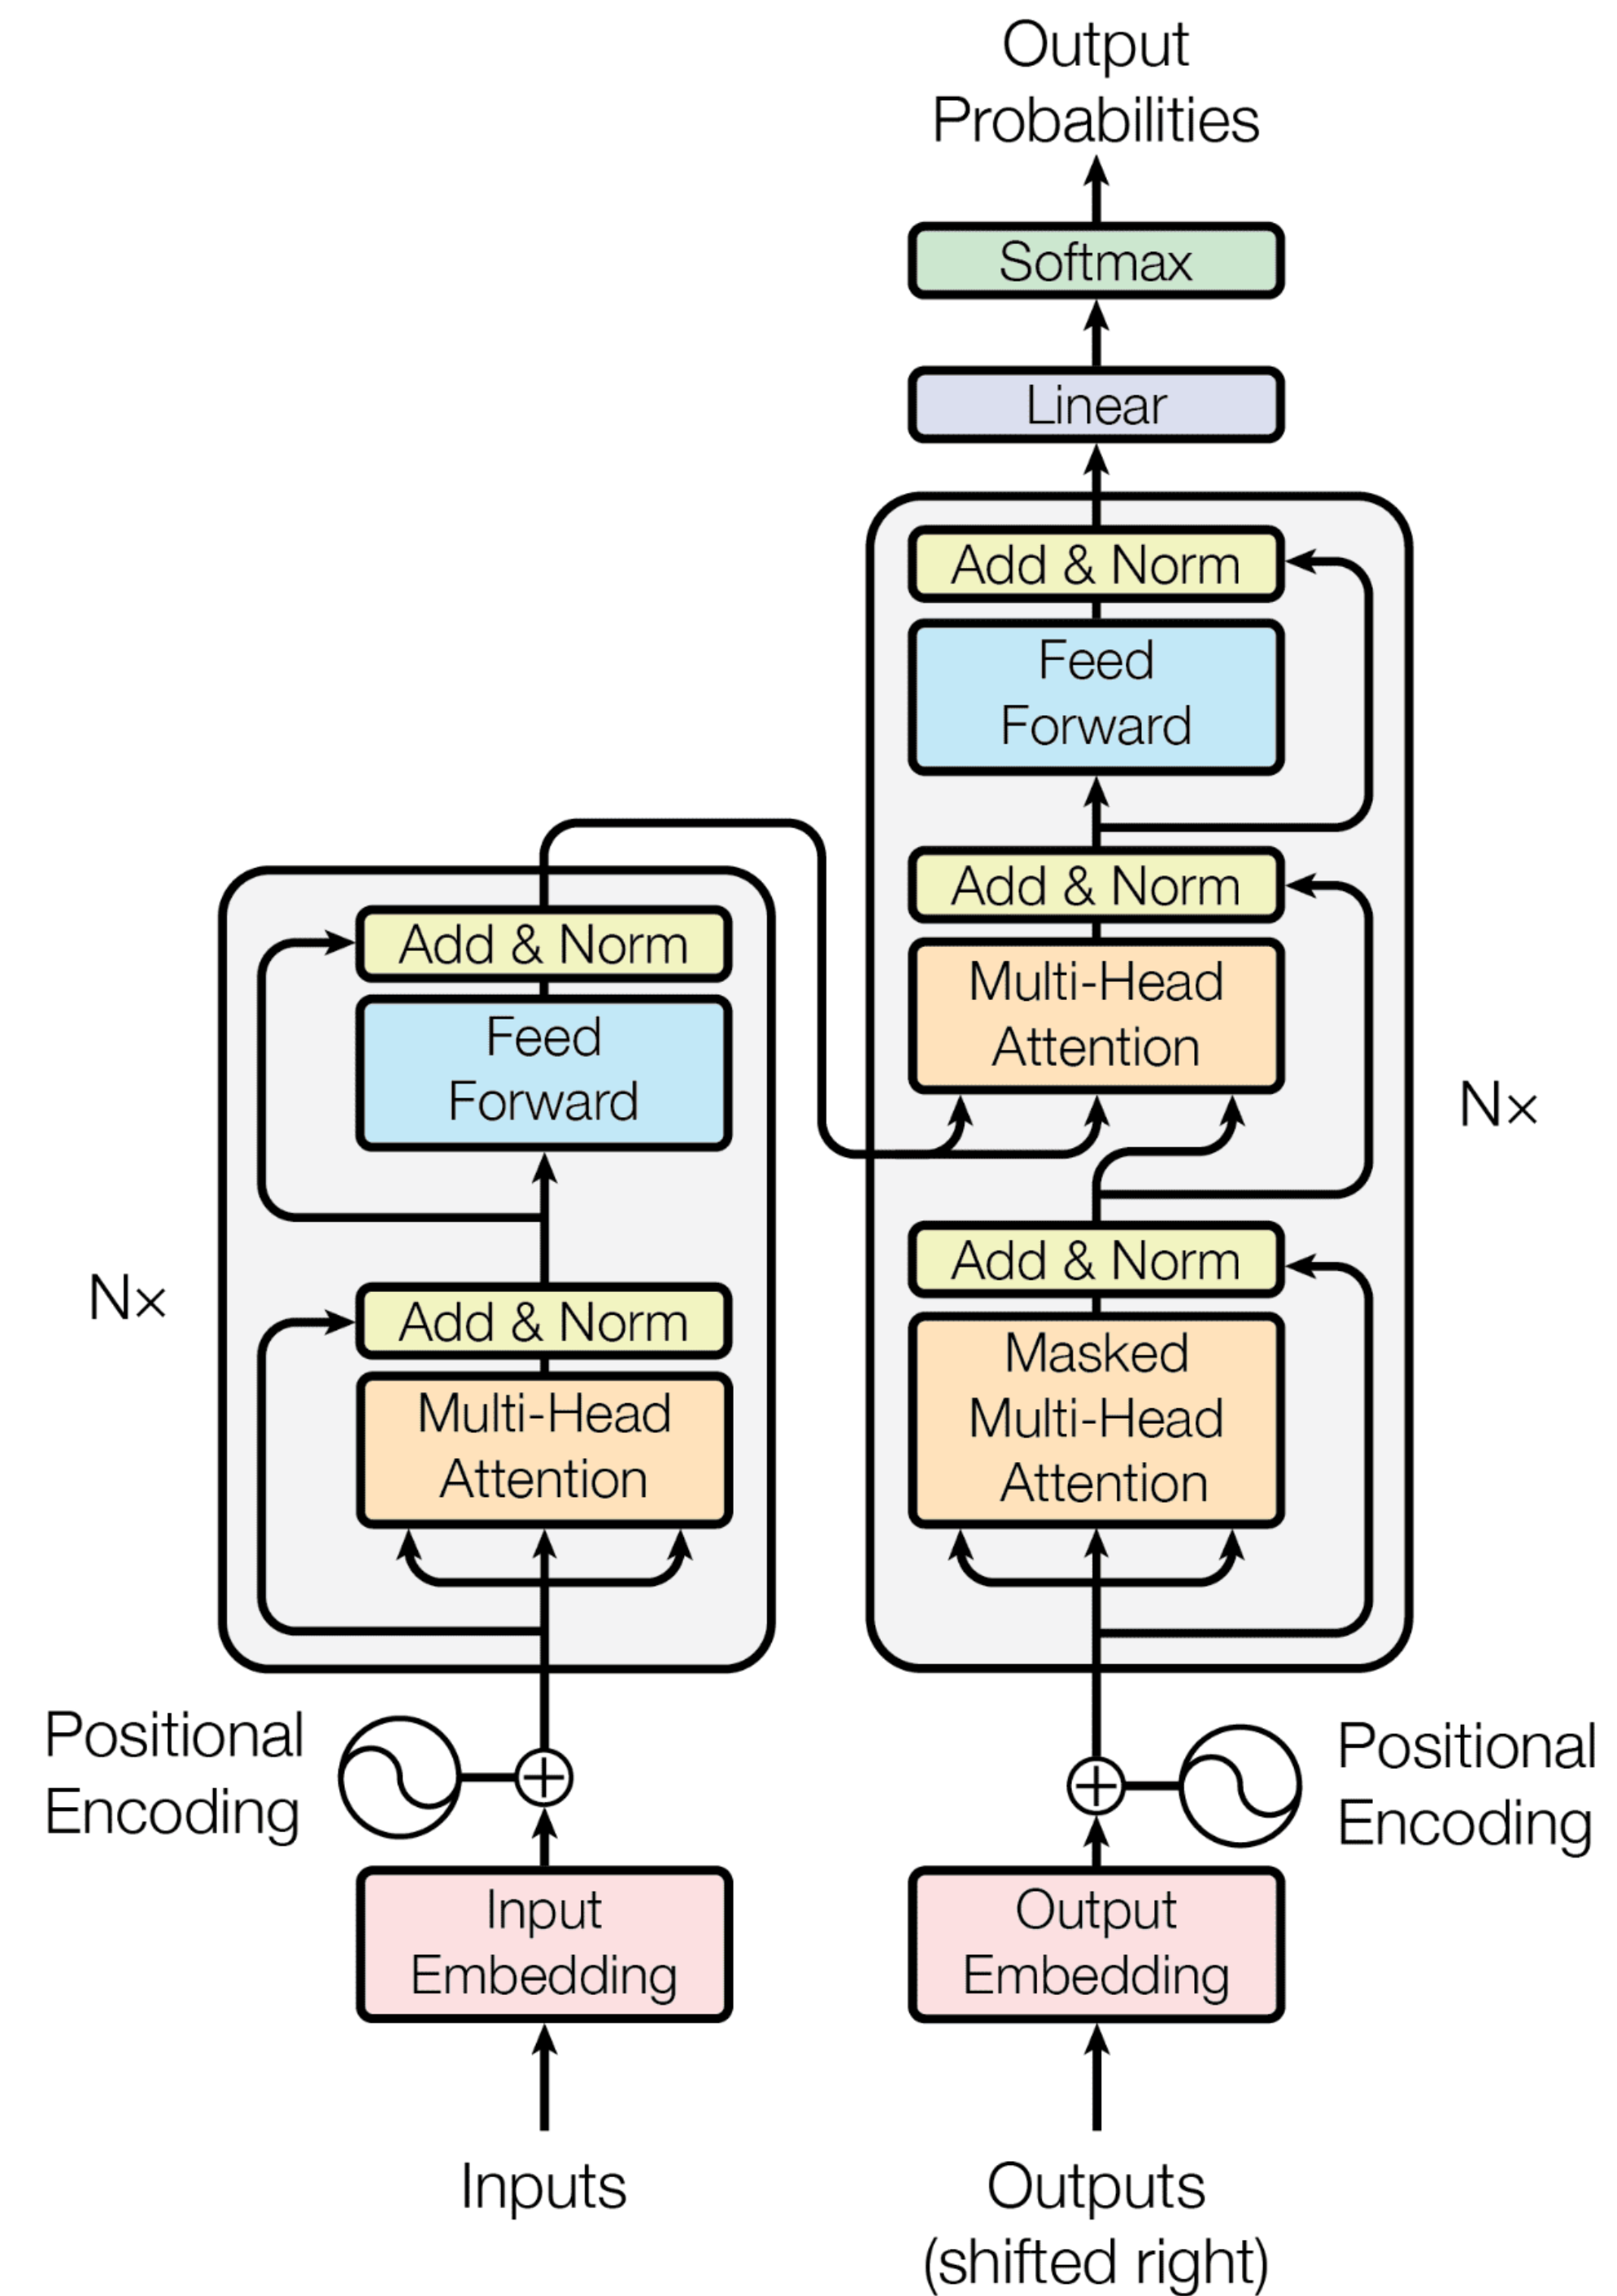


22. Transformers & LLMs (cont'd)

**EECE454 Introduction to
Machine Learning Systems**

Recap

- **Transformers** handle sequence of inputs
 - Tokenization
 - Embedding
 - Transformer blocks
 - Multi-Head Attention
 - Feed-Forward Networks
 - LayerNorm & Residuals



Today

- **How to train** the transformers
 - We will see how **GPTs** evolved, from GPT-1 (2018) to GPT-4 (now)
(caution: not much “theory”)



ChatGPT

Transformers as Universal Backbones

Almost all community uses transformers now...

- **Language.** Transformers (2017)
- **Vision.** Vision Transformer (2020)
- **Speech.** Conformers (2020) + vanilla transformers
- **RL.** Decision Transformers (2021)
+ Time-Series, Graph, Tabular, ...

Note. An important step toward multi-modal AI.

Scalability




Transformers are also known to be much more ***scalable!***

Two things:

- **Parallel Training.** Much easier than RNNs
- **Making a Big ver.** Very easy.
 - Simply increase the channel dim. / num. heads / layers

Scaling Law

We can easily reduce the *test loss of transformer language models*, simply by increasing compute, dataset, and model.



Data Crisis?

Easy. Increasing the compute.

Easy. Increasing the model size.

Hard. Increasing the dataset size

Built-in workflow with Amazon Mechanical Turk

If you use [Amazon Mechanical Turk](#) for labeling, you are charged per object per review instance. We recommend that you use multiple labelers per object to improve label accuracy.

Workflow	Suggested price per labeler
Image classification	\$0.012
Text classification	\$0.012
Named Entity Recognition (NER)	\$0.024
Bounding box	\$0.036
Semantic Segmentation	\$0.84

Data Crisis?

Easy. Increasing the compute.

Easy. Increasing the model size.

