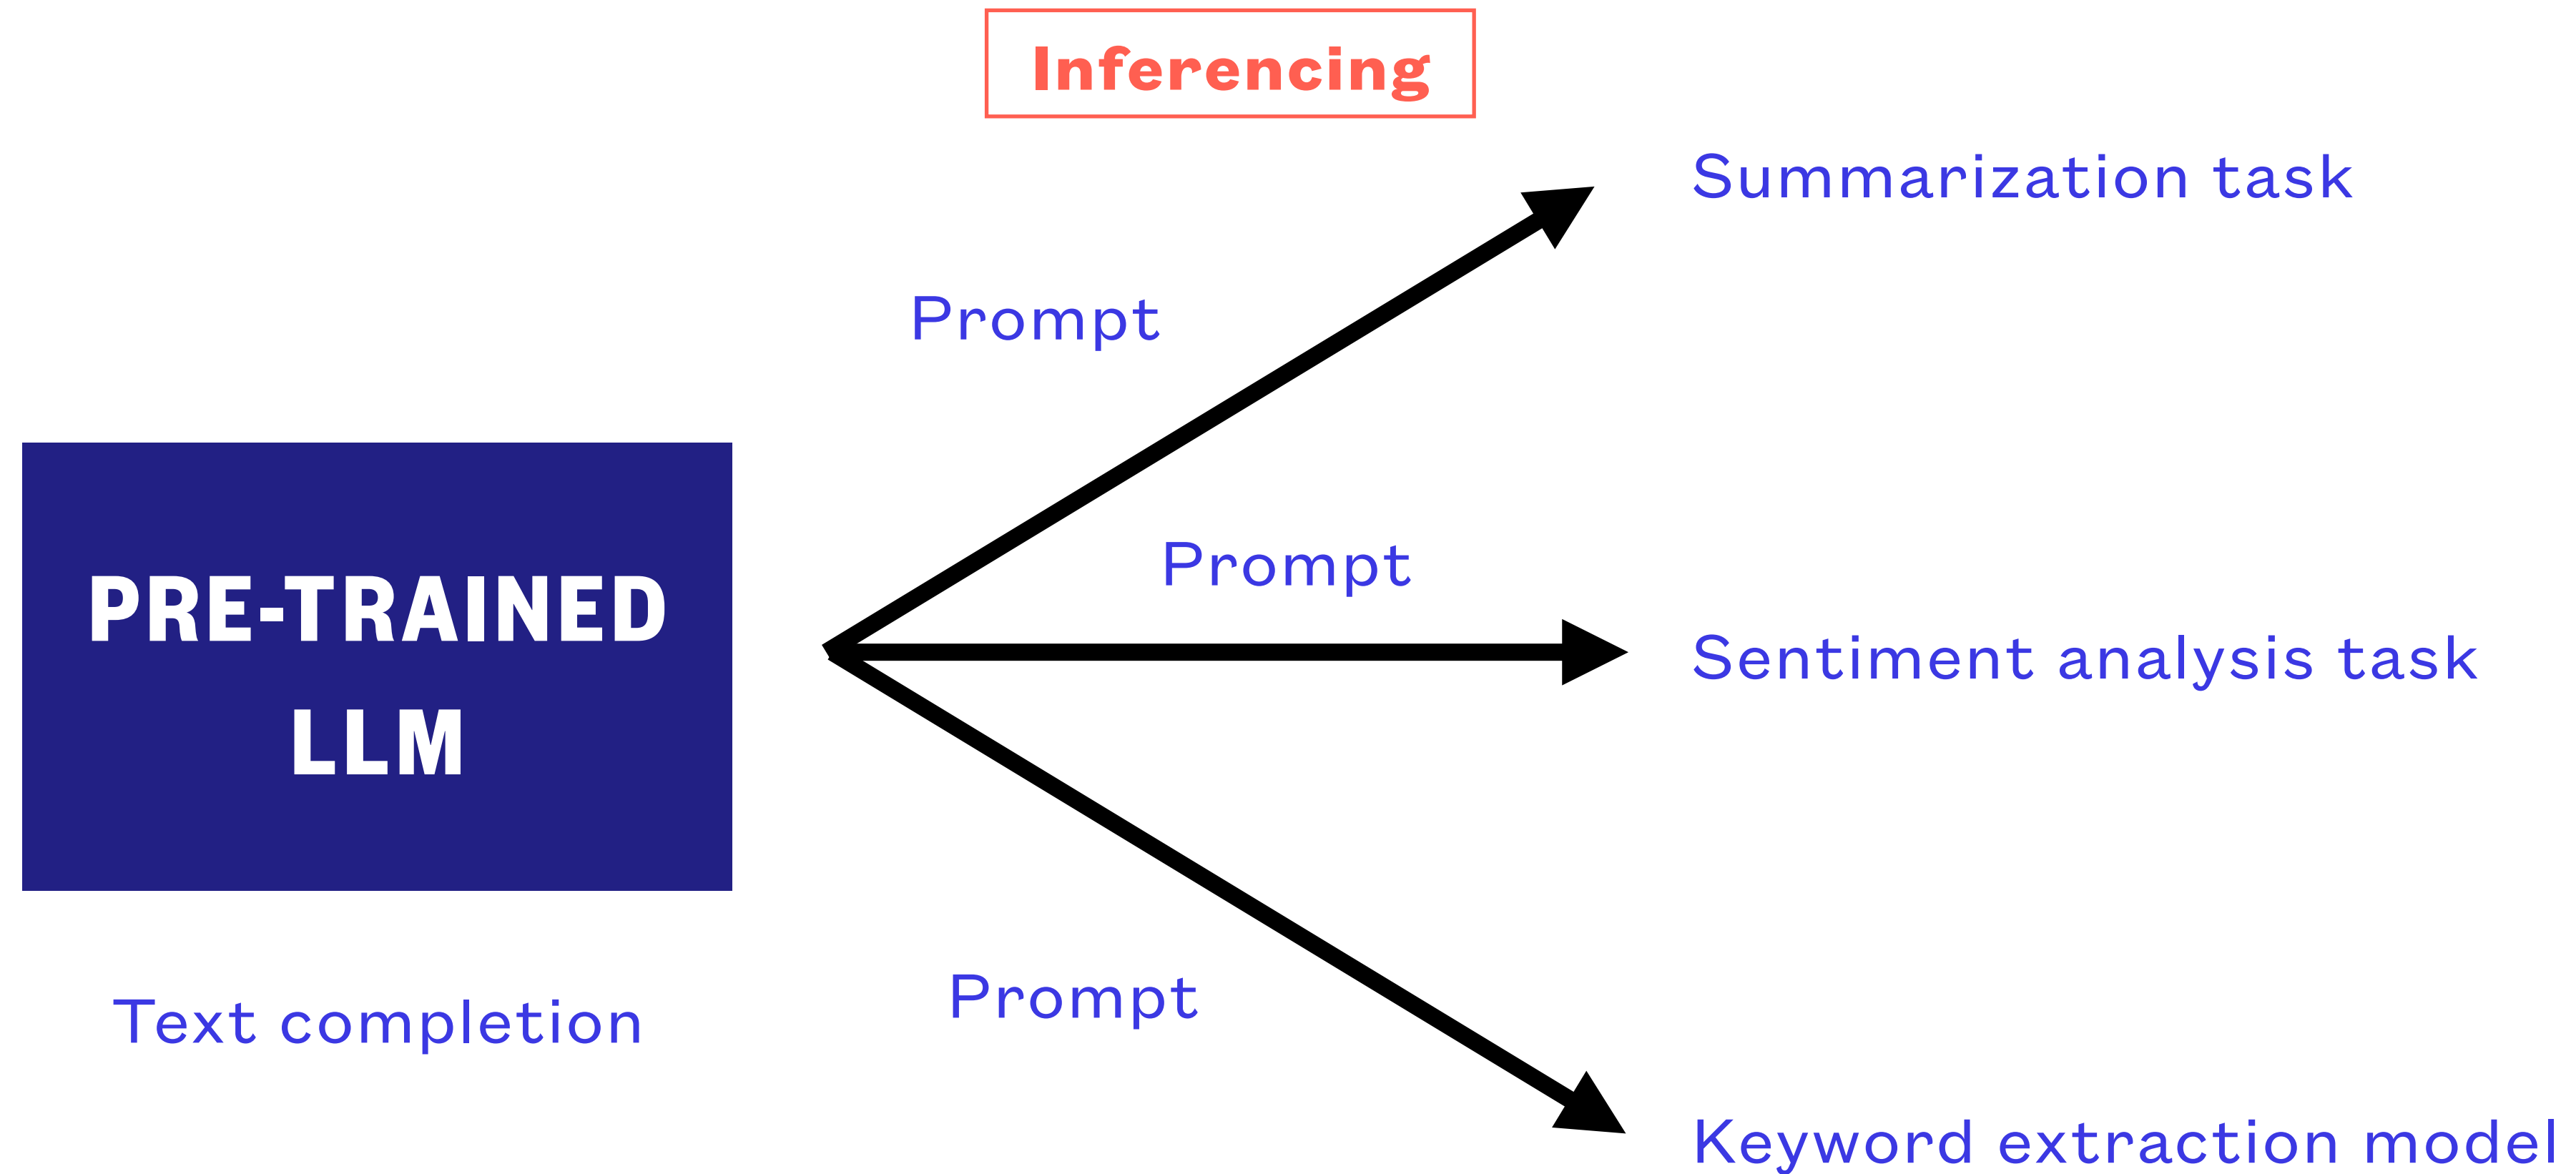
LM vs LLM



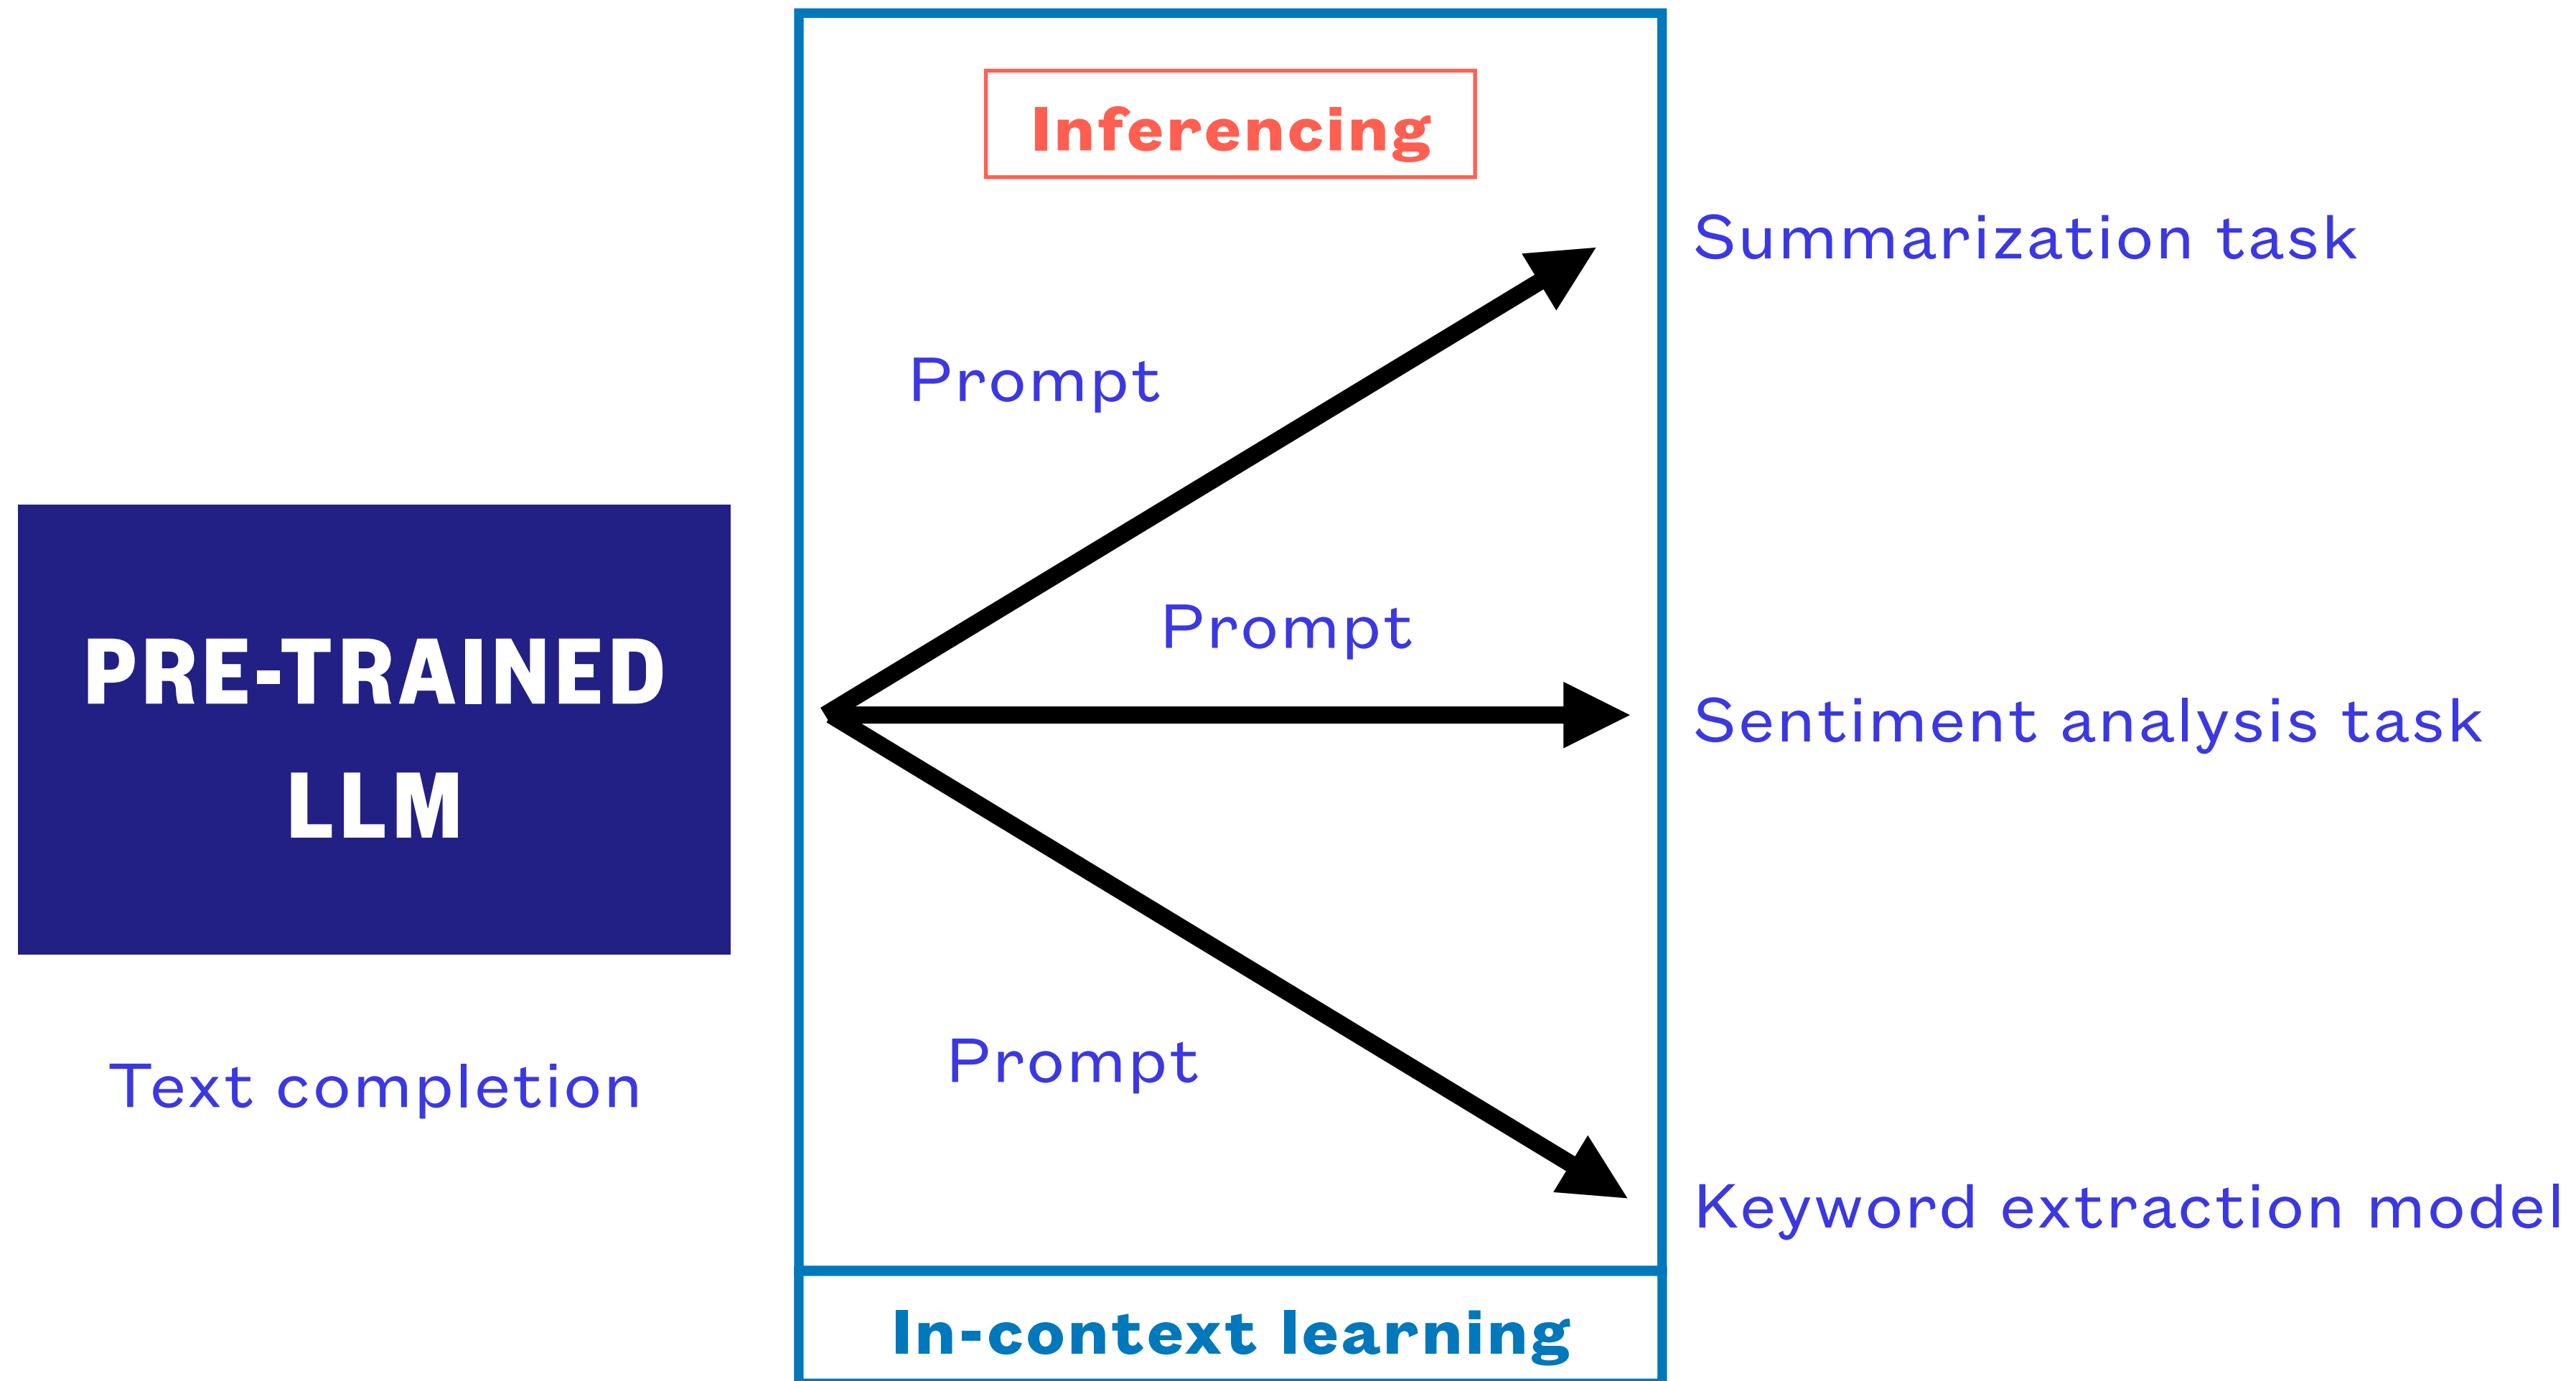
**PRE-TRAINED
LLM**

Text completion

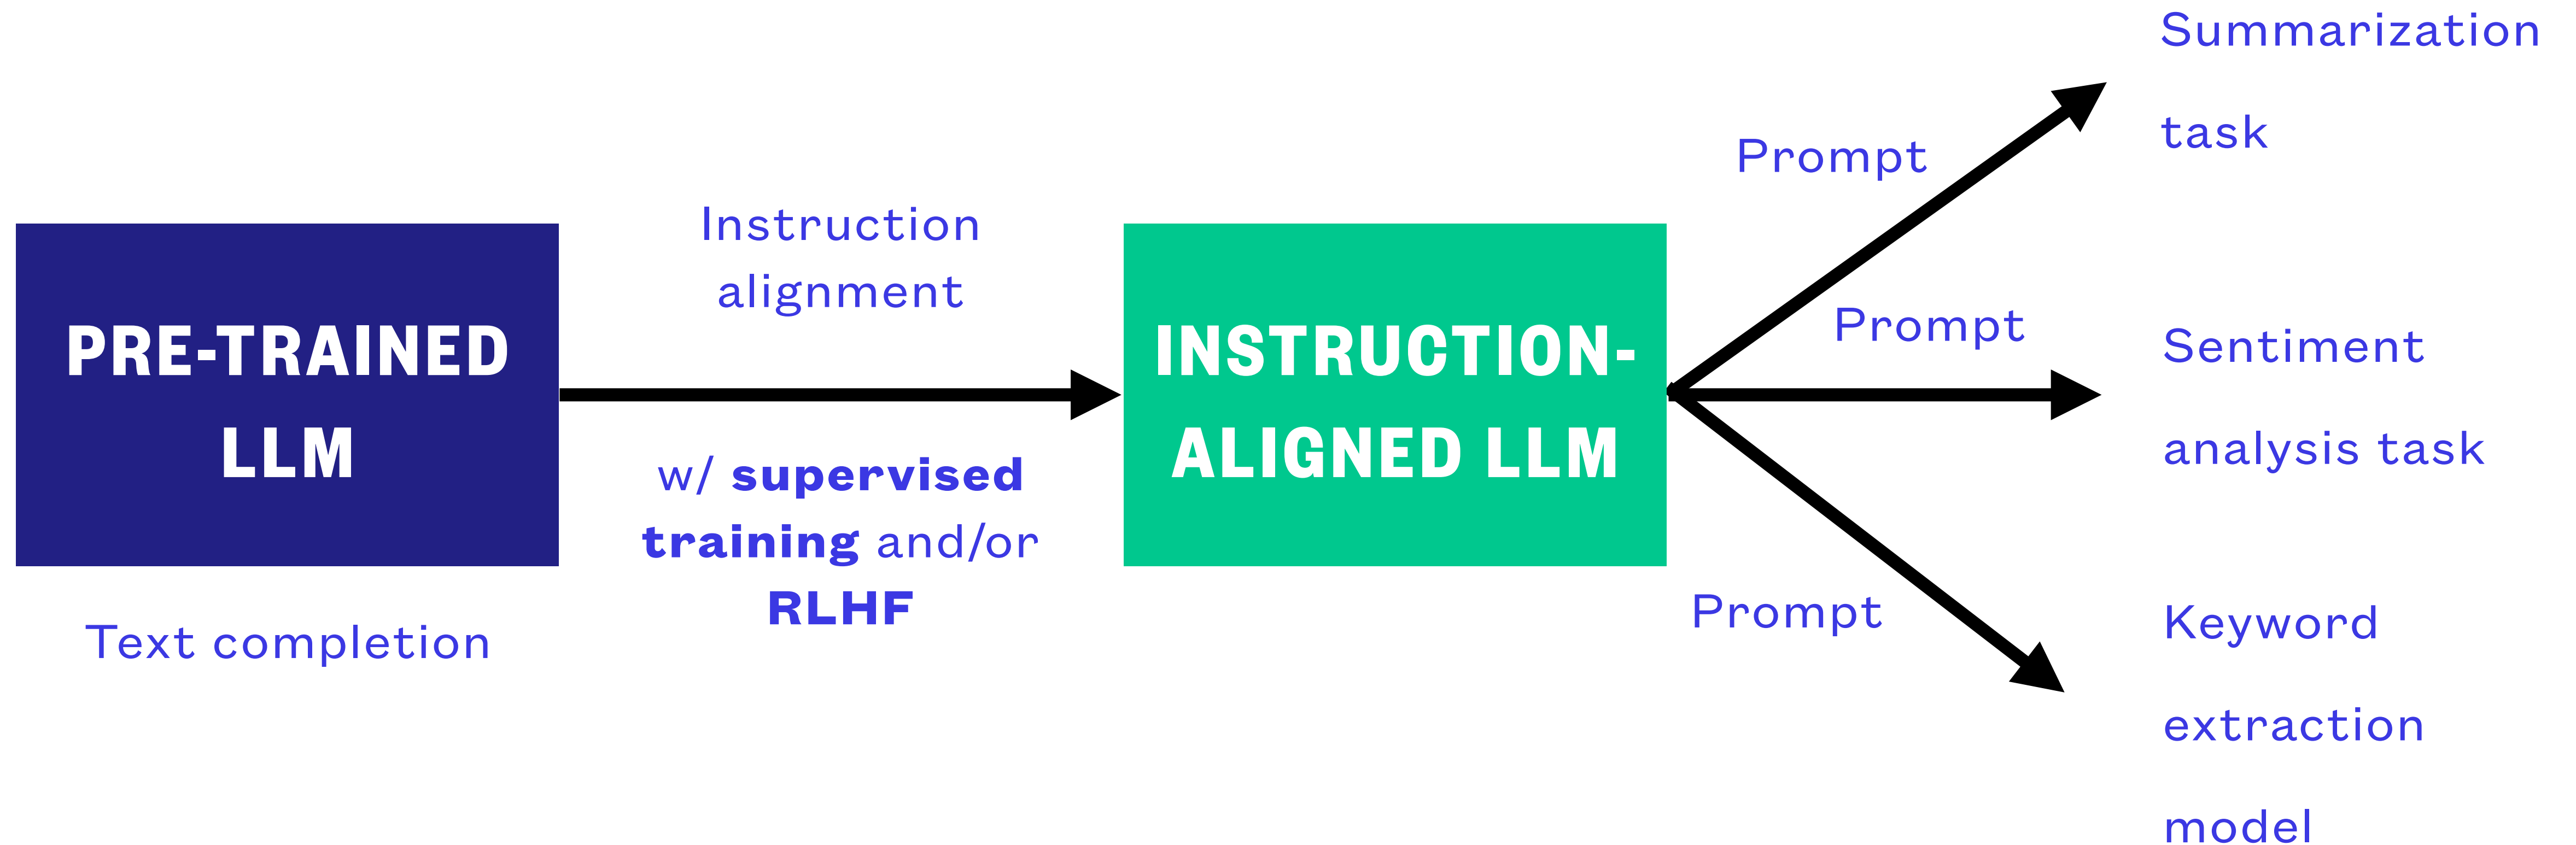
LM vs LLM



LM vs LLM



Instruction-tuned LLM



Prompt (Engineering)

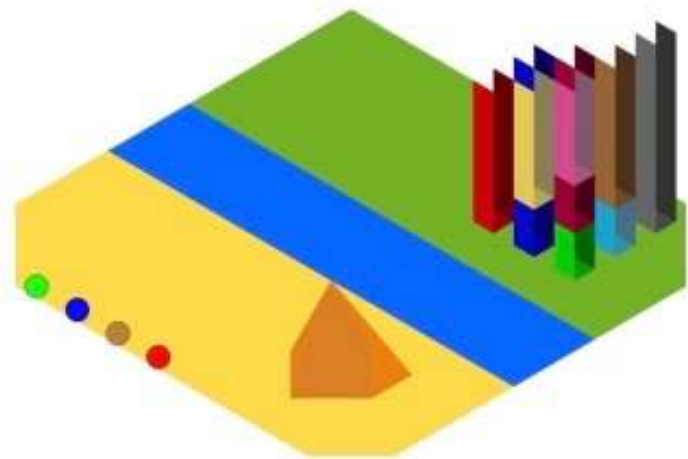
Prompt

- An **input** given to a model
- A **condition** given to a model

GPT-4 and Stable Diffusion