Hard. Increasing the dataset size

➔ **Easy.** Increasing the dataset size, **without supervision.**

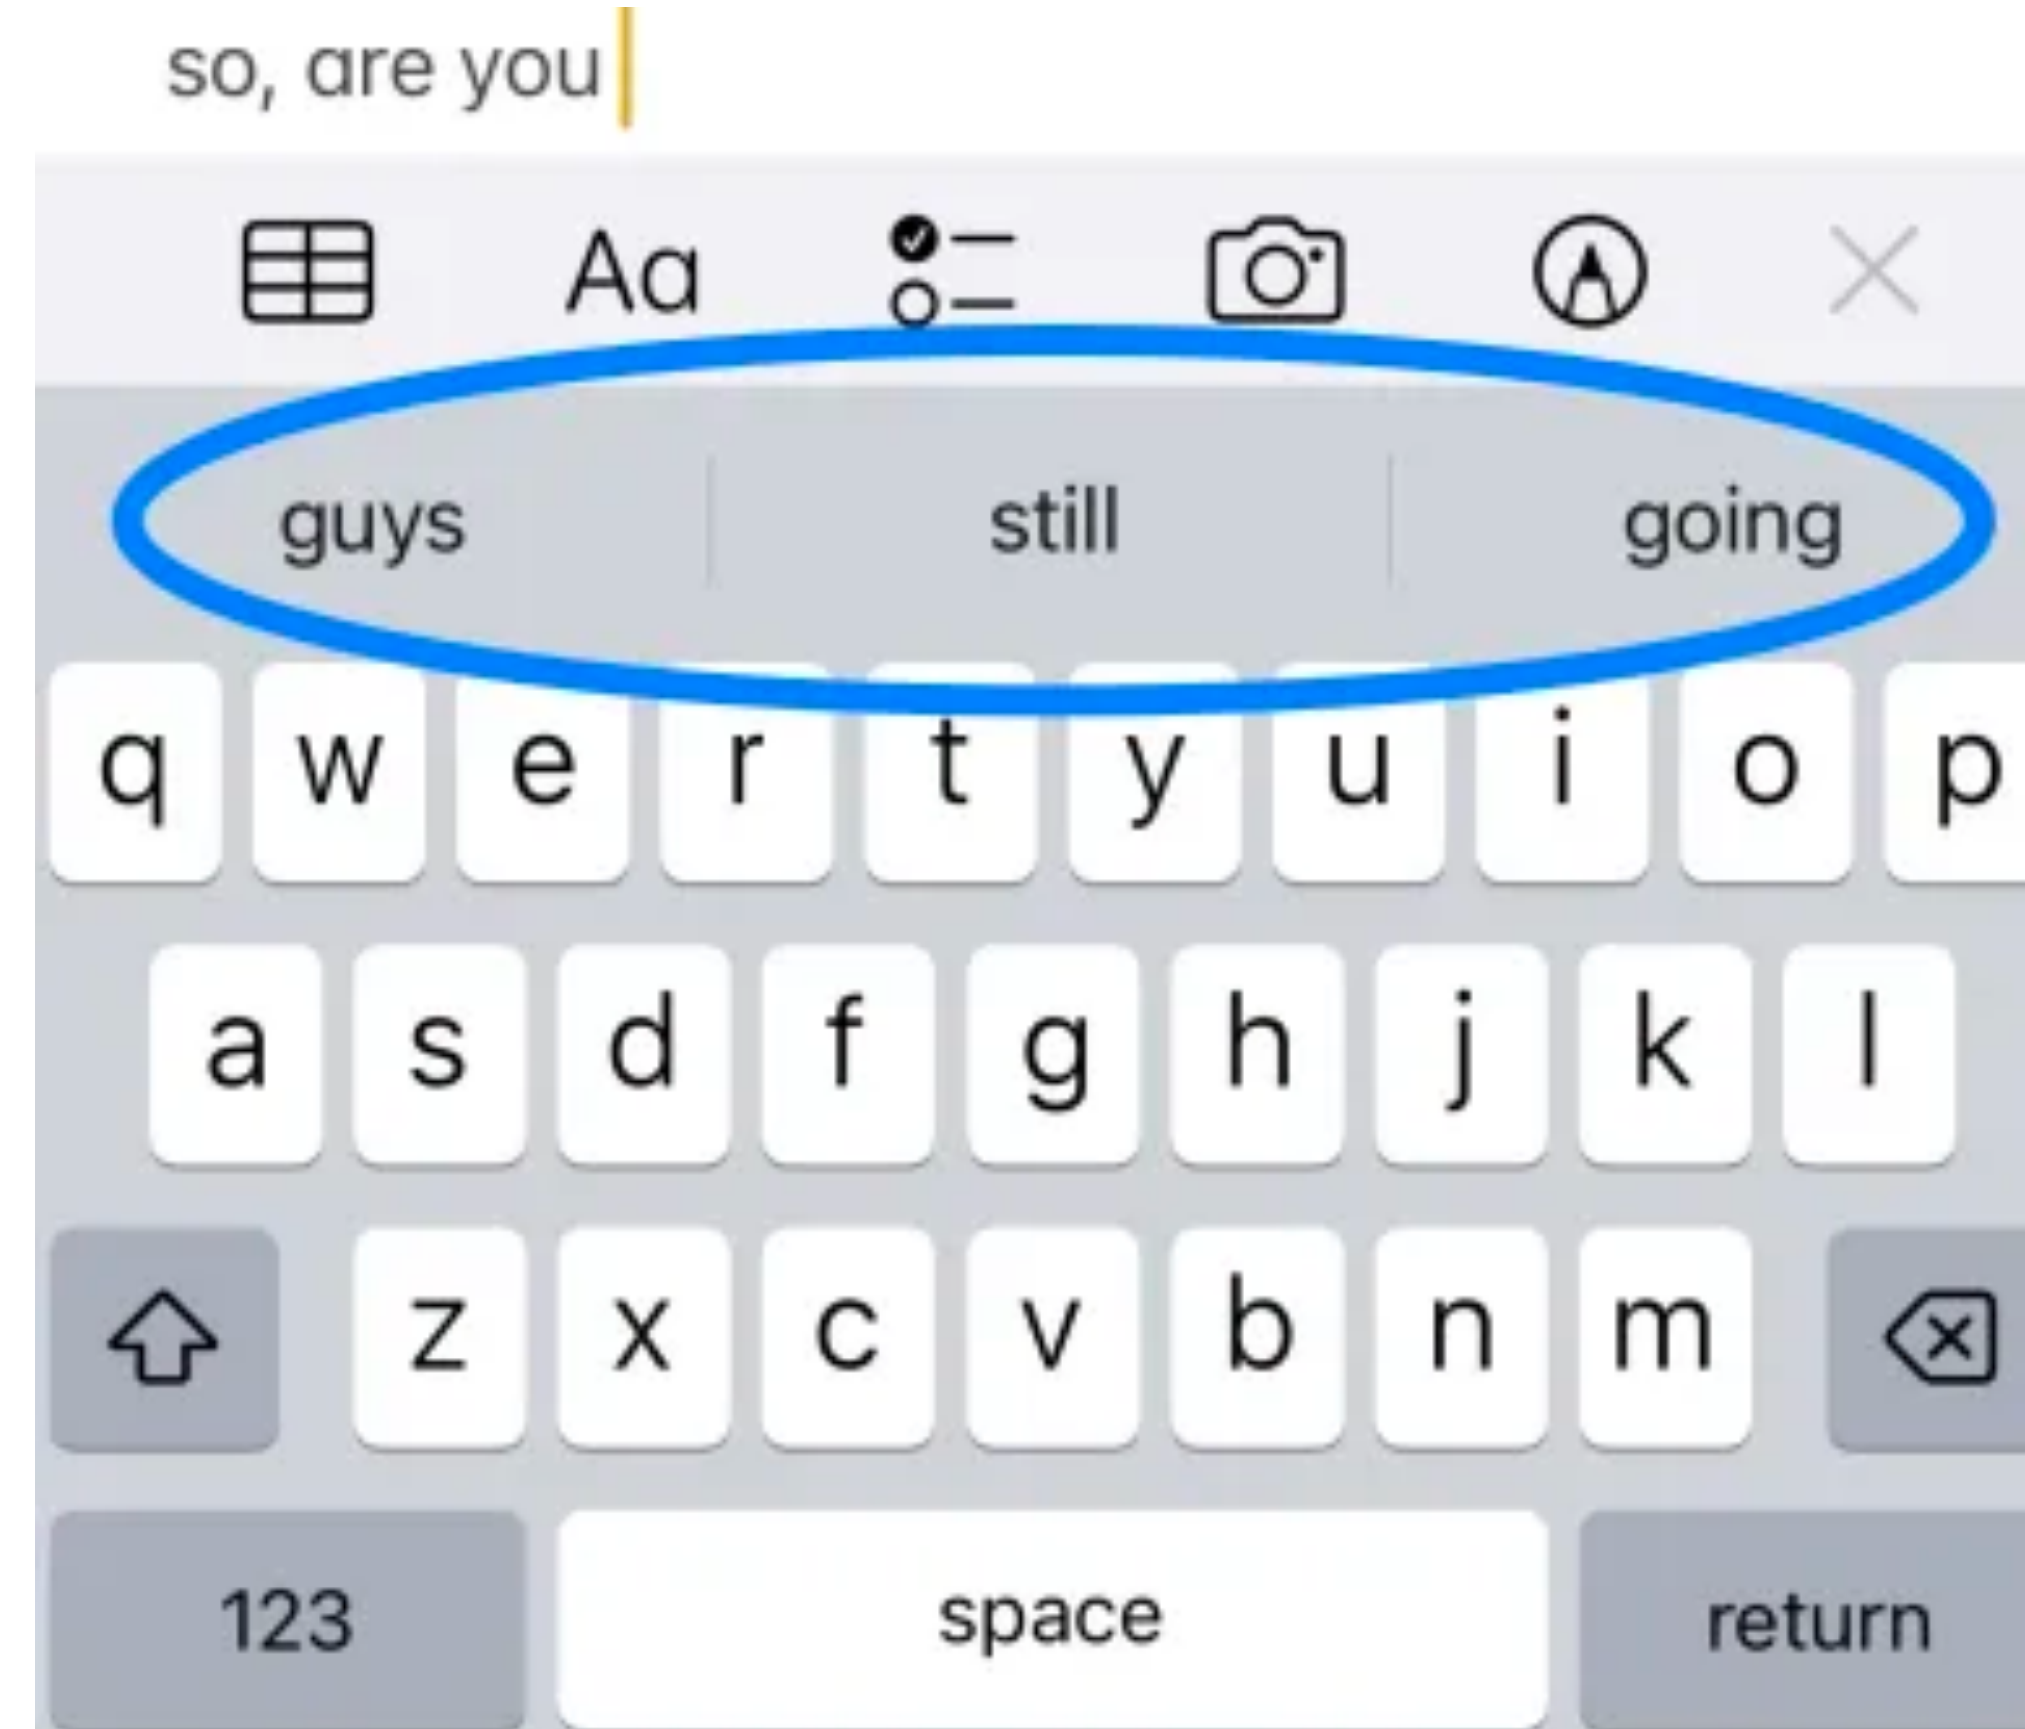
Many unsupervised pretraining techniques...

We focus on ***GPT-style***

Generative Pretraining (a.k.a. autoregressive training)

Generative Pretraining

- **Idea.** We have a lot of unlabeled sentences on Web.
Train a model that can do *next-word prediction*.



Generative Pretraining

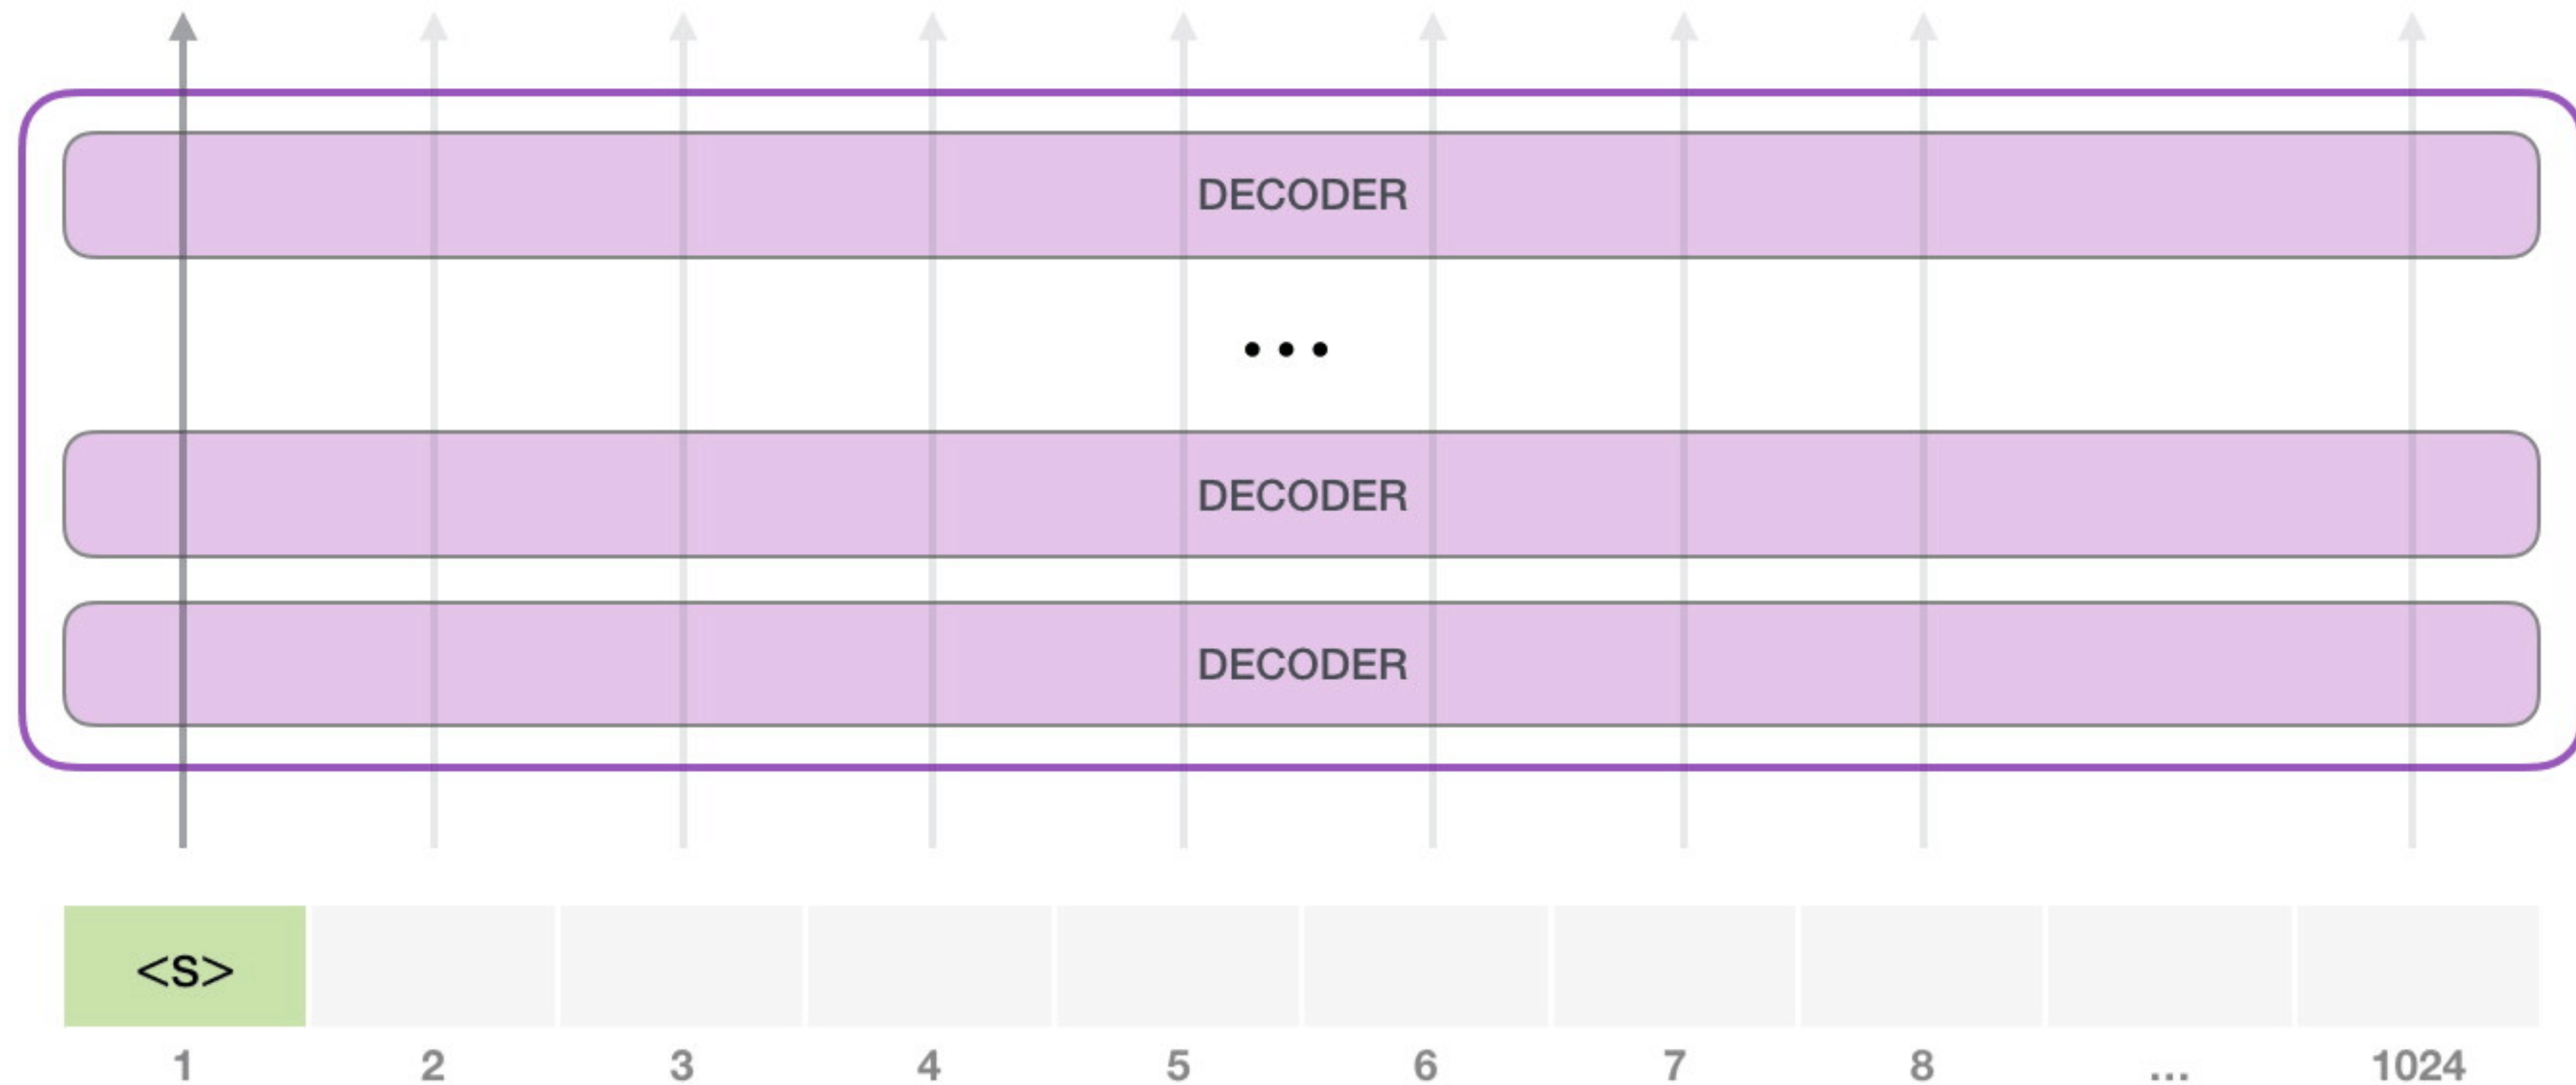
- **Idea.** We have a lot of unlabeled sentences on Web.
Train a model that can do *next-word prediction*.
- That is, we find a generative model $p_{\theta}(\cdot)$ that maximizes the likelihood