Prompt: A screenshot of a city-building game in 3D. The screenshot is showing a terrain where there is a river from left to right, there is a desert with a pyramid below the river, and a city with many highrises above the river. The bottom of the screen has 4 buttons with the color green, blue, brown, and red respectively.

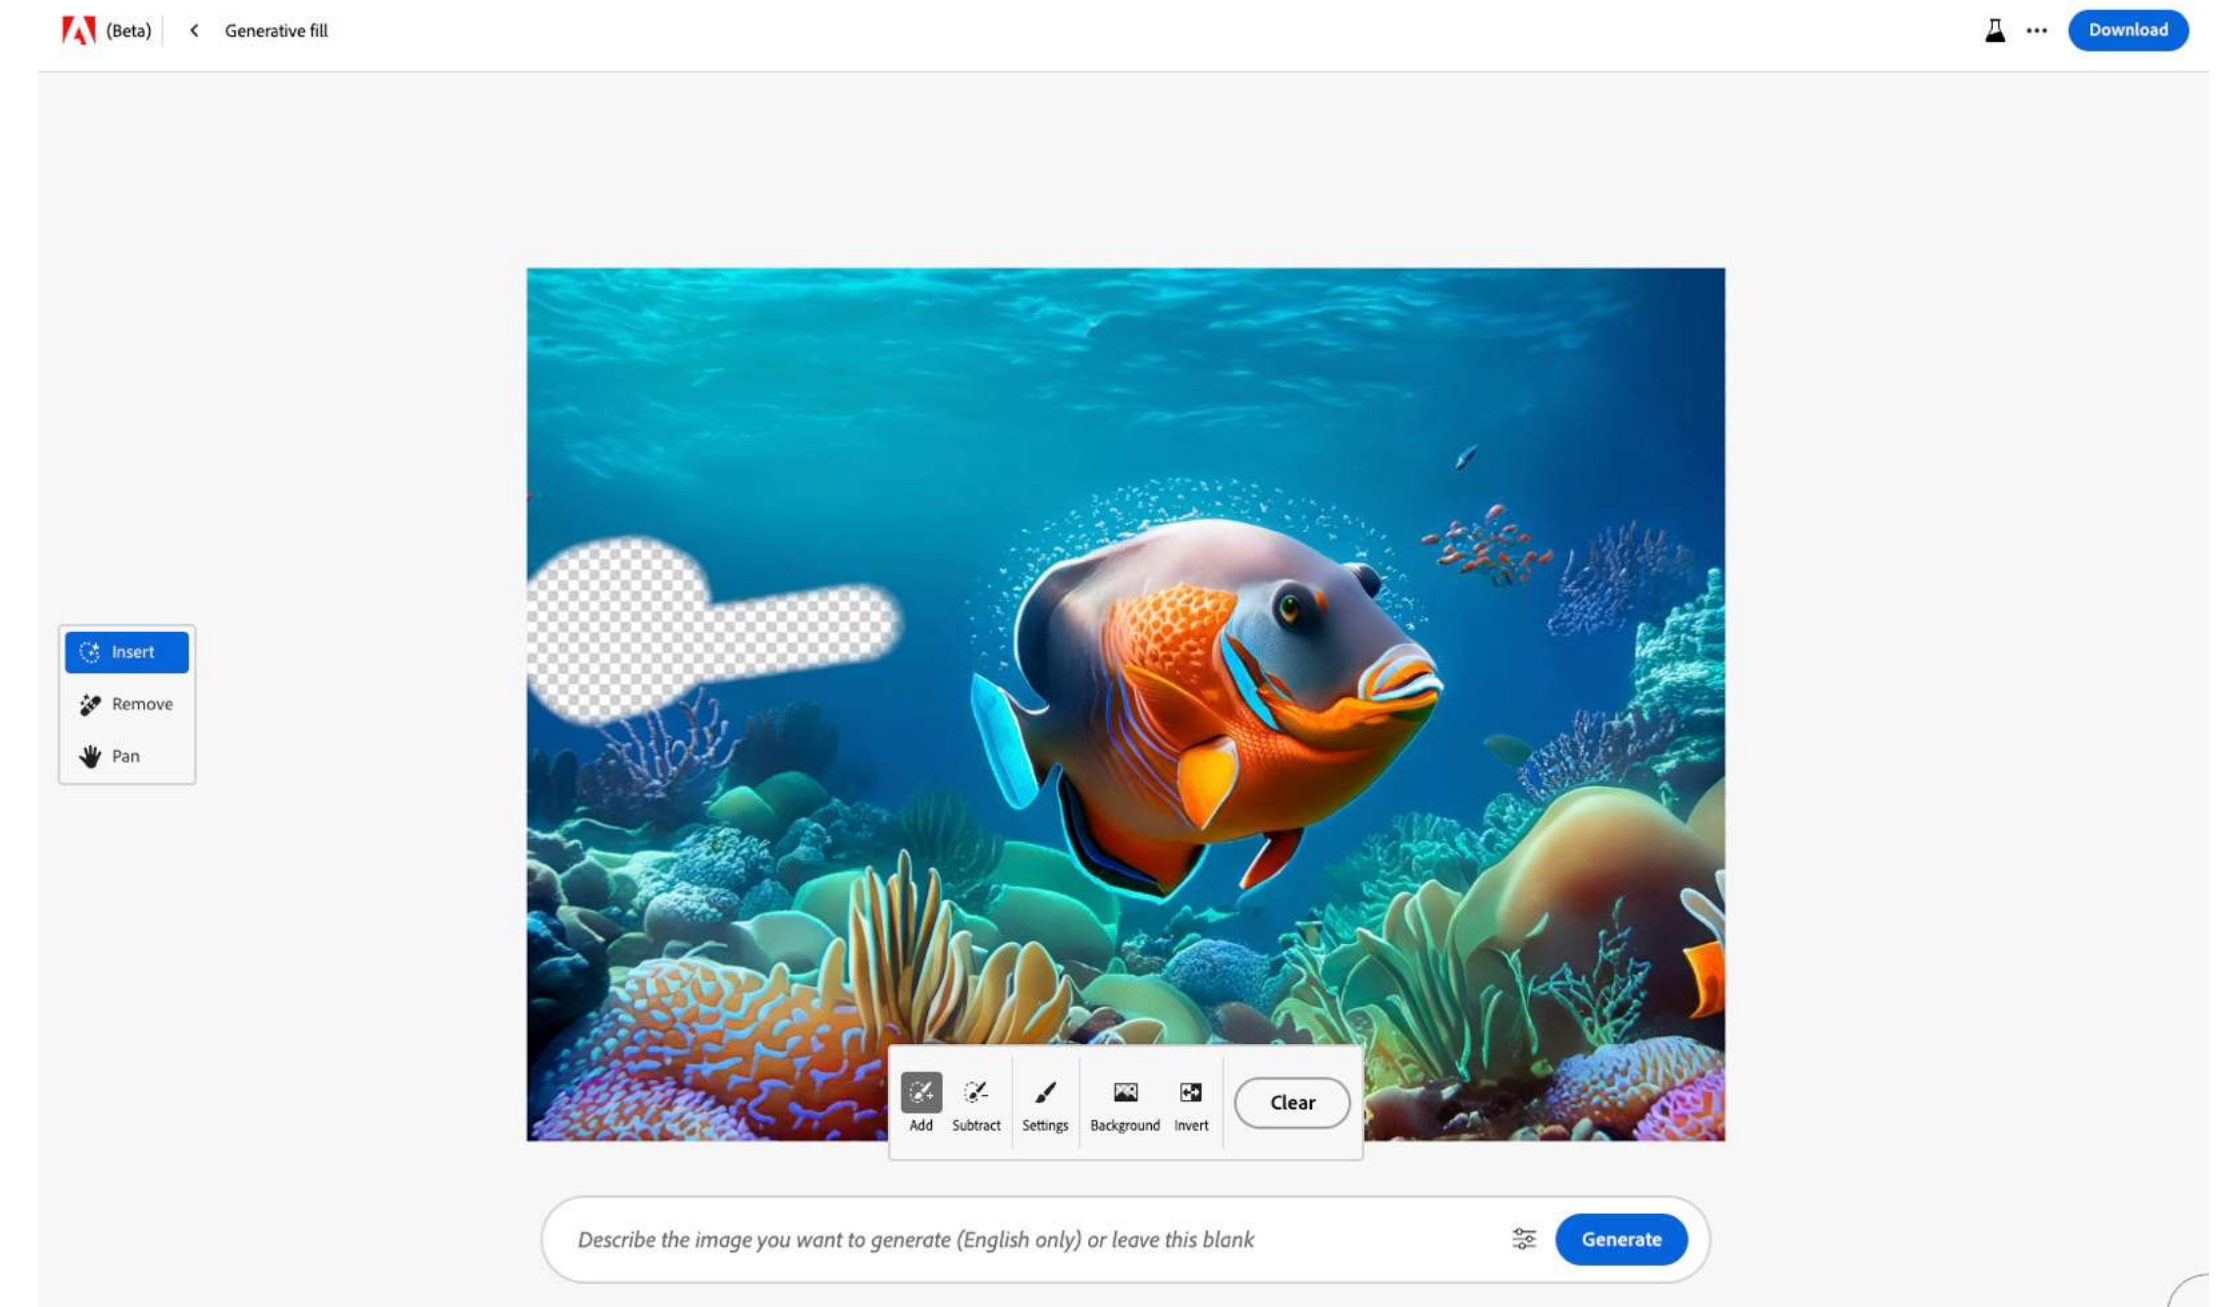
GPT-4:



Stable Diffusion v2.1:

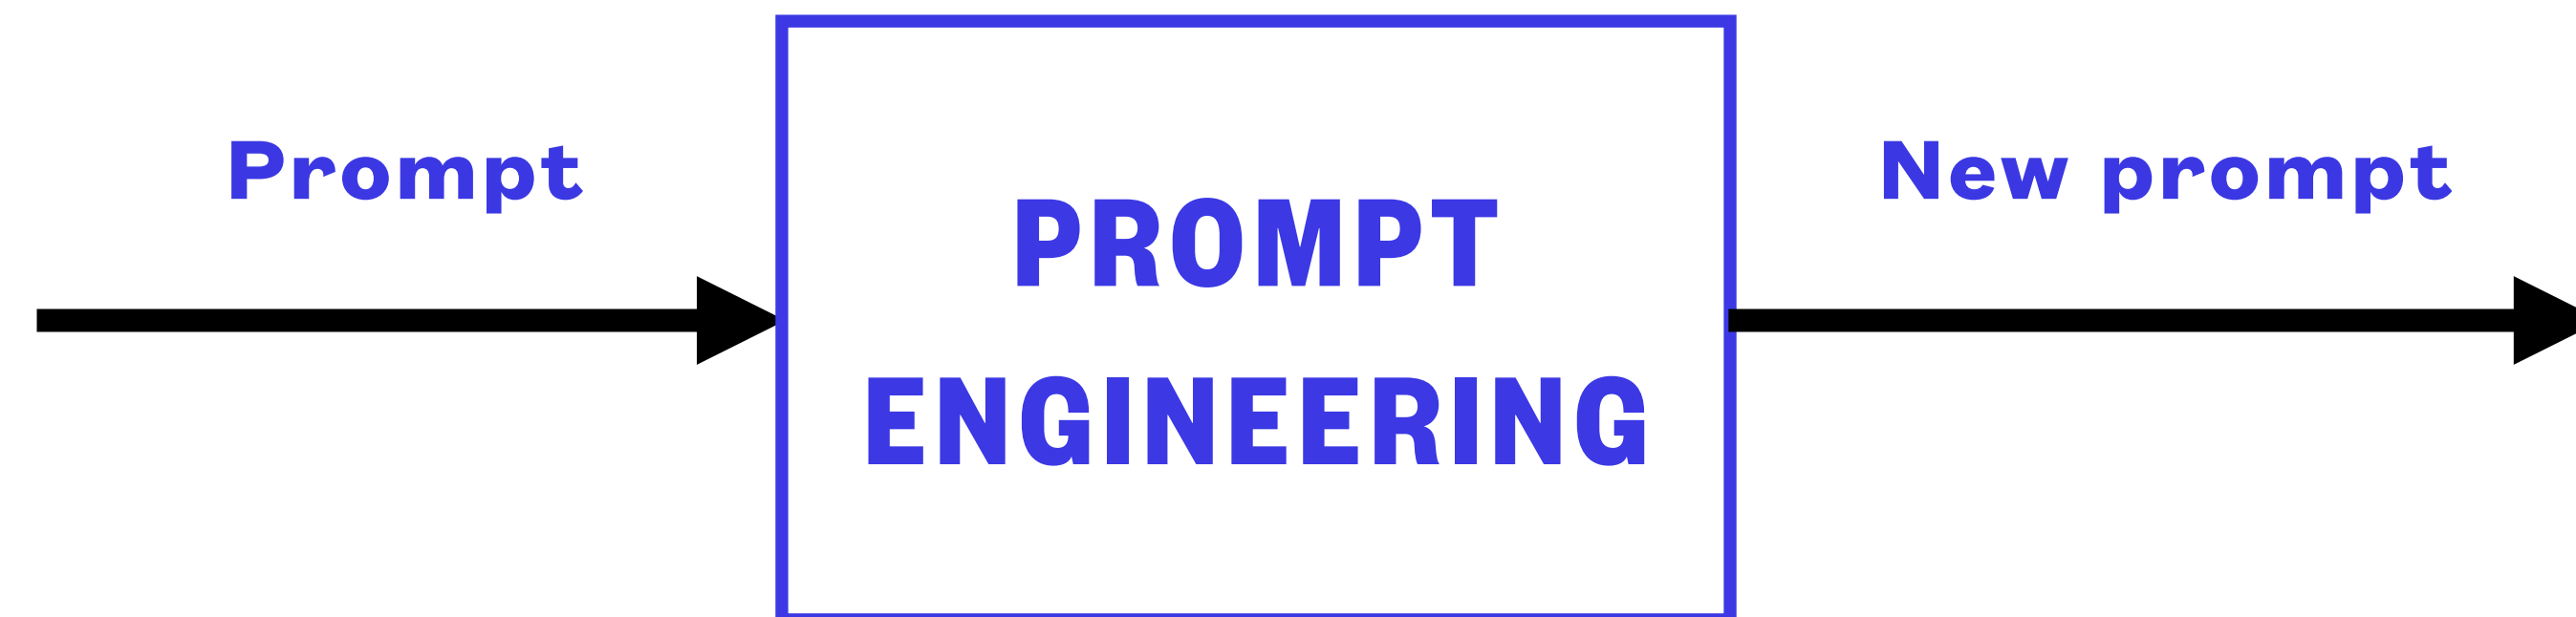


Stable Diffusion v2.1
with GPT-4's sketch:



Prompt Engineering (PE)

- A process of **composing/improving** a prompt given to a model to get desired outcomes



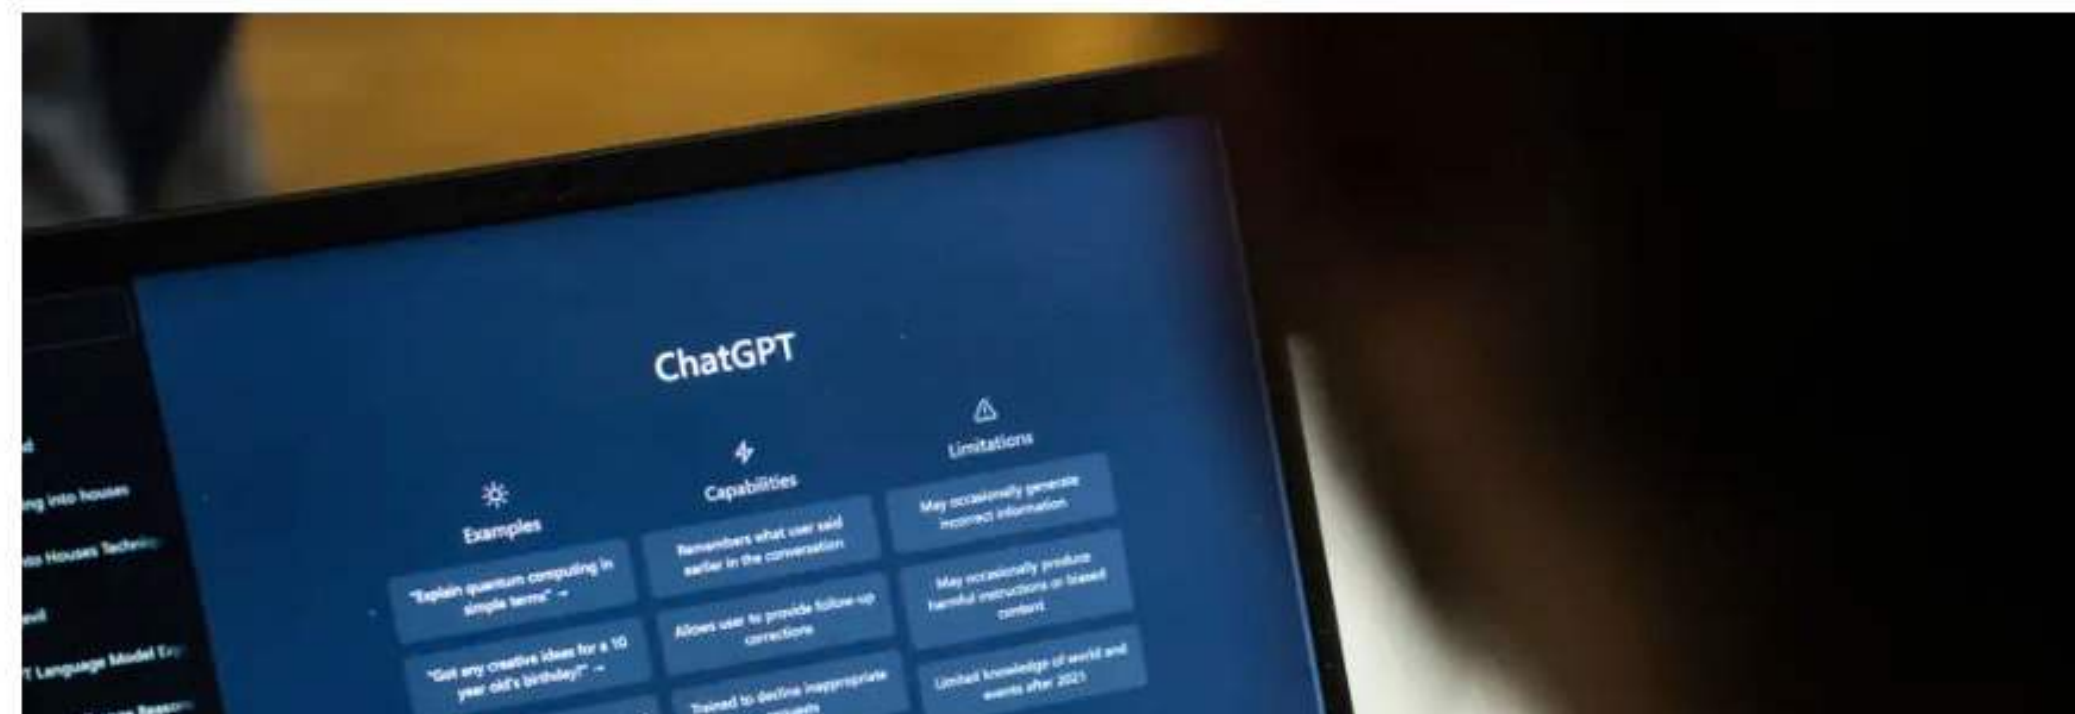
PE is a New Google Skills

- Late 2000s: How to make a search query for Google to get desired results
- **Now:** How to make a prompt for a model to get desired results

HOME > TECH

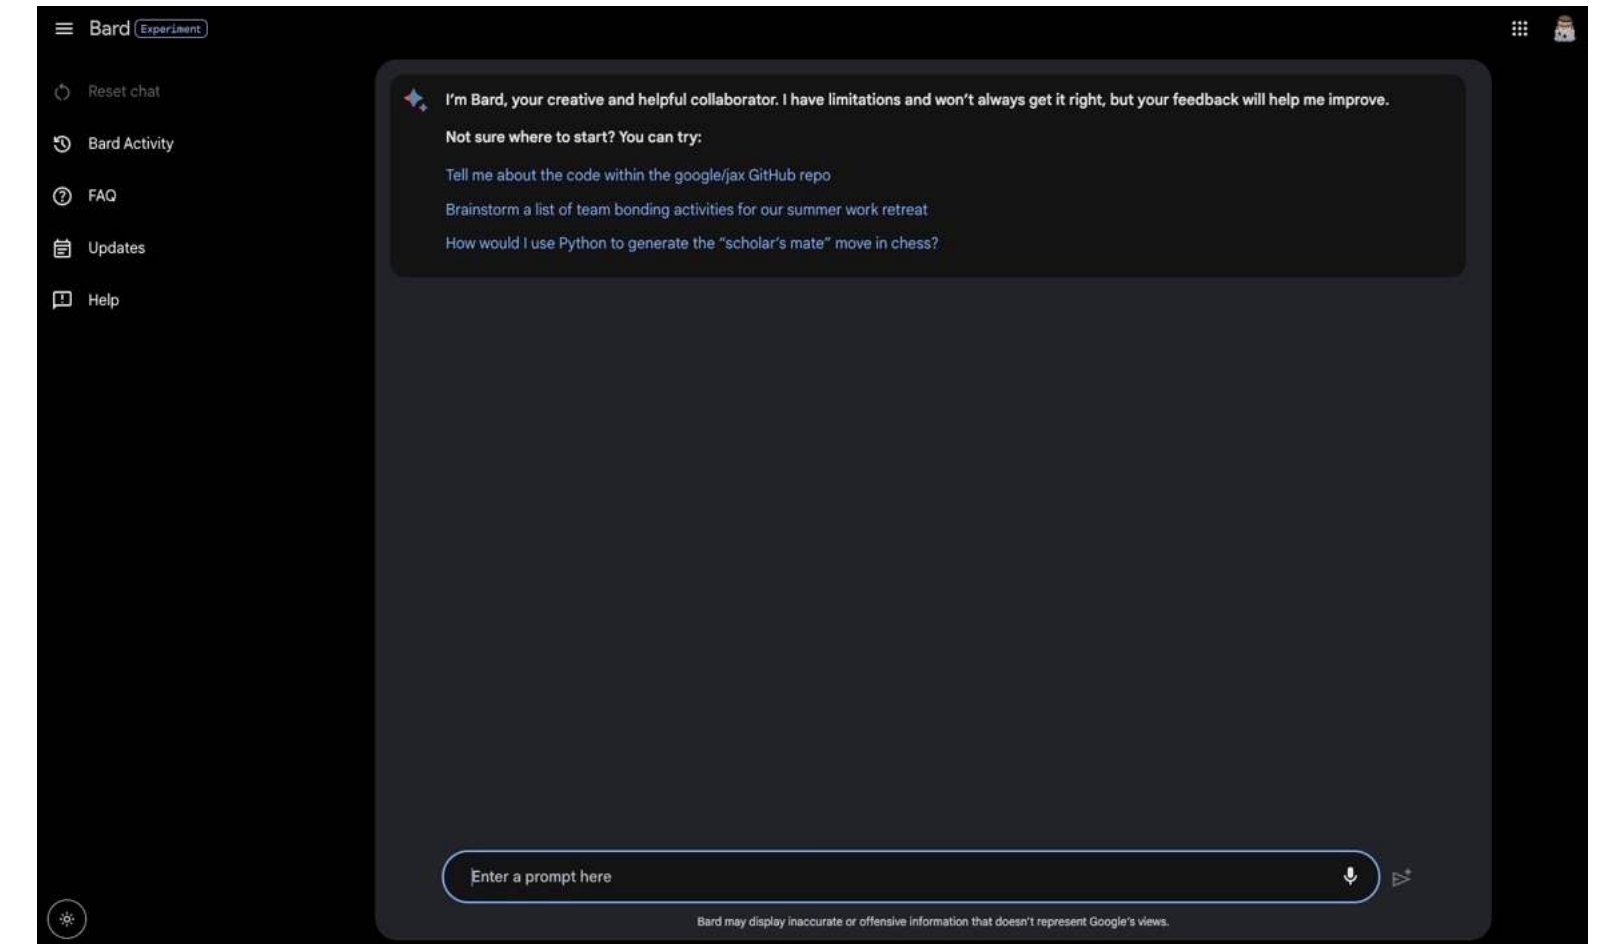
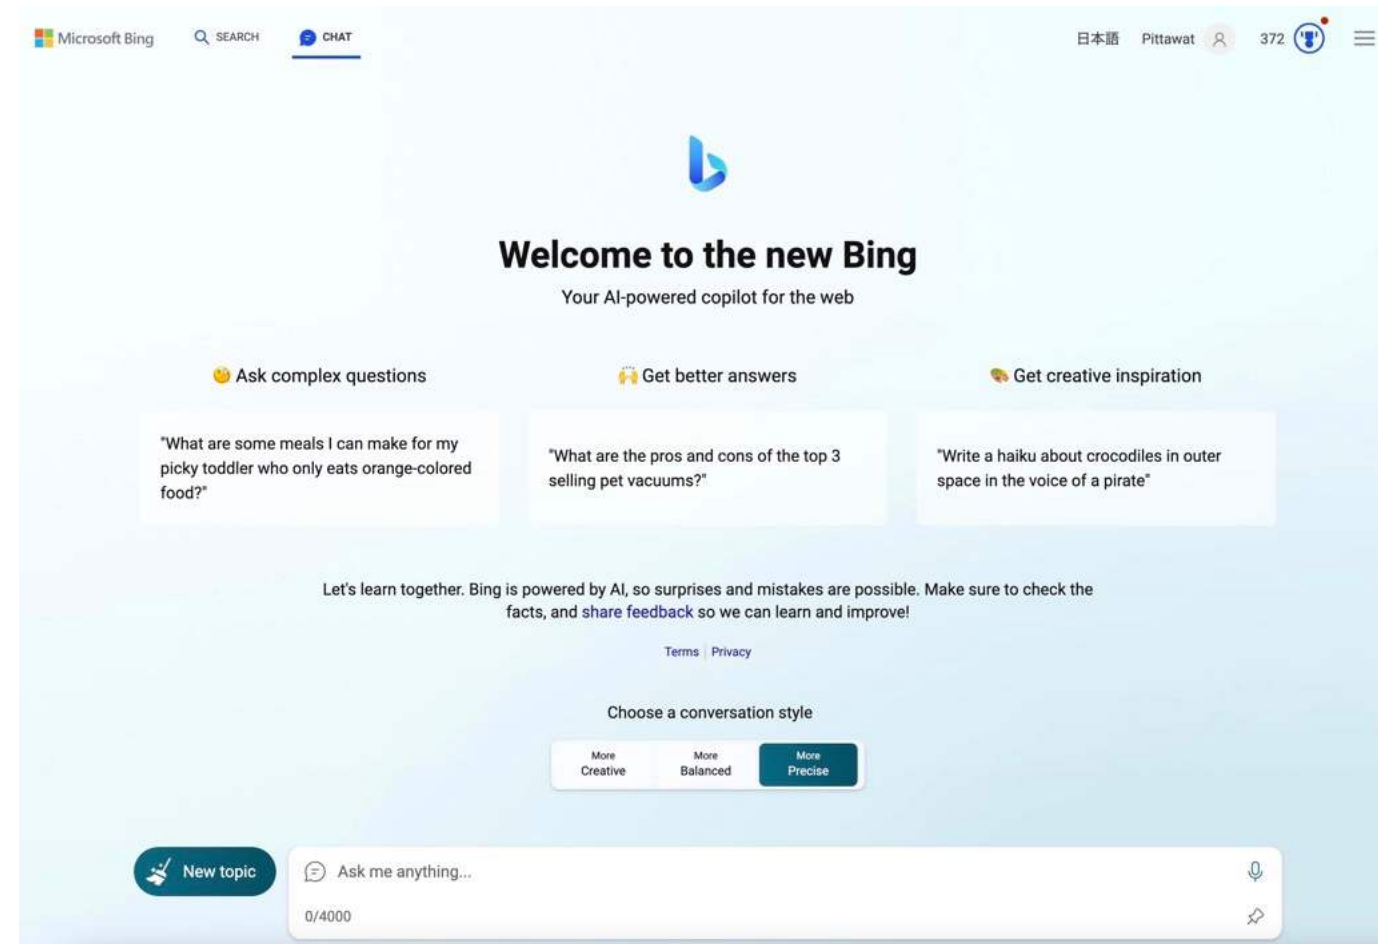
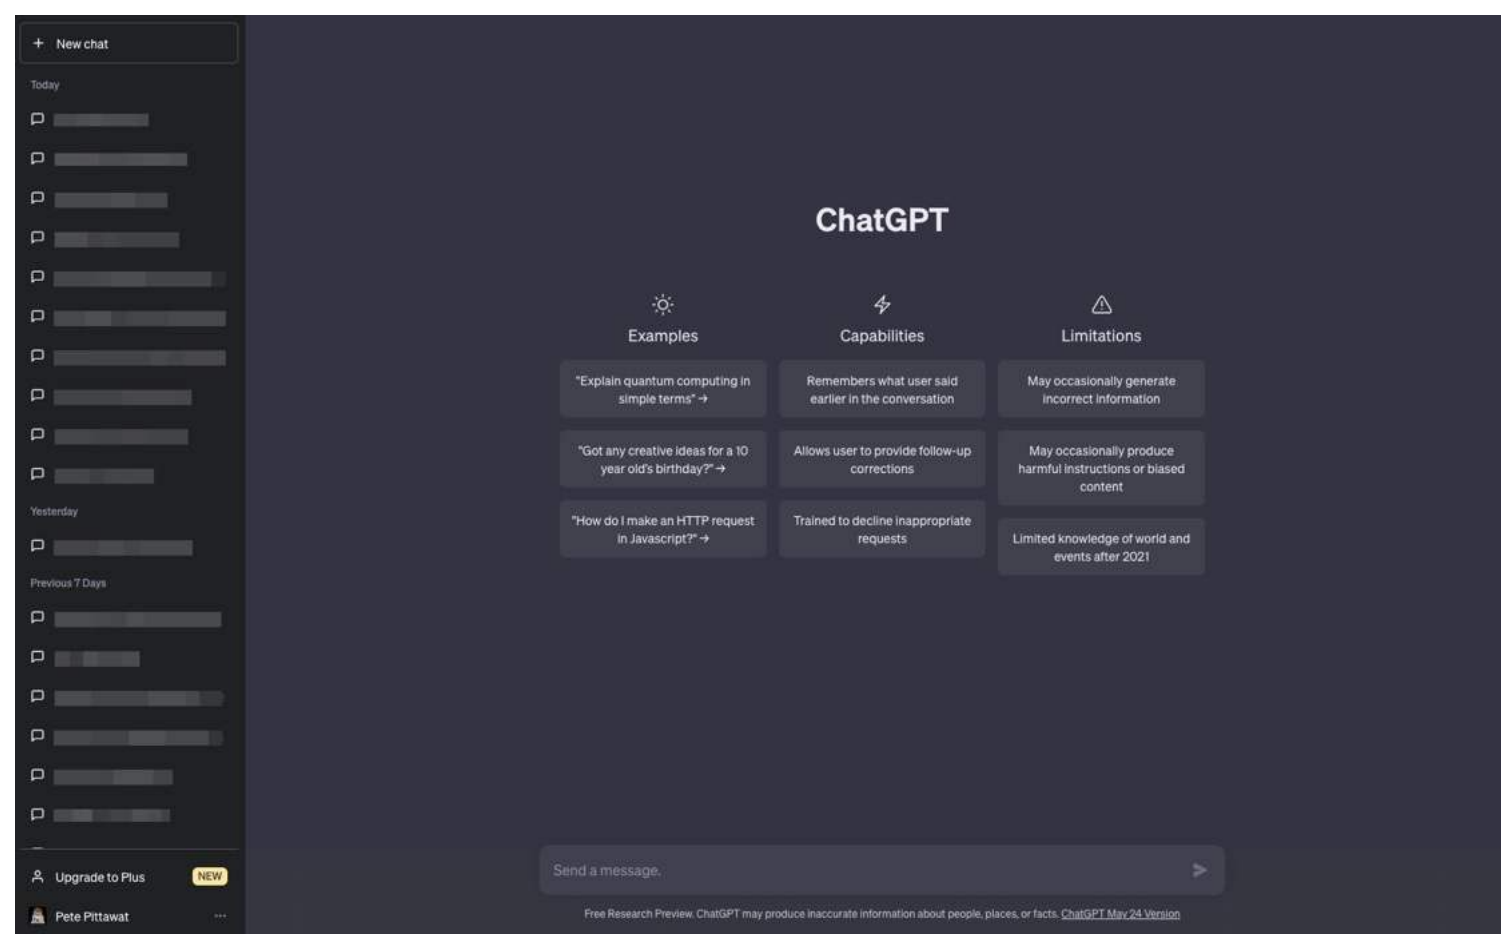
AI 'prompt engineer' jobs can pay up to \$375,000 a year and don't always require a background in tech

Britney Nguyen May 2, 2023, 12:34 AM GMT+9



PE for Instruction-aligned LLMs

- Focus on crafting **text** for chat-based LLMs, for example, ChatGPT, Bing Chat, Google Bard



"The simplest explanation is usually the best one."

Occam's razor

General Prompt Structure

- Instruction/Question
- Output indicator
- Context
- (Optional) Input
- (Optional) Example

Instruction/Question

- Instruction
 - Start with **verb**: think of it like a teacher designing a question for a student
 - Be **concise**: think of it like explaining to a child
 - Be **explicit**: what do you want and don't want
- Question
 - Start with **5W1H**
 - End with **question mark (?)** may help

Output Indicator

- Be clear about what do you **want** in which **format**
- Modifier: formal, modern, funny, etc.
- Constraint: one sentence, within code block (` ` `), etc.
- Format: JSON, XML, email, a list of 10 keywords, etc.
- Output **example** might be useful

Context

- The most **important** part
- Explain it in a sense that LLMs know **nothing** about a topic
 - Explain all the **constraints**, what **possible** and not possible, and **relevant** information, etc.
- **Role** prompting: Make LLMs become something/someone
 - Ex. teacher, Steve Jobs, Enigma, computer science professor
 - **Expert** role prompting: Ask the LLM to assume the role of an expert in order to obtain more **factual** and **accurate** results.

General Tips

- **Experiment** a lot!
- Some time it is useful to ask it in **code** (output/input)
- **Multi-turn** conversation might be useful
- Use uncommon **delimiters** to signify important information or separate sections, e.g., #####

Example 1: Information Retrieval

<https://chat.openai.com/share/9e03fcdd-8952-48fc-887a-2563d6ea51f1>

User Message 1

Model: ChatGPT

Prompt

What is a language model?

Response

A language model is a type of artificial intelligence (AI) model designed to understand and generate human language. It uses statistical techniques and pattern recognition to learn the structure, grammar, and semantics of natural language. Language models are trained on large amounts of text data and are capable of predicting the probability of a given sequence of words or generating coherent and contextually relevant text.

Language models are widely used in various natural language processing (NLP) tasks, such as machine translation, text summarization, sentiment analysis, question answering, and chatbot systems. They provide a way to process and understand human language in a computationally efficient manner.

With the advent of deep learning and neural networks, more advanced language models have been developed. These models, such as OpenAI's GPT (Generative Pre-trained Transformer), are capable of capturing complex dependencies and long-range context in text, resulting in more accurate and contextually aware language generation.

Language models like GPT-3.5, upon which I am based, have been trained on a vast amount of diverse text from the internet and other sources, allowing them to generate human-like responses and provide useful information on a wide range of topics.

User Message 2

Model: ChatGPT

Prompt

Summarize it

Response

A language model is an AI system that understands and generates human language. It learns from large amounts of text data and can predict word sequences or generate coherent text. Language models are used in various applications like translation, summarization, and chatbots. Advanced models, such as GPT-3.5, capture complex language patterns and context. They provide accurate responses and information on diverse topics.

User Message 3

Model: ChatGPT

Prompt

Explain it like I'm five years old

Response

Sure! Imagine a language model as a very smart robot that can talk and understand the things we say. It has learned a lot from reading many books and listening to people talking. This robot can help us by answering our questions and even having conversations with us. It's like having a really clever friend who knows a lot about many different things. The more the robot learns, the better it becomes at understanding and talking like a real person.

Example 2: Information Extraction

<https://chat.openai.com/share/06dc11b6-ee34-4d41-9fca-3b1dacad07c1>

User Message 1

Model: ChatGPT

Prompt

What is a company being mentioned in this content?

As Apple prepares its long-rumored jump into augmented reality on Monday, doubts have shadowed every step of the way. There are reports of frequent changes in direction and skepticism inside Apple's ranks. The device has allegedly been hard to manufacture and required numerous compromises. The process has taken years longer than Apple expected. And at a rumored \$3,000, even Apple reportedly expects slow short-term sales.

But among AR professionals, the mood is jubilant. "This is the single greatest thing that could happen to this industry," says Jay Wright, CEO of VR / AR collaboration platform Campfire 3D. "Whether you make hardware or software. We're excited about it."

Response

The company being mentioned in the given content is Apple.

User Message 1

Model: ChatGPT

Prompt

What is a company being mentioned in this content?

Question

Input

As Apple prepares its long-rumored jump into augmented reality on Monday, doubts have shadowed every step of the way. There are reports of frequent changes in direction and skepticism inside Apple’s ranks. The device has allegedly been hard to manufacture and required numerous compromises. The process has taken years longer than Apple expected. And at a rumored \$3,000, even Apple reportedly expects slow short-term sales.

But among AR professionals, the mood is jubilant. “This is the single greatest thing that could happen to this industry,” says Jay Wright, CEO of VR / AR collaboration platform Campfire 3D. “Whether you make hardware or software. We’re excited about it.”

Response

The company being mentioned in the given content is Apple.

User Message 2

Model: ChatGPT

Prompt

What is the main topic in of this excerpt?

Response

The main topic of this excerpt is Apple's rumored entry into the augmented reality (AR) market.

Example 3: Code Generation

<https://chat.openai.com/share/06dc11b6-ee34-4d41-9fca-3b1dacad07c1>

User Message 1

Model: ChatGPT

Prompt

Write a TypeScript code to validate format of email string. Output only code without explanation.

Response

Sure! Here's the TypeScript code to validate the format of an email string:

```
...  
  
function validateEmail(email: string): boolean {  
    const emailRegex = /^[A-Za-z0-9._%+-]+@[A-Za-z0-9.-]+\.[A-Za-z]{2,}$/;  
    return emailRegex.test(email);  
}  
...
```

Please note that this code uses a regular expression to validate the email format. The regular expression `/^[A-Za-z0-9._%+-]+@[A-Za-z0-9.-]+\.[A-Za-z]{2,}$/` ensures that the email string contains alphanumeric characters, dots, underscores, percent signs, plus signs, and hyphens in the local part before the "@" symbol. The domain part should contain alphanumeric characters, dots, and hyphens, followed by a dot and at least two characters for the top-level domain.

The `validateEmail` function takes an email string as an argument and returns a boolean value indicating whether the email format is valid.

User Message 1

Model: ChatGPT

Prompt

Instruction

Output

Write a TypeScript code to validate format of email string. Output only code without explanation.

Response

Sure! Here's the TypeScript code to validate the format of an email string:

```
...  
  
function validateEmail(email: string): boolean {  
  const emailRegex = /^[A-Za-z0-9._%+-]+@[A-Za-z0-9.-]+\.[A-Za-z]{2,}$/;  
  return emailRegex.test(email);  
}  
  
...
```

Please note that this code uses a regular expression to validate the email format. The regular expression `/^[A-Za-z0-9._%+-]+@[A-Za-z0-9.-]+\.[A-Za-z]{2,}$/` ensures that the email string contains alphanumeric characters, dots, underscores, percent signs, plus signs, and hyphens in the local part before the "@" symbol. The domain part should contain alphanumeric characters, dots, and hyphens, followed by a dot and at least two characters for the top-level domain.

The validateEmail function takes an email string as an argument and returns a boolean value indicating whether the email format is valid.

Prompt Engineering (PE)

PE Techniques: A Summary

- **Basics**

- Zero-shot prompting
- One-shot prompting
- Few-shot prompting

- **Chain-of-Thought**

- Chain-of-Thought (CoT) prompting
- Zero-shot CoT prompting
- Self-consistency
- Tree of Thoughts (ToT) prompting

- Active-Prompt

- Multi-modal CoT

- Least-to-Most (LtM) prompting

- **Others**

- Generated Knowledge prompting

- Automatic Prompt Engineer (APE)

- Directional Stimulus prompting

- ReAct

PE: Basics

Zero-shot Prompting

- Instruct a LLM to perform a task that it has **not been seen** in training set

Prompt

Classify the text into neutral, negative or positive.

Text: I think the vacation is okay.

Sentiment:

Response

The sentiment of the text "I think the vacation is okay" is neutral.

One-shot Prompting

- Instruct a LLM to perform a task that it has **not been seen** in training set by providing **one example**

Prompt

Classify the text into neutral, negative or positive.

Text: I think the vacation is okay.

Sentiment: Neutral

Text: I love this hotel!

Sentiment:

Response

Positive

One-shot Prompting

- Instruct a LLM to perform a task that it has **not been seen** in training set by providing **one example**

Prompt

Classify the text into neutral, negative or positive.

Text: I think the vacation is okay.

Sentiment: Neutral

Example

Text: I love this hotel!

Sentiment:

Response

Positive

Few-shot Prompting

- Instruct a LLM to perform a task that it has **not been seen** in training set by providing **a few examples**

Prompt

Classify the text into neutral, negative or positive.

Text: I think the vacation is okay.

Sentiment: Neutral

Text: I love this hotel!

Sentiment: Positive

Text: I don't like this food.

Sentiment:

Response

Negative

Few-shot Prompting

- Instruct a LLM to perform a task that it has **not been seen** in training set by providing **a few examples**

Prompt

Classify the text into neutral, negative or positive.

Text: I think the vacation is okay.

Sentiment: Neutral

Text: I love this hotel!

Sentiment: Positive

Text: I don't like this food.

Sentiment:

Examples

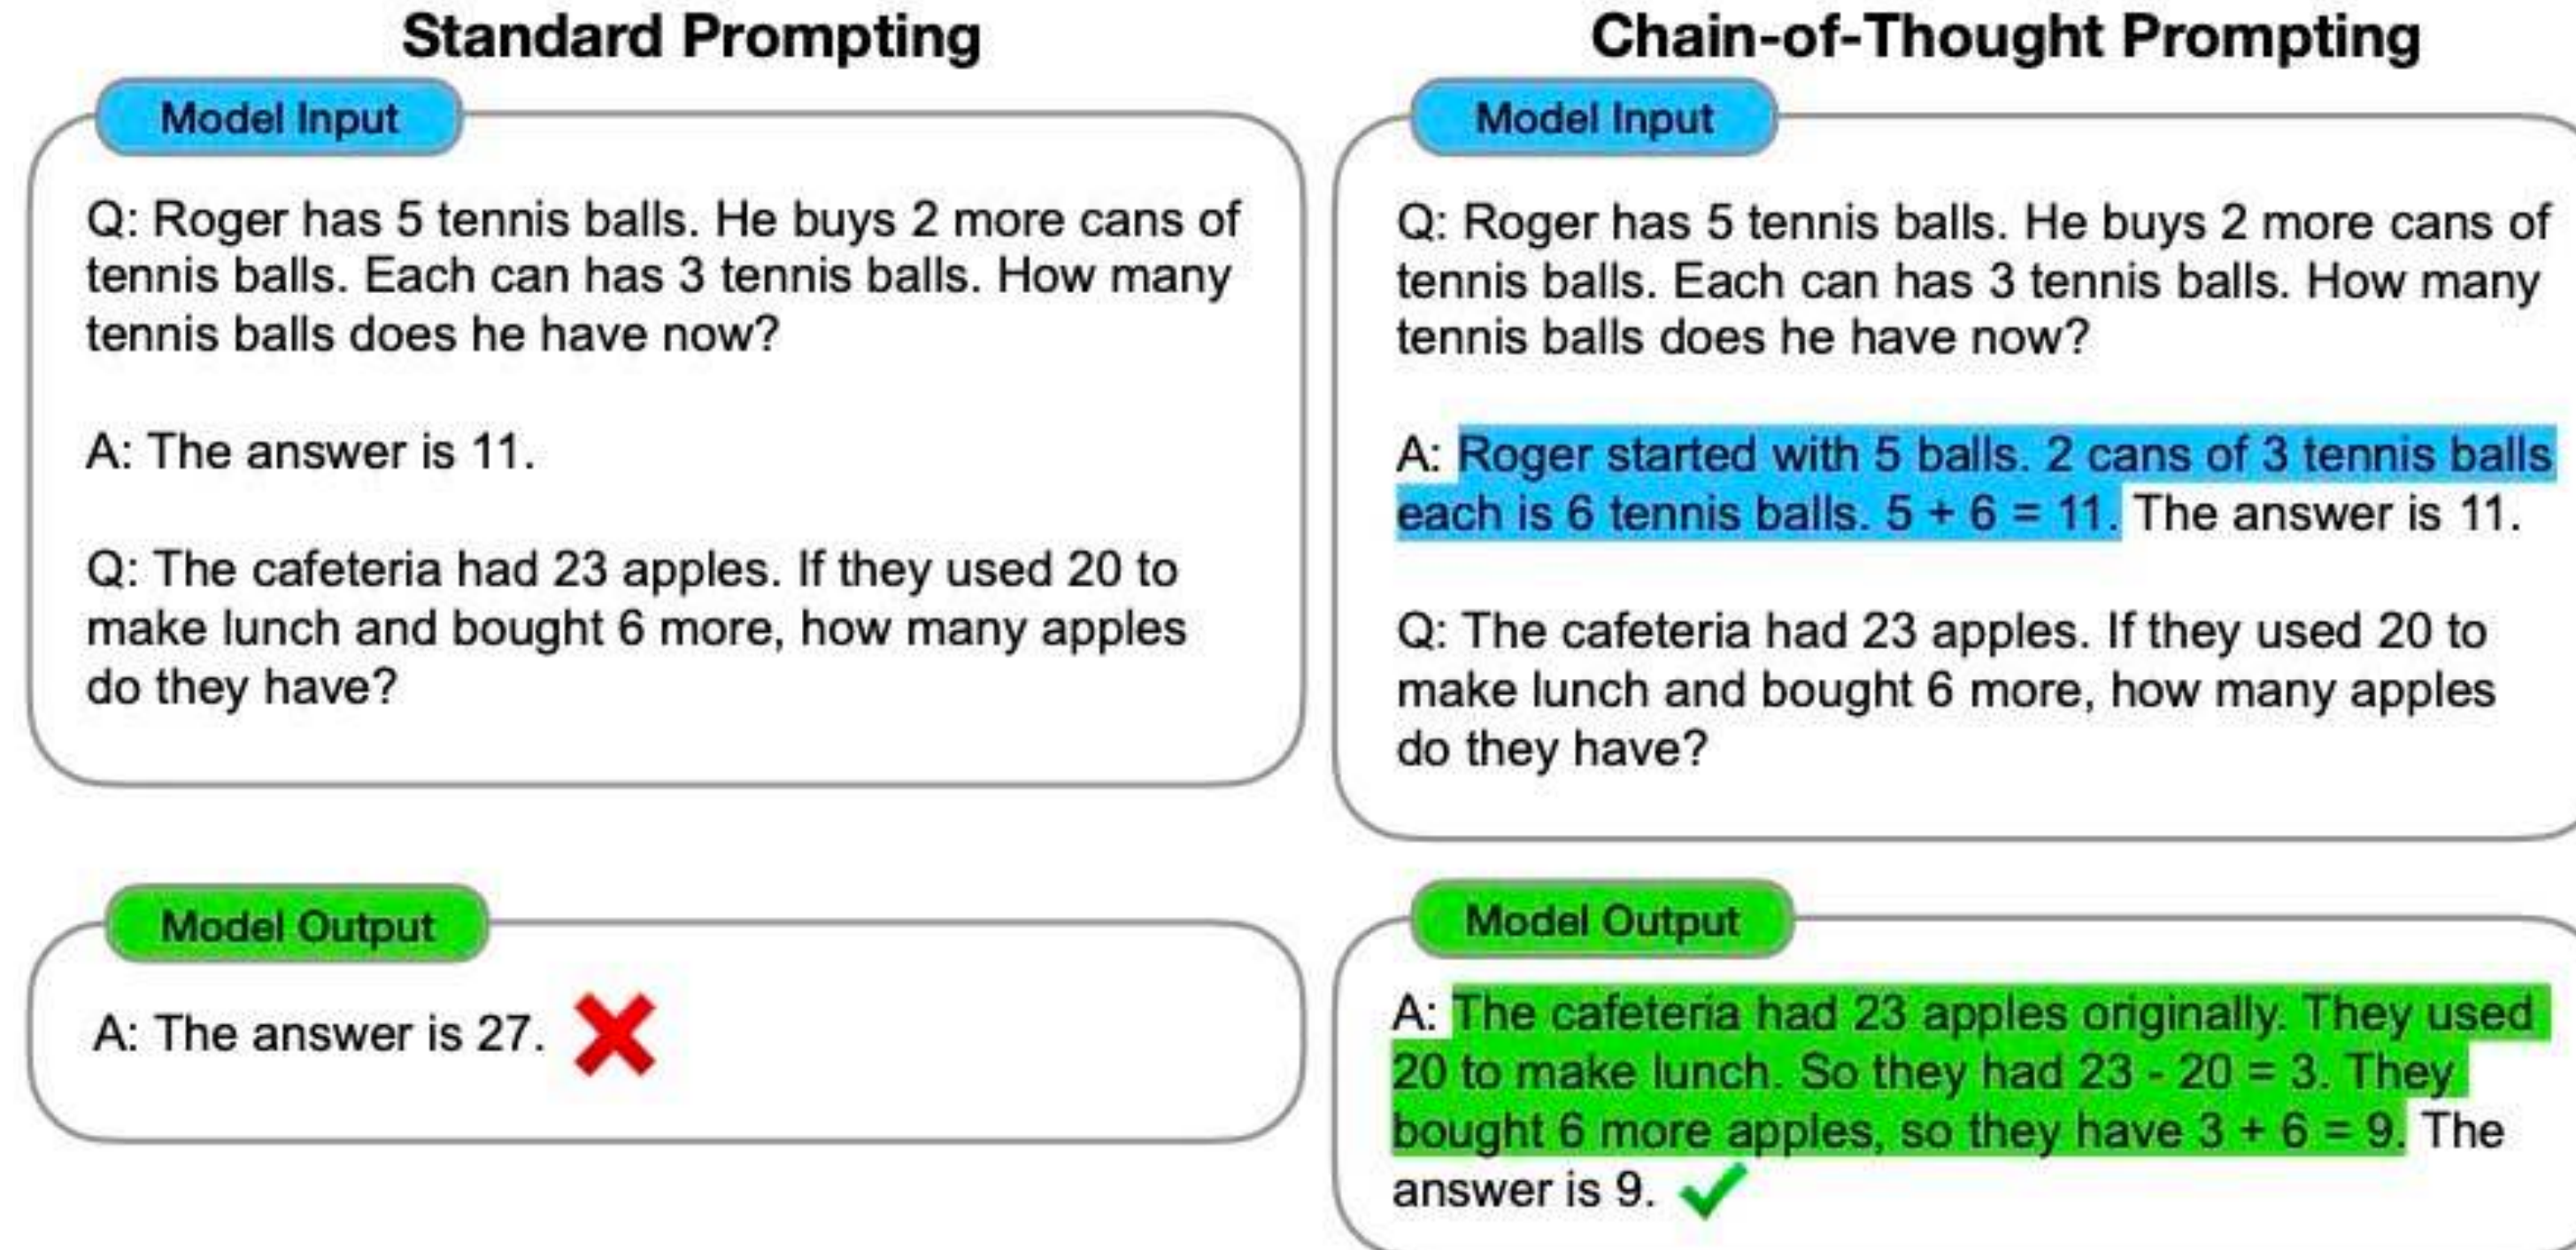
Response

Negative

PE: Chain-of-Thoughts

Chain-of-Thought (CoT) Prompting

- Improve performance through **reasoning steps**



Zero-shot CoT Prompting

- Improve performance through reasoning steps **without** providing any examples

(d) Zero-shot-CoT (Ours)

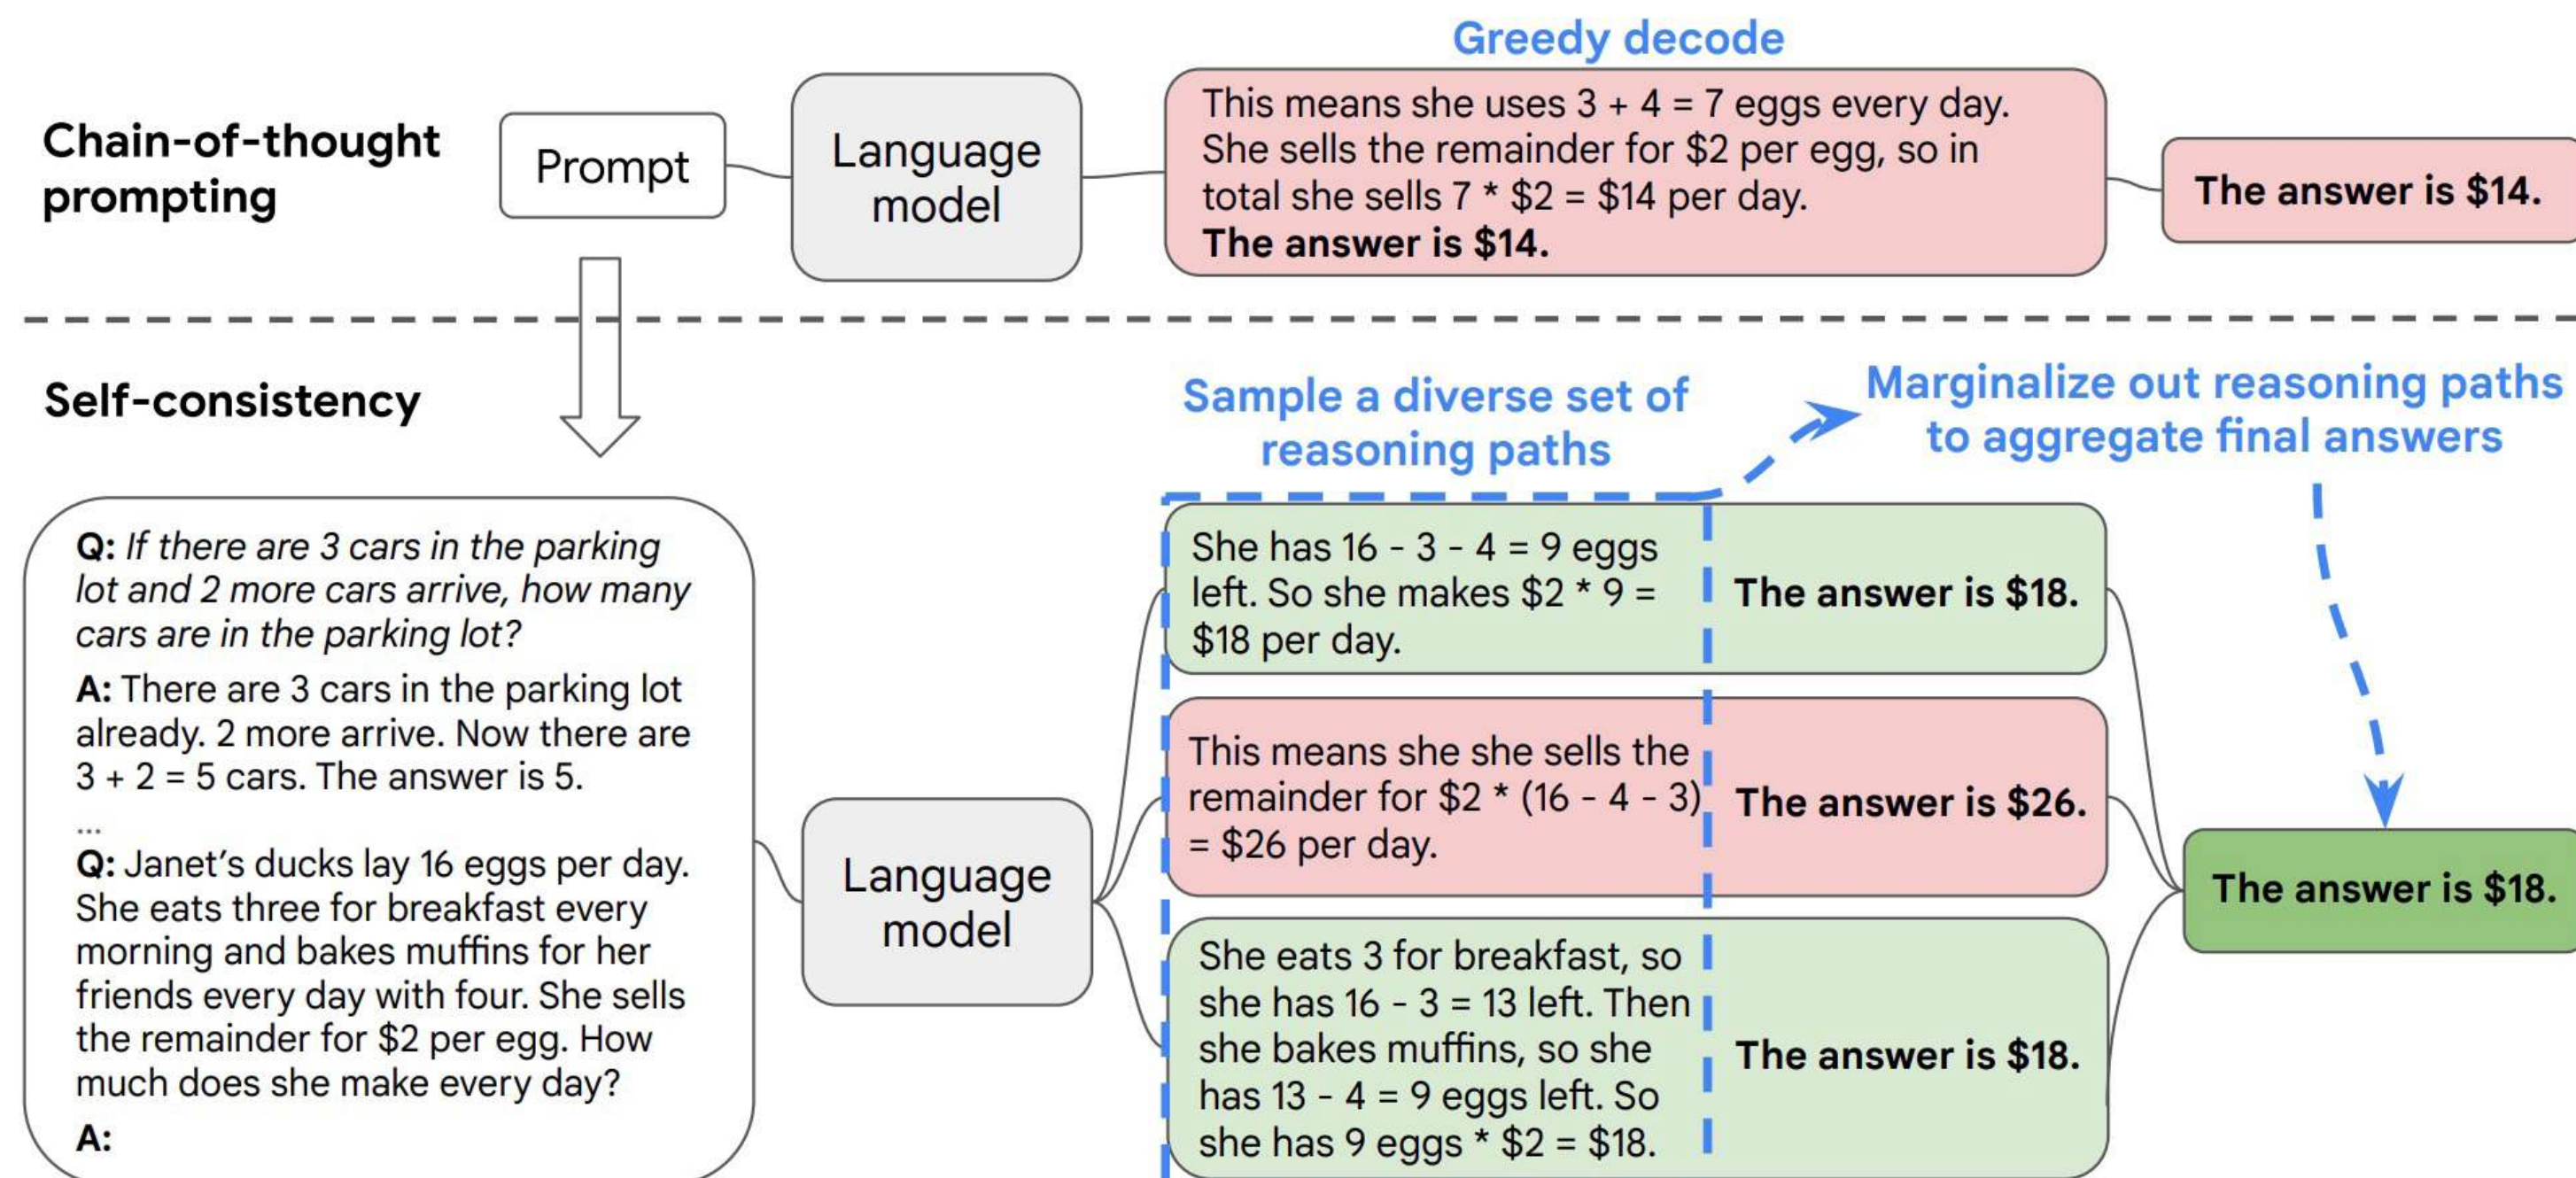
Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A: **Let's think step by step.**

(Output) *There are 16 balls in total. Half of the balls are golf balls. That means that there are 8 golf balls. Half of the golf balls are blue. That means that there are 4 blue golf balls. ✓*

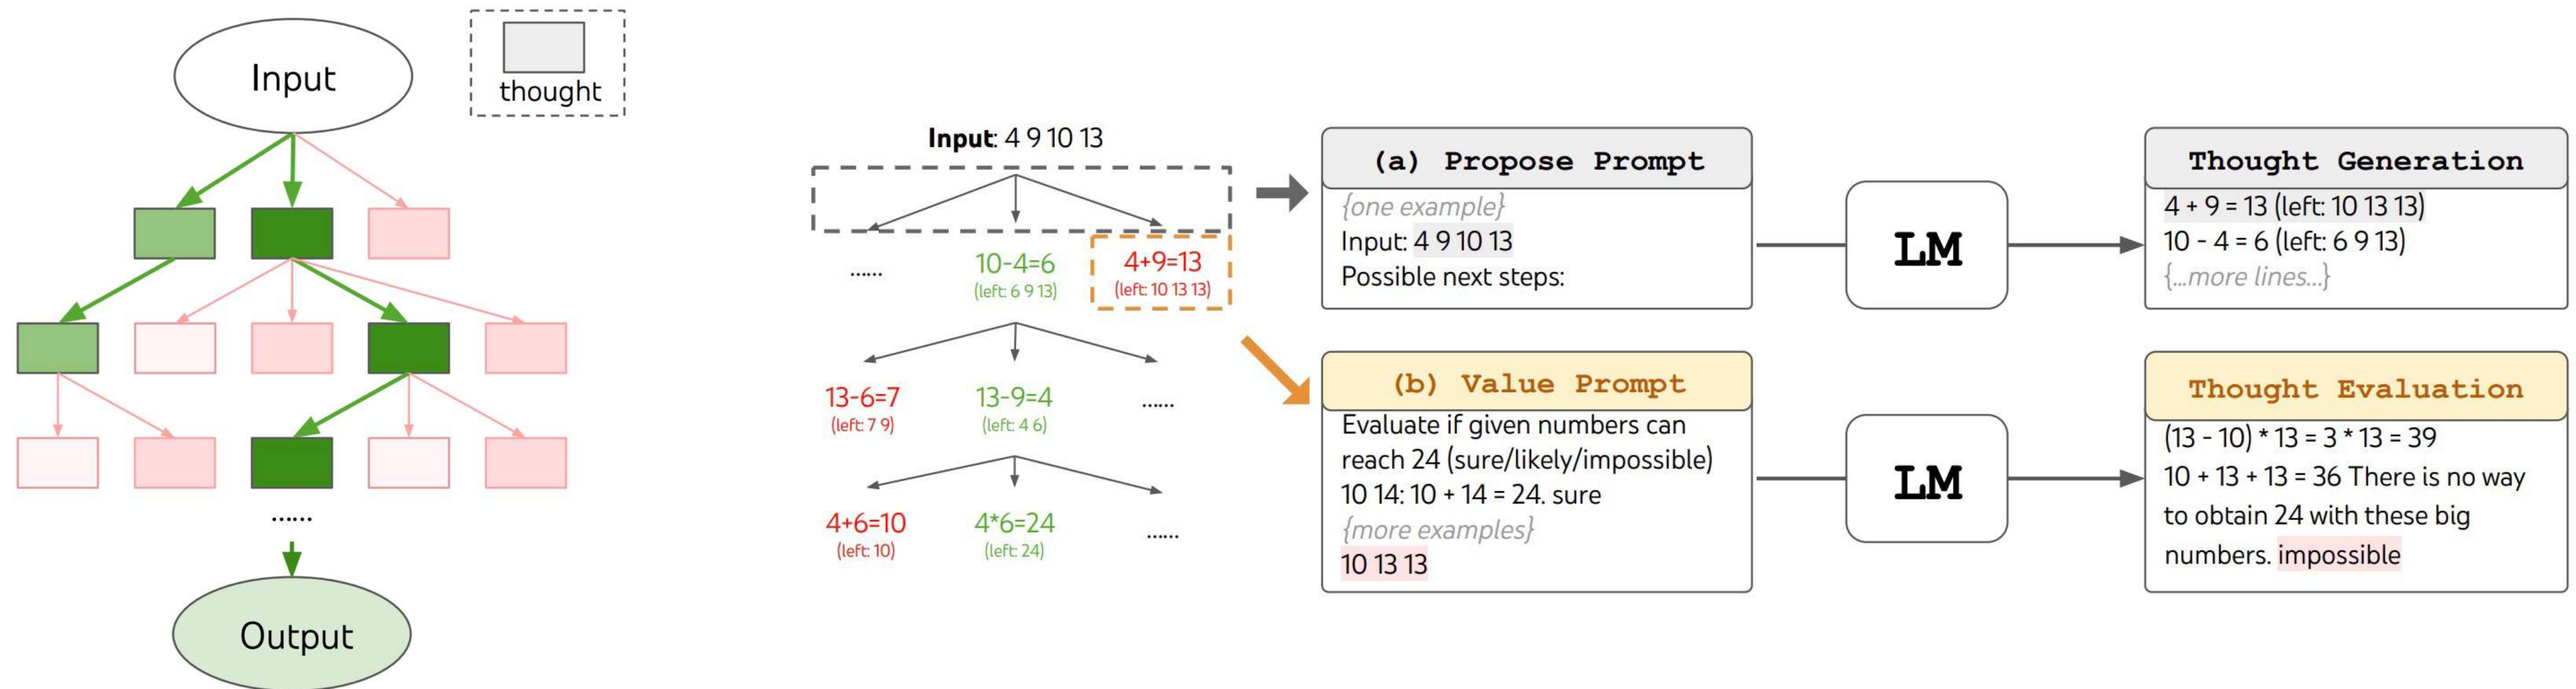
Self-consistency

- Since LLMs are stochastic, sample responses **multiple times** and perform **majority vote** may be useful



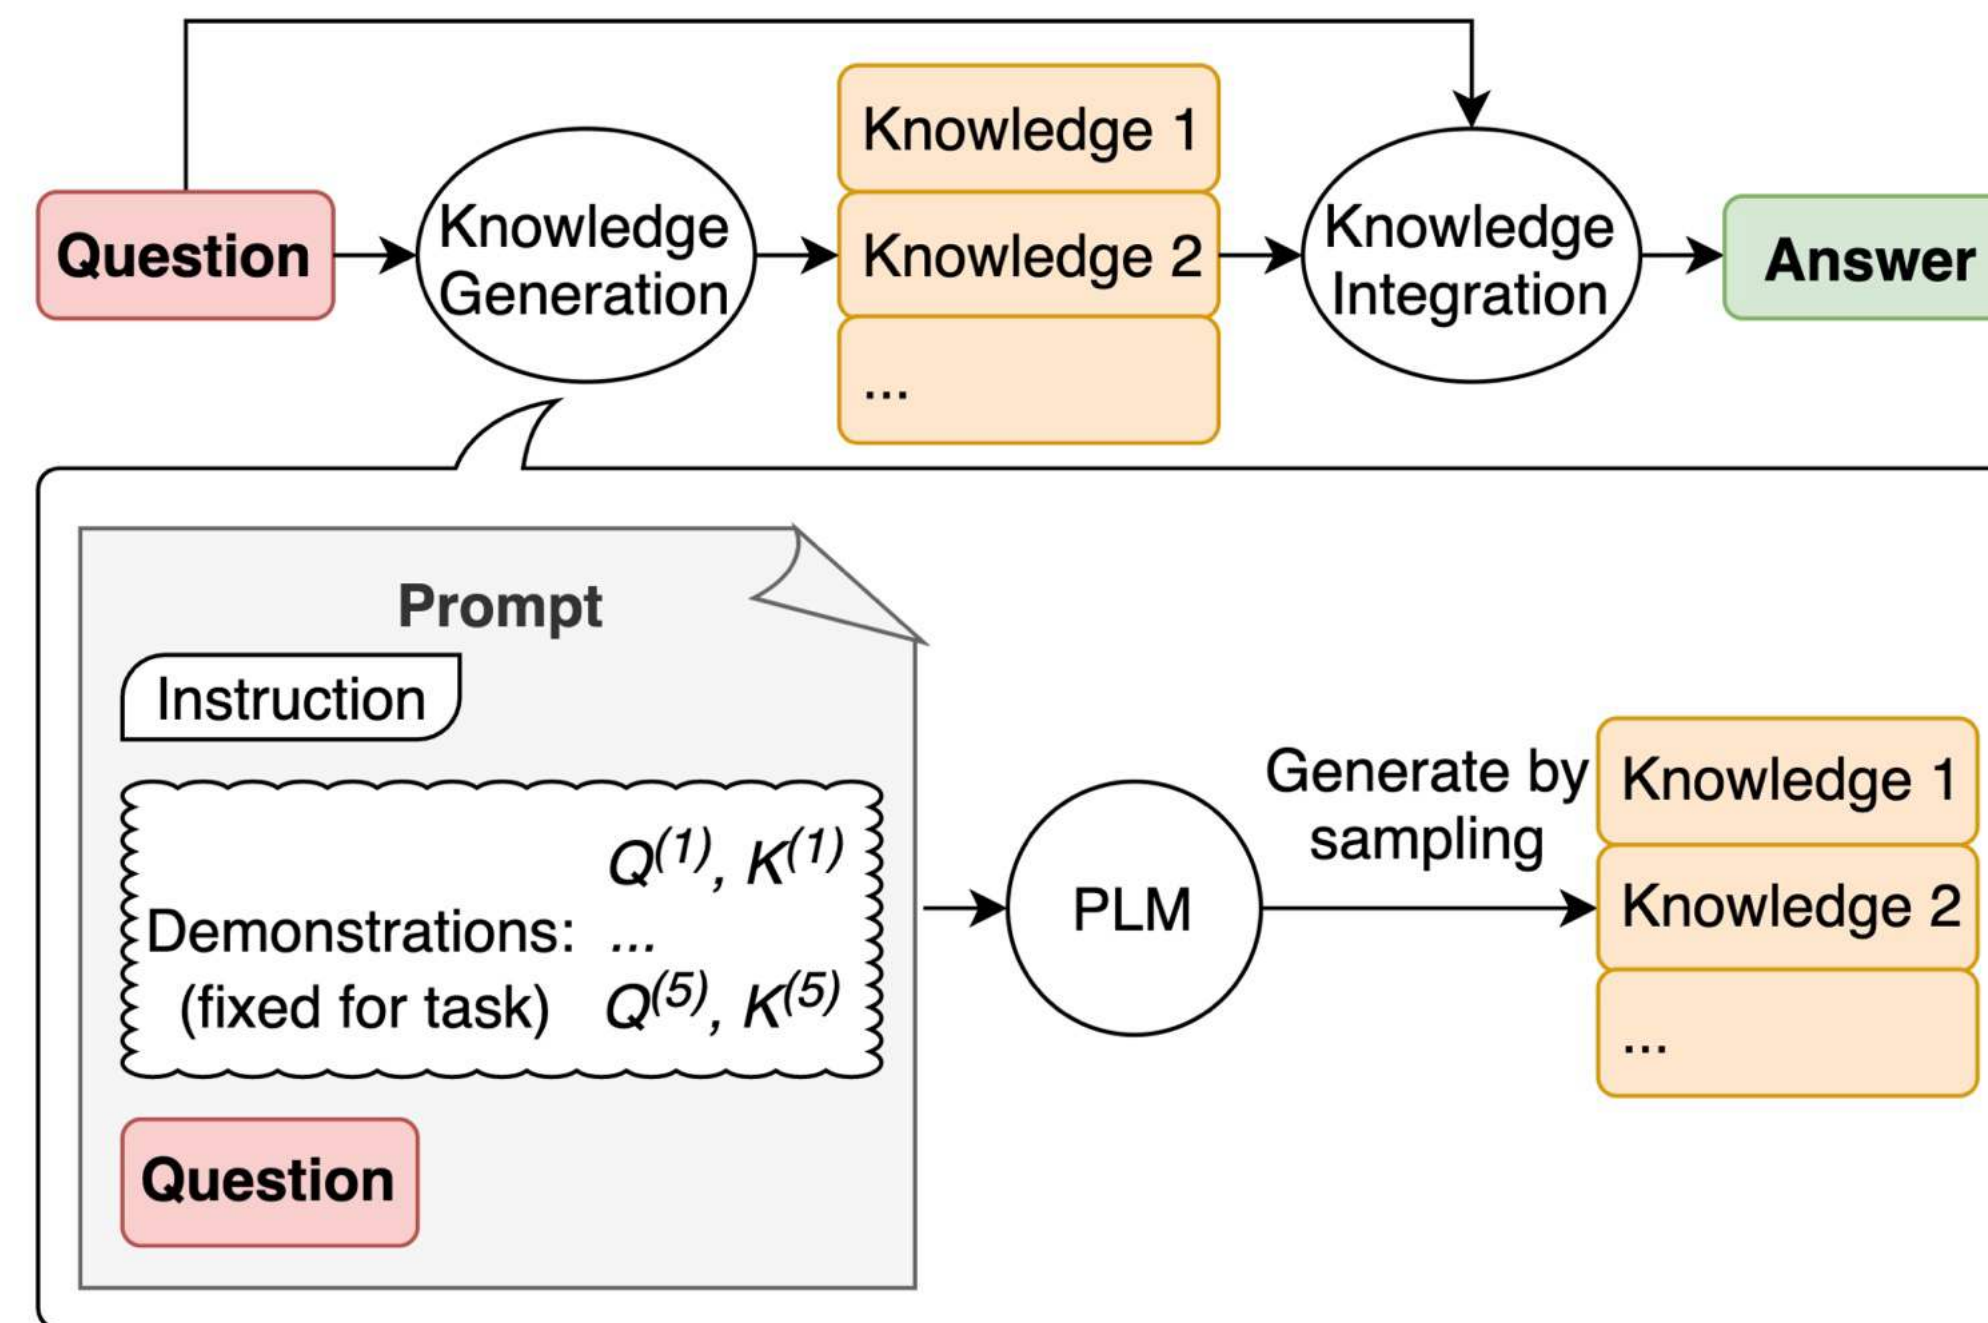
Tree of Thoughts (ToT) Prompting

- Since CoT's reasoning steps may take a **diverse path** due to the stochastic nature of LLMs, exploring **many strategies** might be useful



Generated Knowledge Prompting

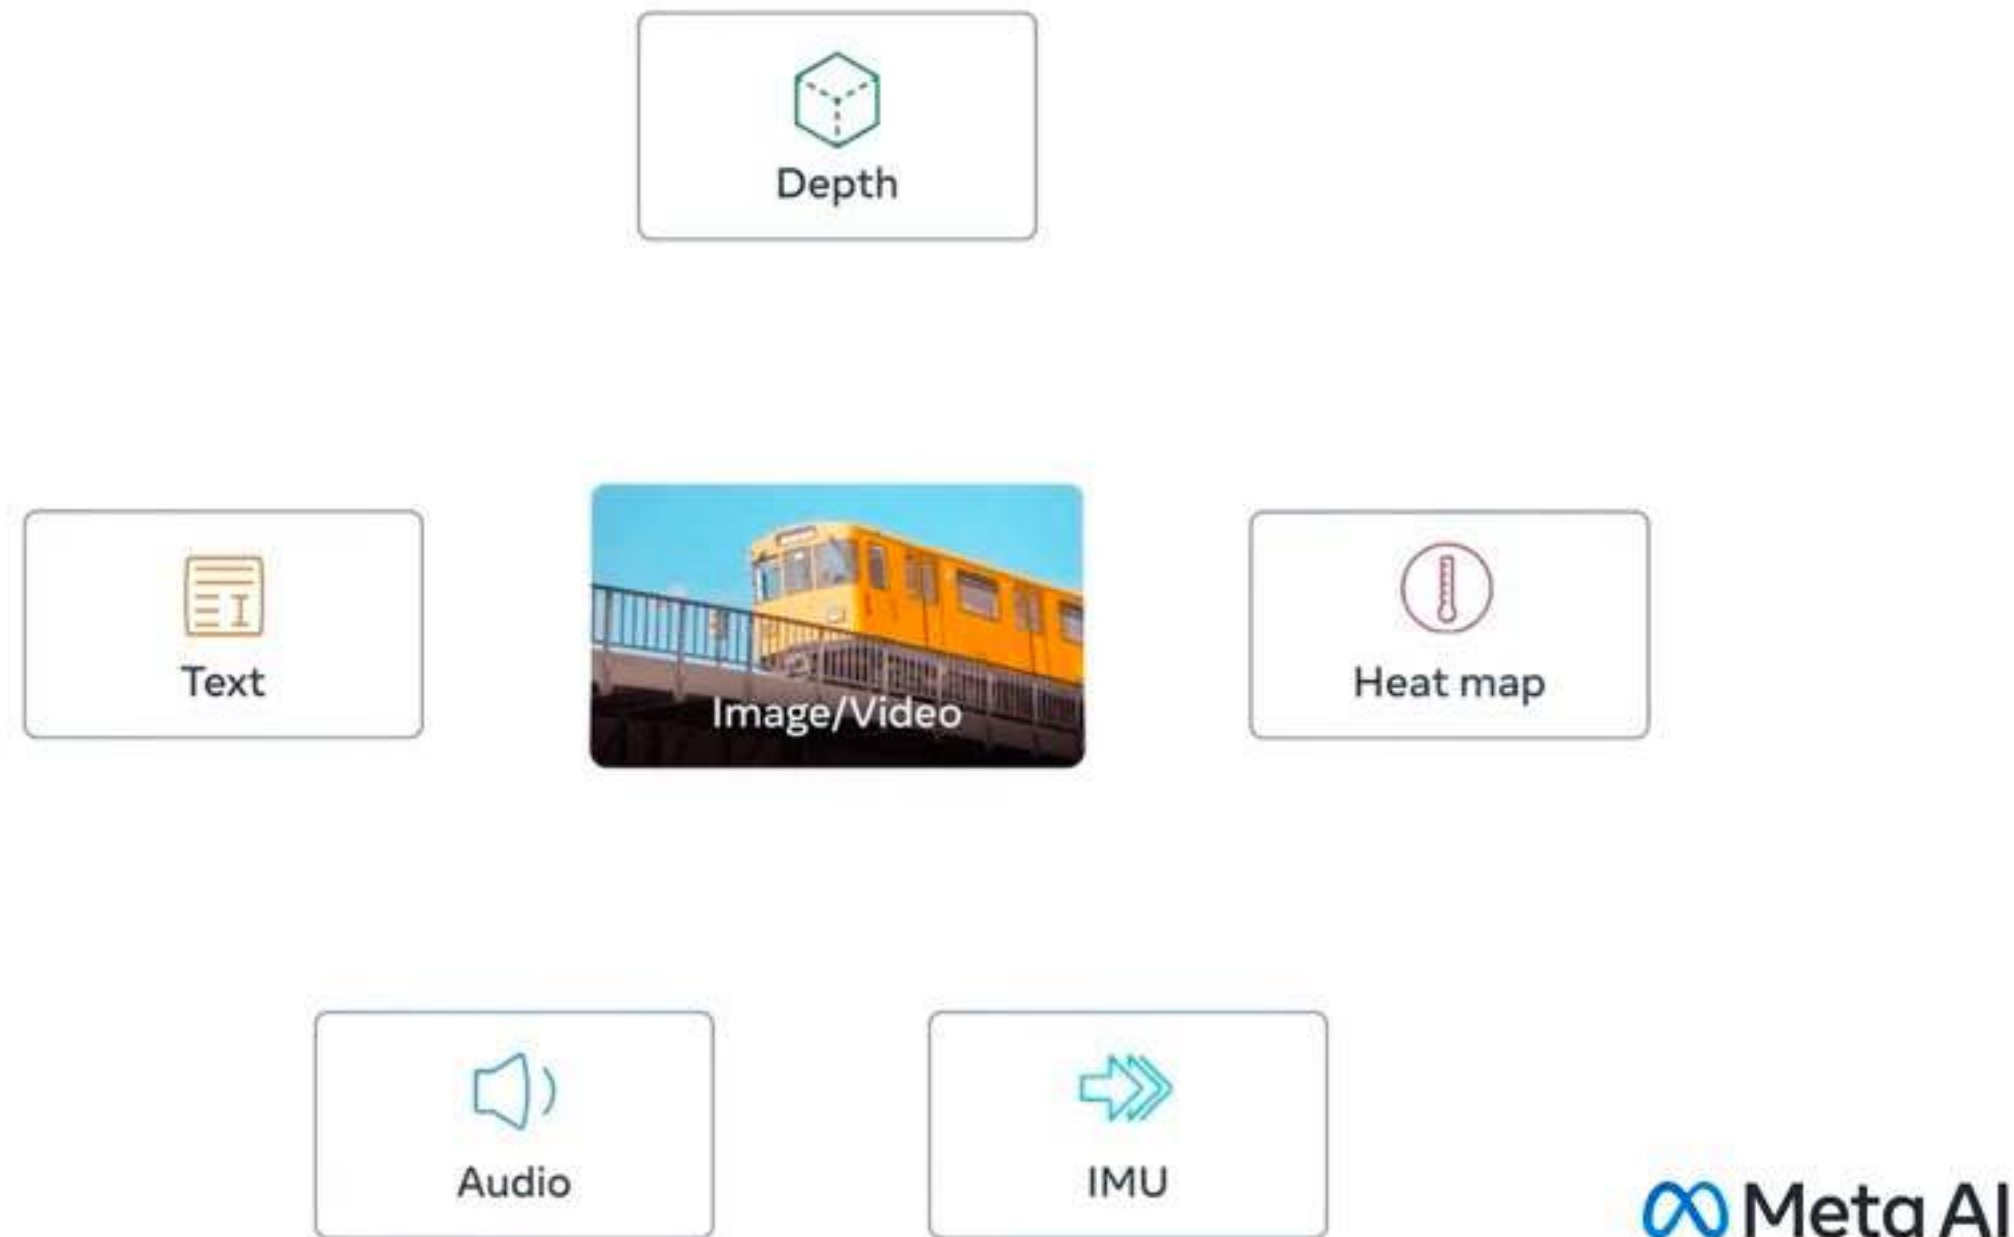
- Ask a model to **generate new knowledge** based on facts, e.g., summary, and use generated facts to improve performance



Closing Remarks

A Glimpse into the Future

- Multi-modal LLMs, e.g., PandaGPT



A recording of thunder and pouring rain.

Write a story to a child about what you see and hear.

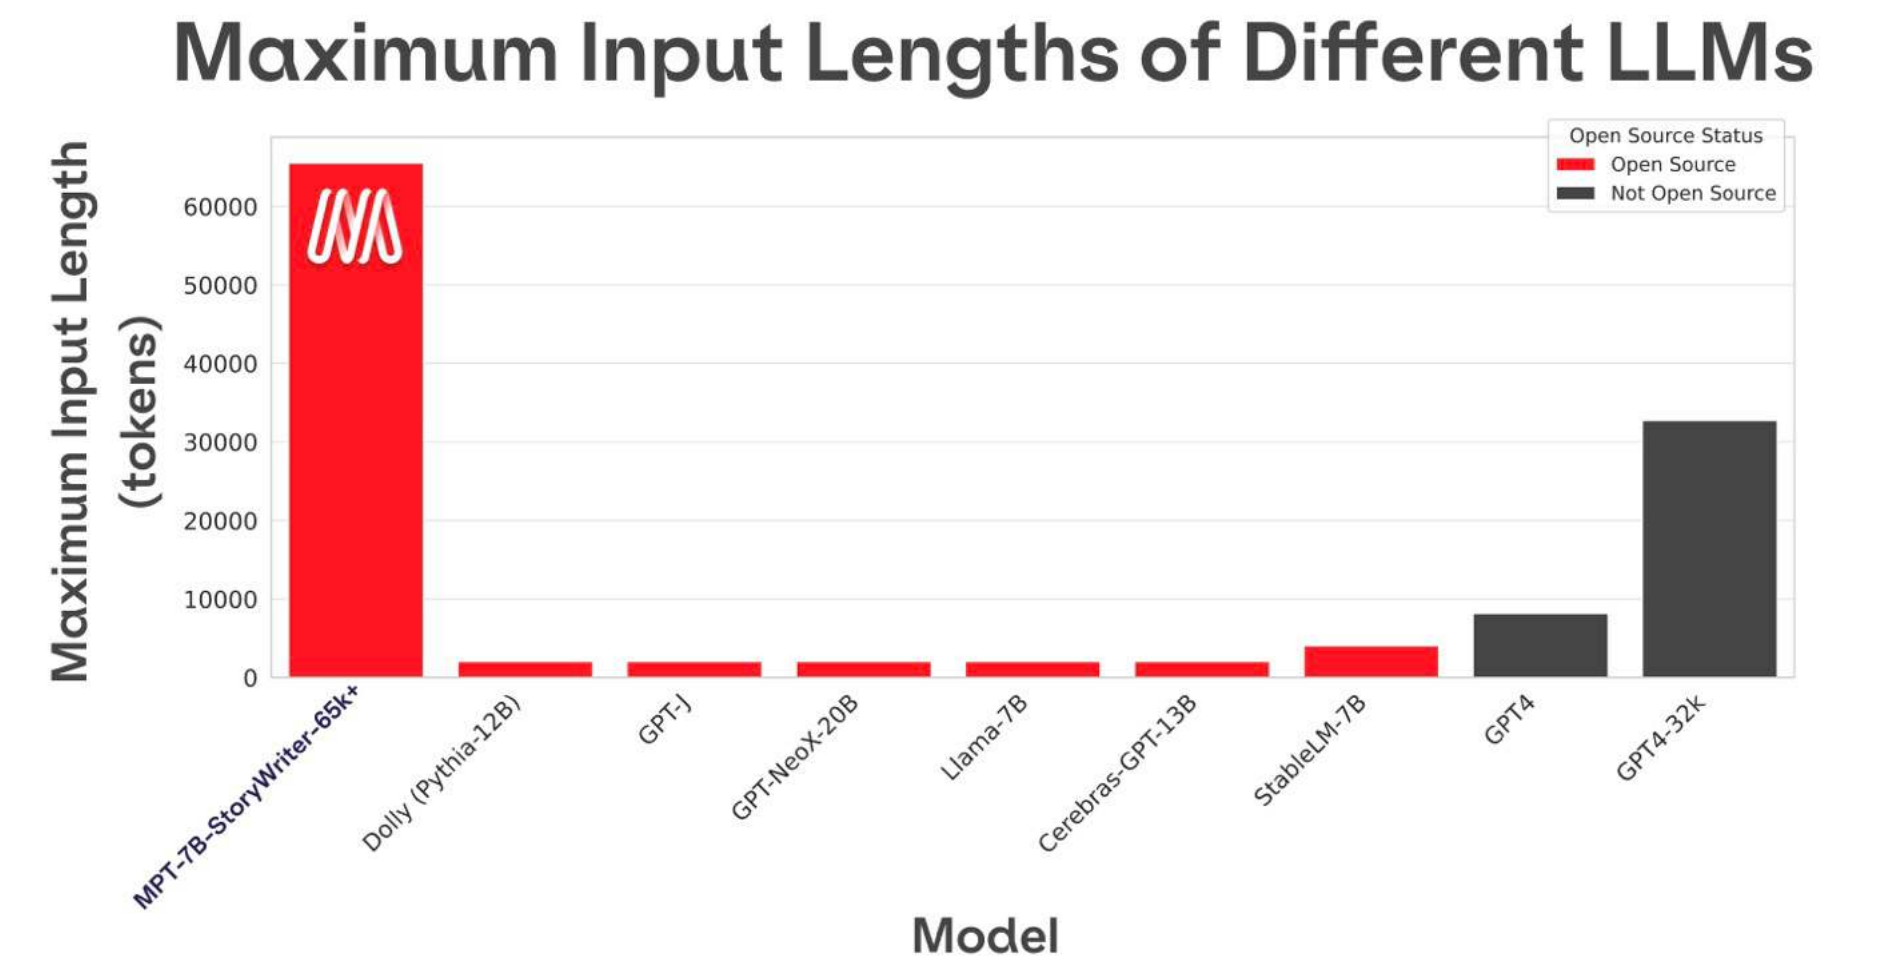
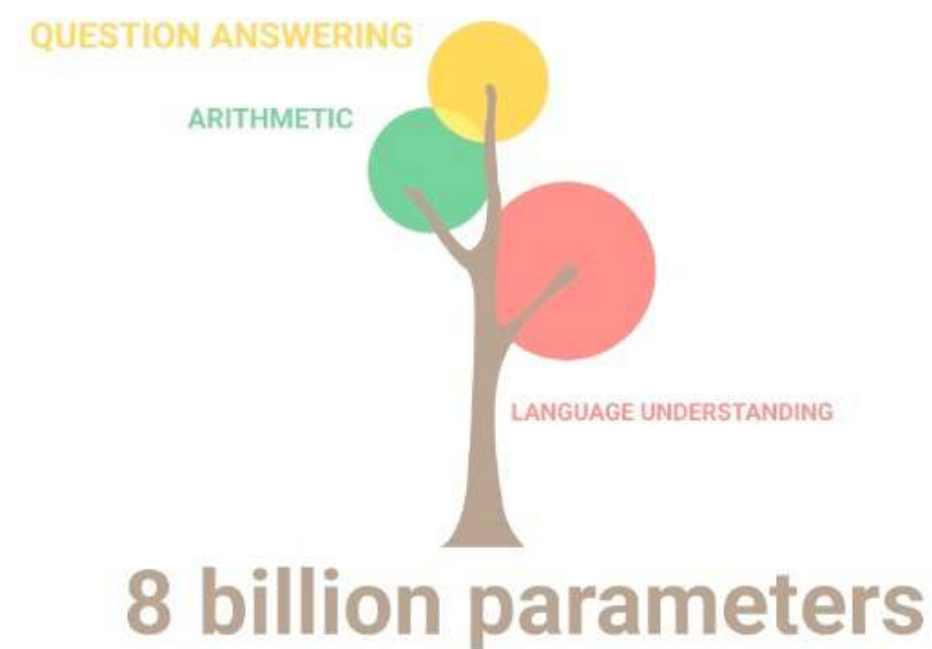


As I look out the window, I see a beautiful scene of a person walking on the road in the distance. The person is walking under a cloudy sky, which is dark and stormy. The sound of the wind is blowing, and I can hear the distant sound of thunder rumbling in the background. The person is walking with a backpack, and it seems like they are on a journey or adventure. The stormy weather might make the journey a bit more challenging, but the person seems determined to continue their journey despite the conditions.

Meta AI

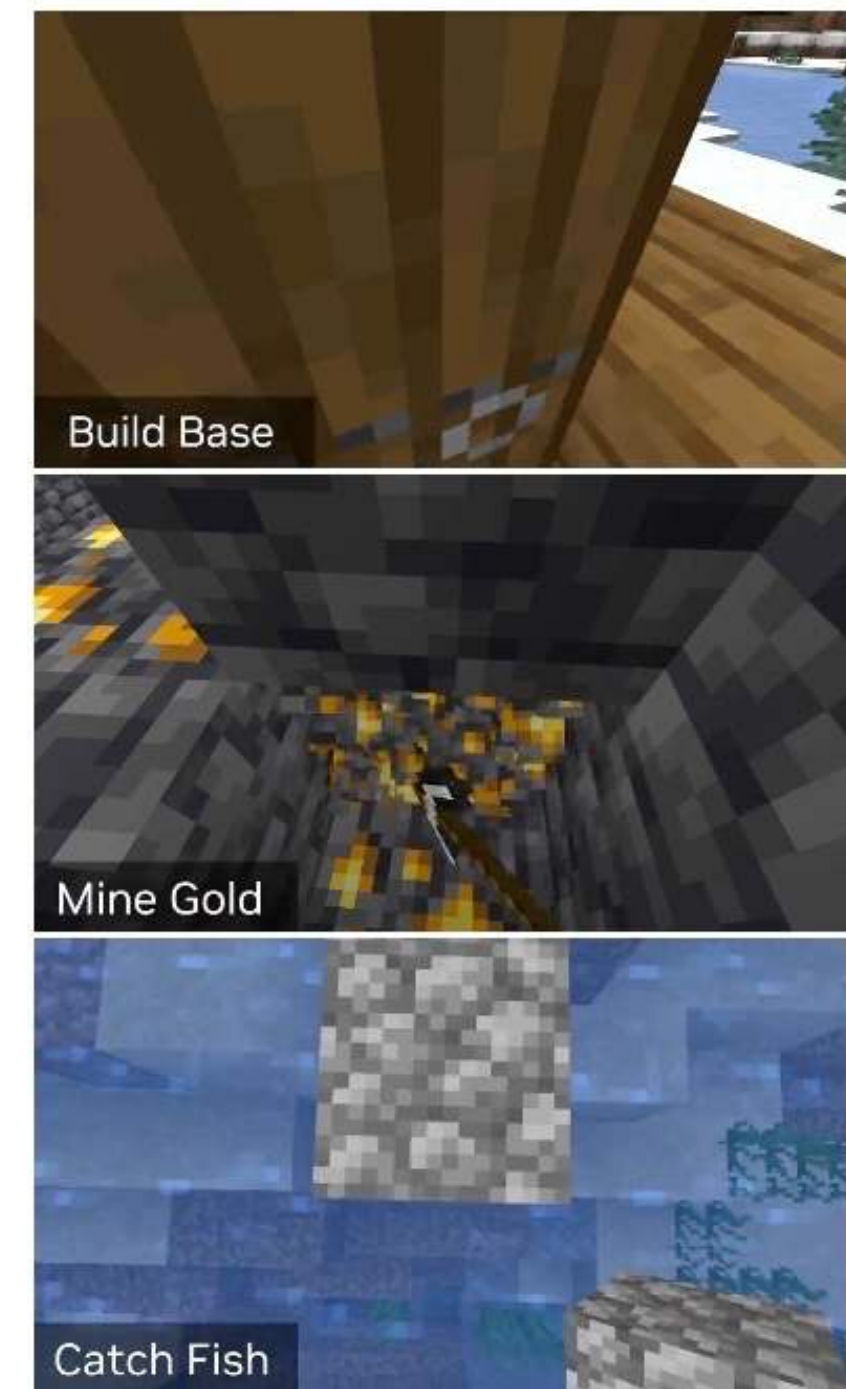
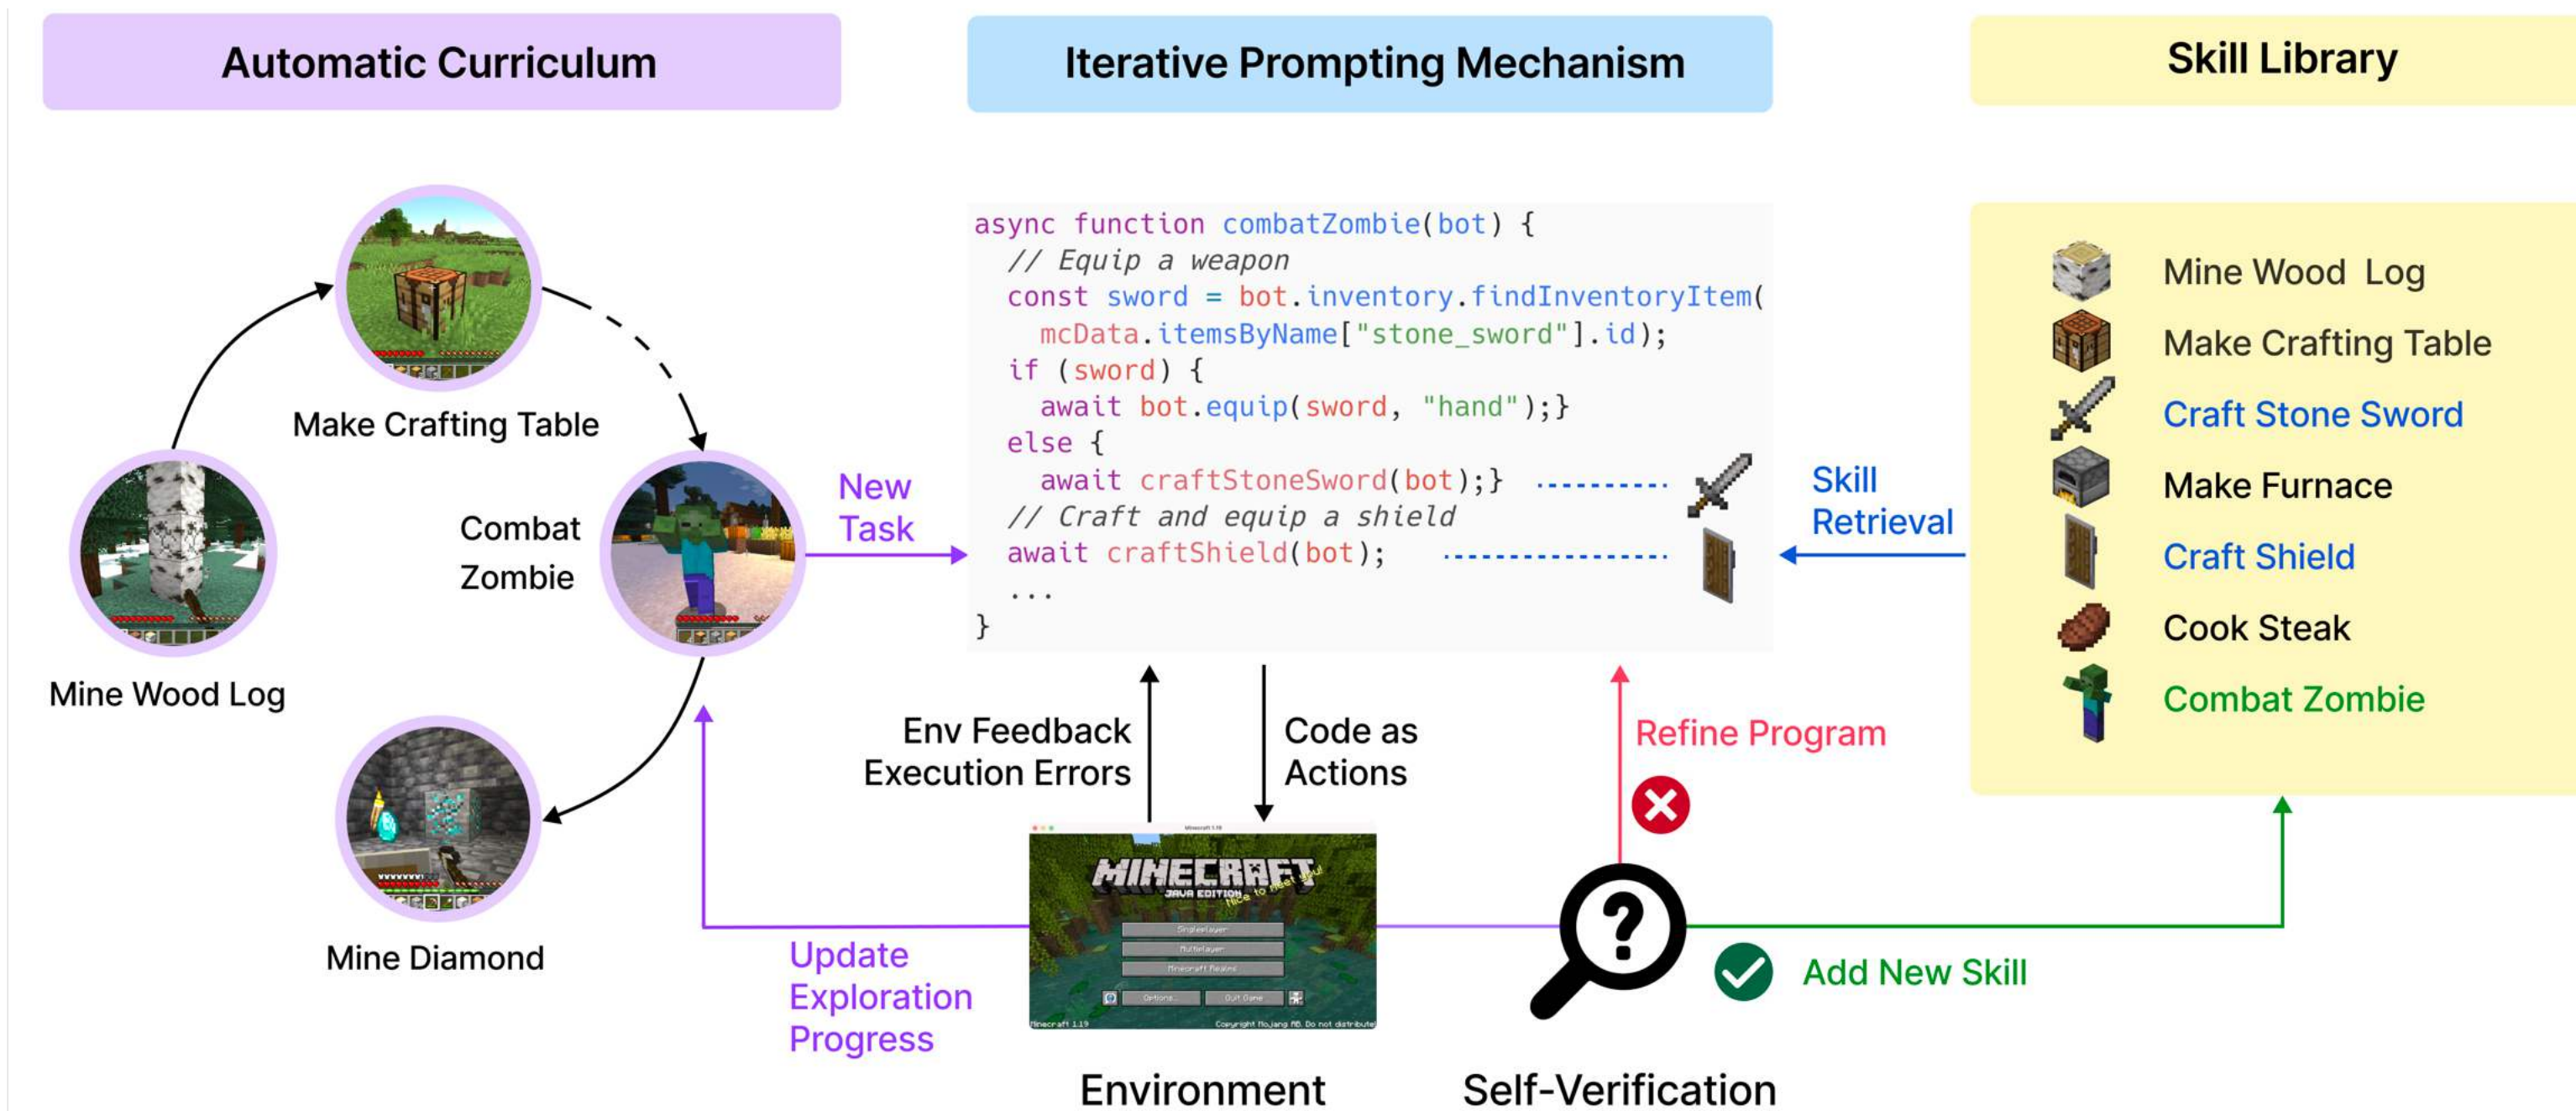
A Glimpse into the Future

- Larger context and larger LLMs, e.g., PaLM **540B**, MPT-7B-StoryWriter-**65k+**



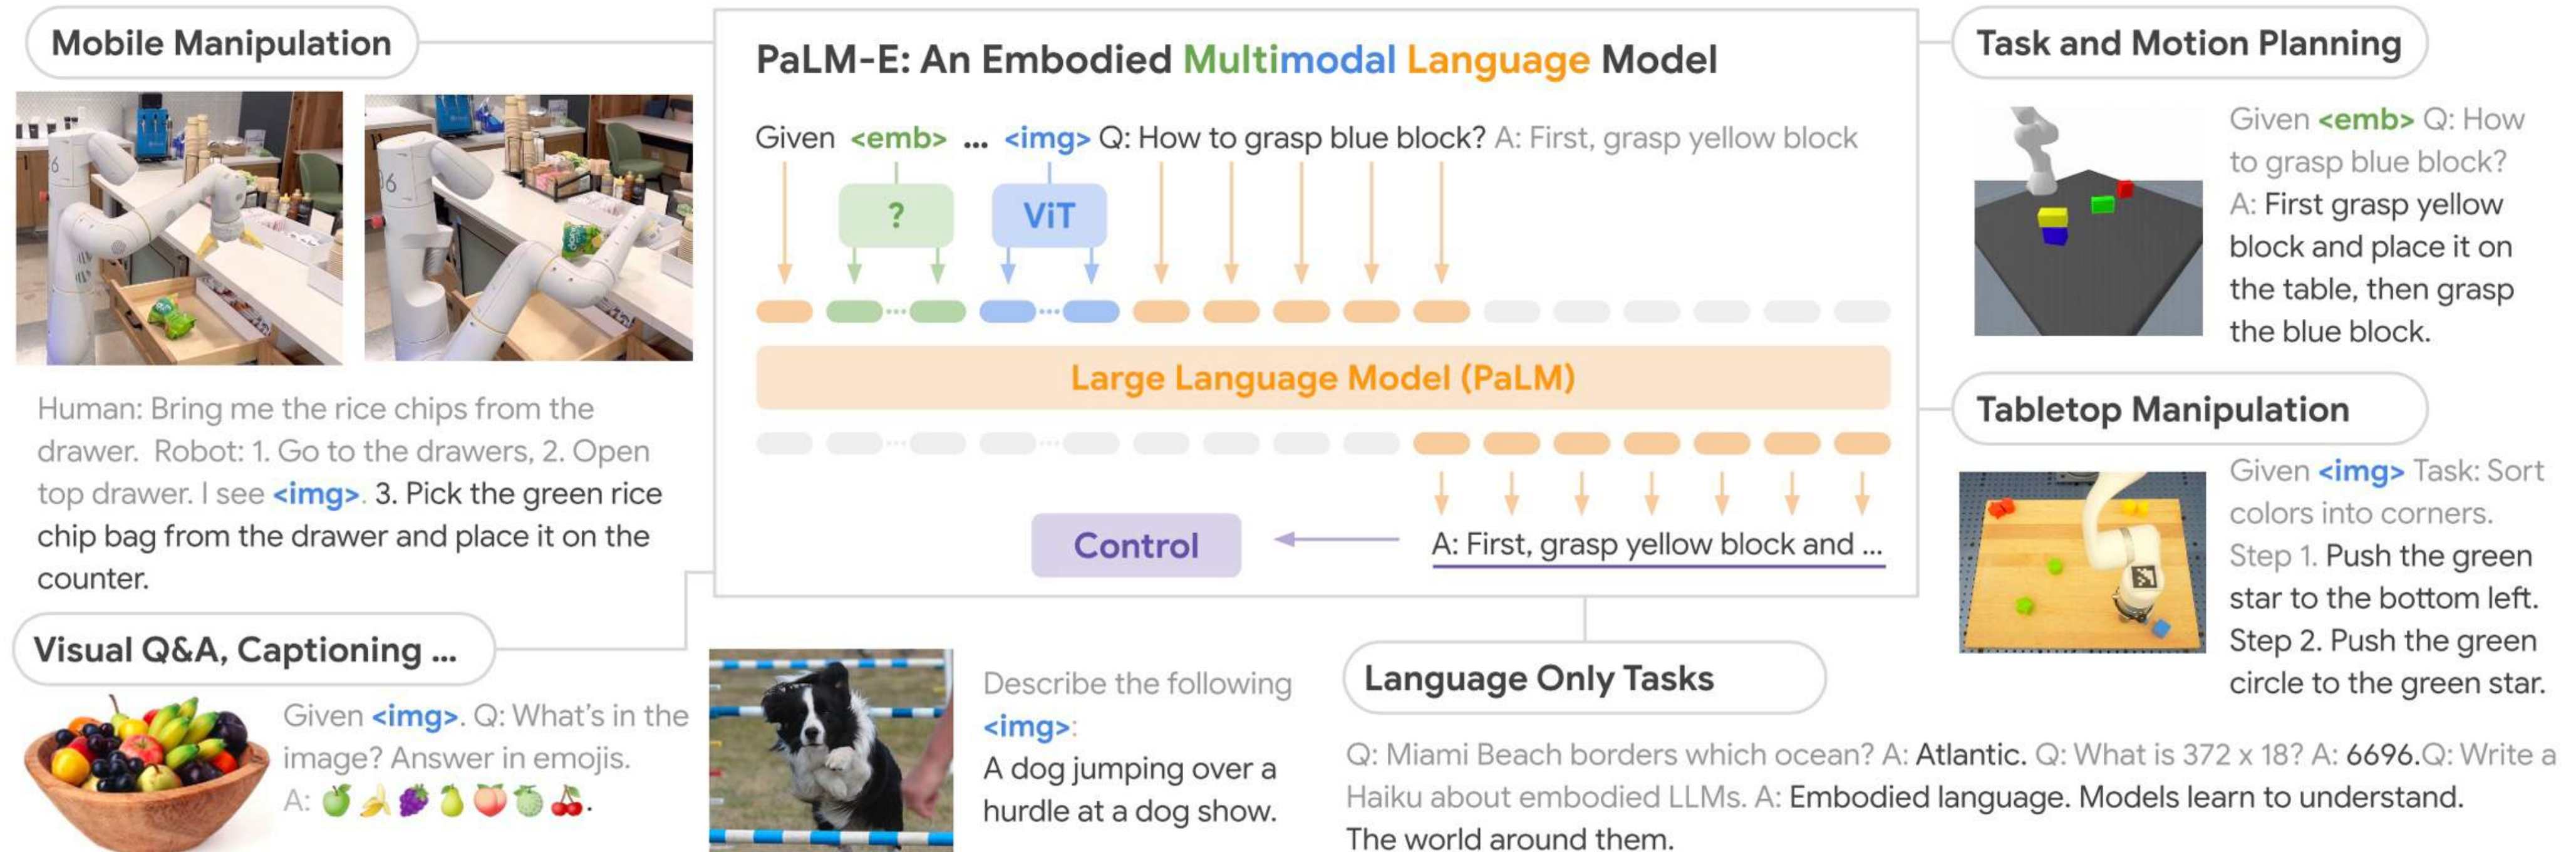
A Glimpse into the Future

- LLMs as a **tool maker/user**, e.g., Voyager, LATM



A Glimpse into the Future

- Embodied LLMs, LLMs that can affect real world, e.g., PaLM-E



Need a Place to Practice PE Skills?



chatgpt4pcg.github.io



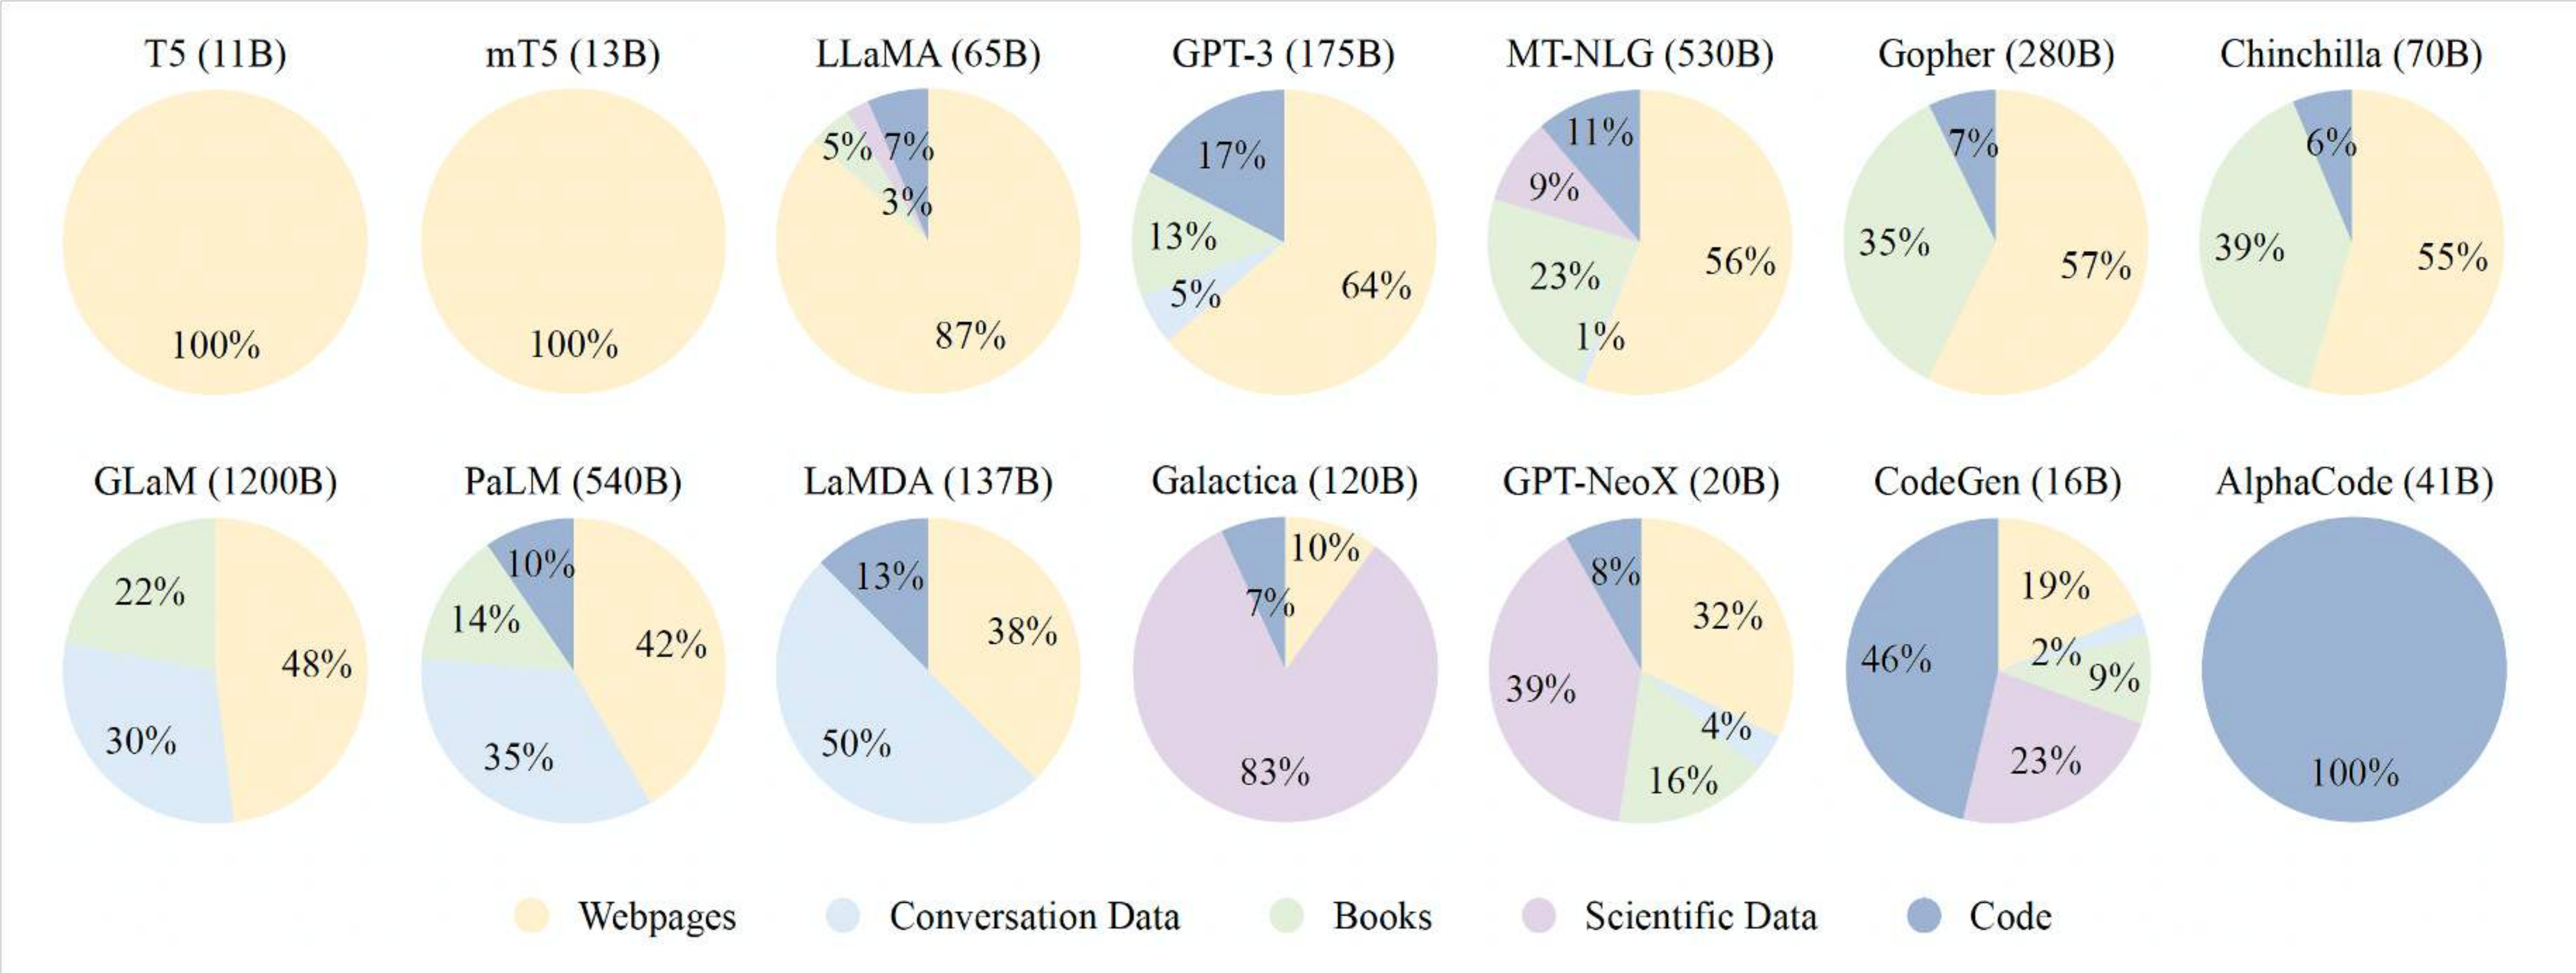


Coding Time!

Bonus: How to Choose LLMs?

- Based on these factors
 - **Context length:** Its memory
 - **Training dataset:** Its prior knowledge
 - **Model size:** Its capabilities and hardware requirements
 - **License:** Research-only or business-ok

Bonus: How to Choose LLMs?



Bonus: More Topics to Explore

- Prompt patterns (SE for prompt)
- Adversarial Prompting
 - Prompt injection
 - Prompt leaking
 - Jailbreaking
- Soft prompts
 - Interpretable soft prompts
- AI-generated content detection
- Prompt engineering for diffusion-based/GAN-based models
 - Negative prompts
 - Non-text input as a prompt
- LLM: Grounding, Alignment, Distillation