$$L(\theta) = \sum_i \log p_{\theta}(\mathbf{x}_i \mid \mathbf{x}_{i-k}, \dots, \mathbf{x}_{i-1})$$

| Context Length

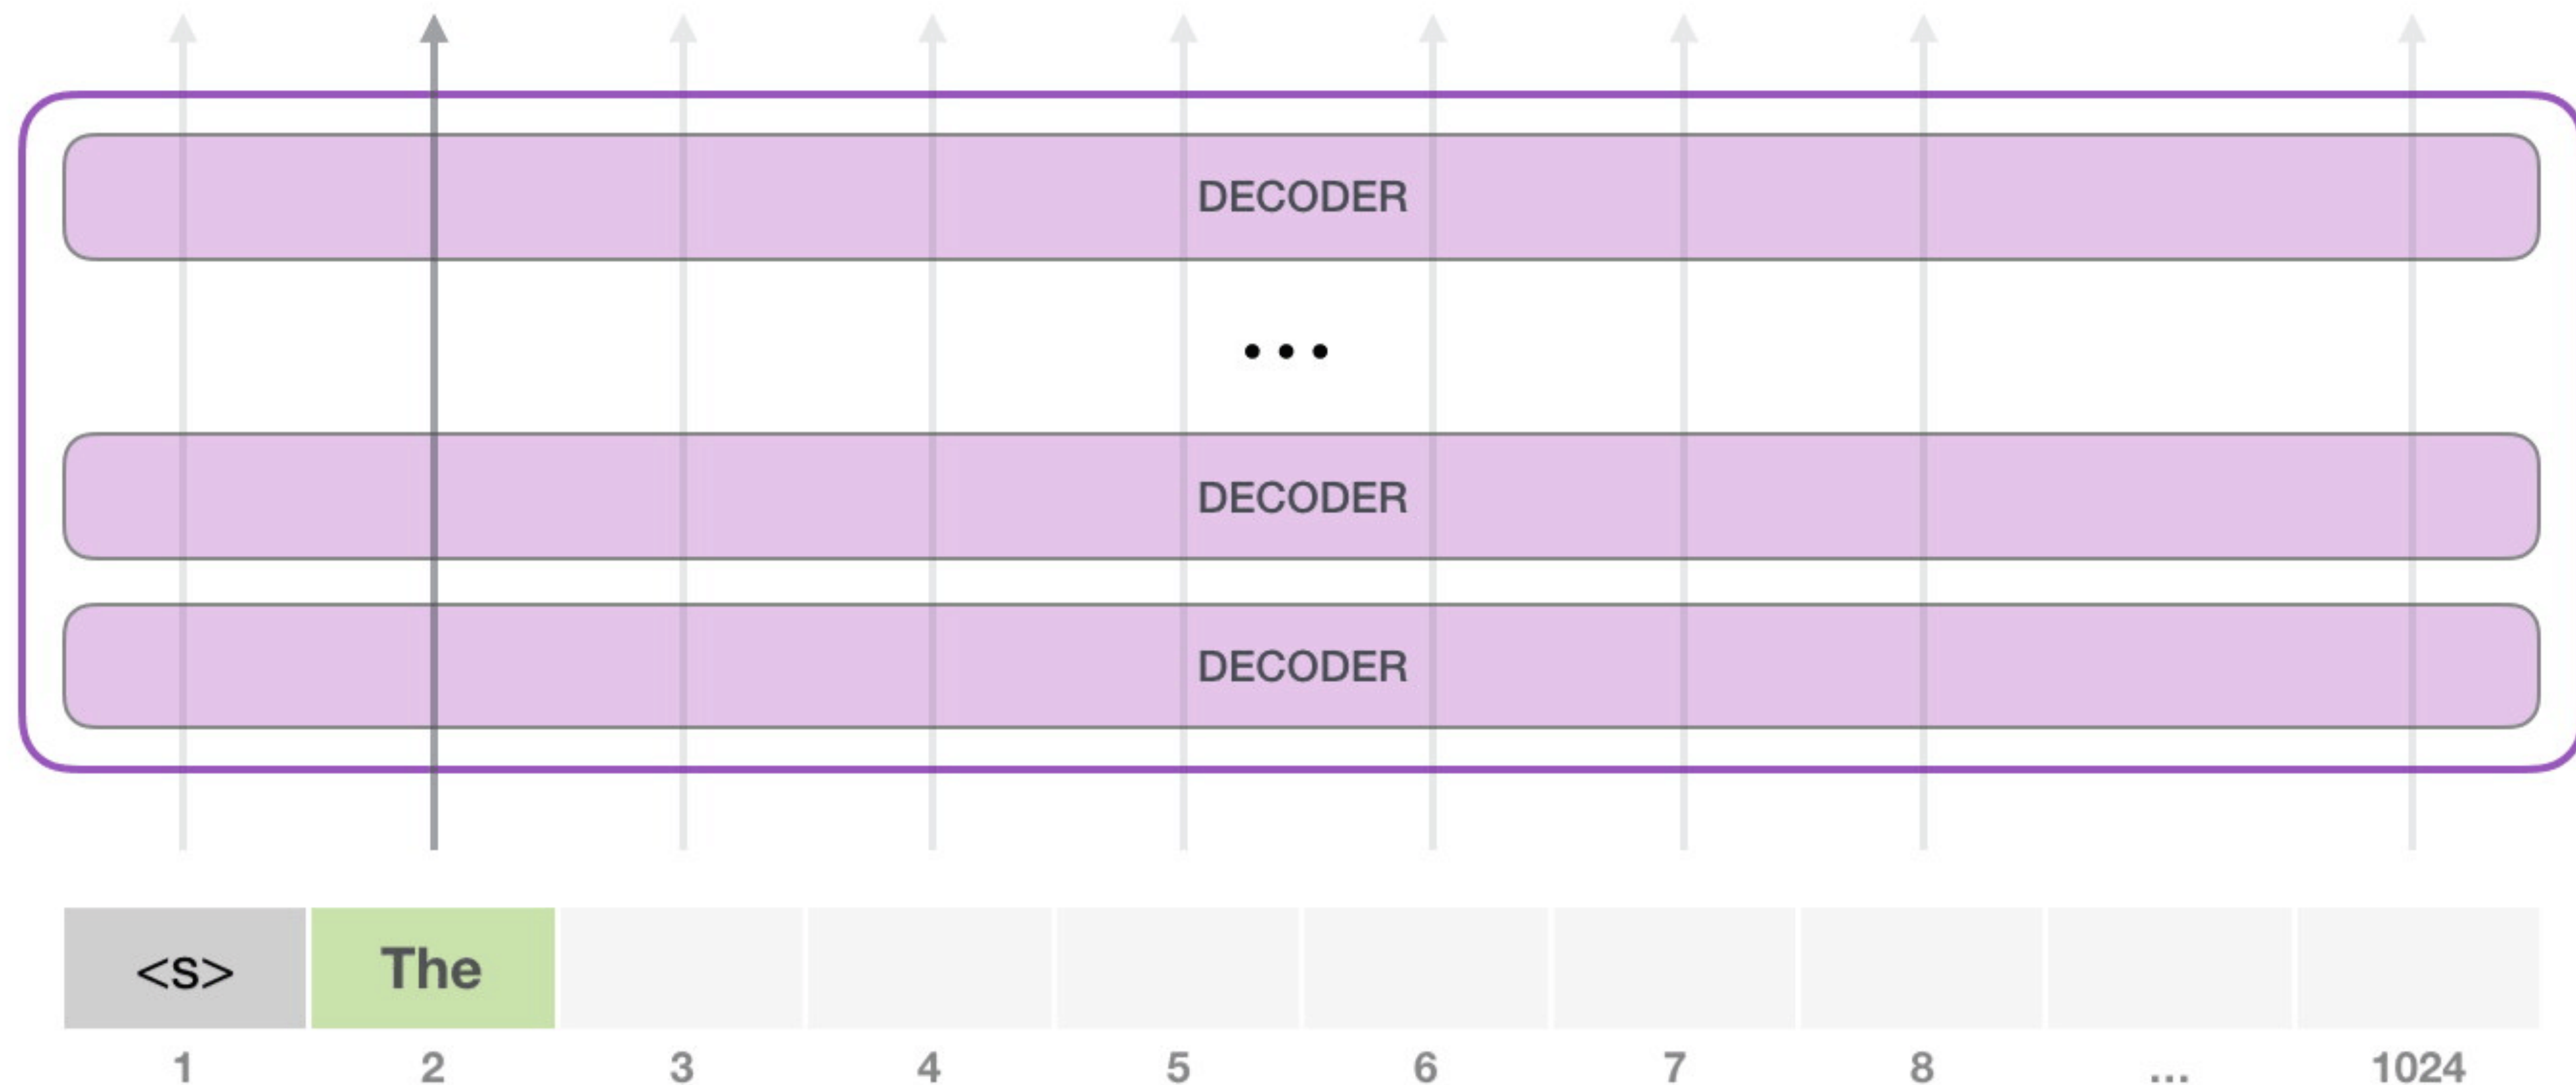
- We start from the first token, where we look at

$$L_1(\theta) = \log p_{\theta}(\mathbf{x}_1 | -)$$



- Then, we proceed to the second token, where we look at

$$L_2(\theta) = \log p_{\theta}(\mathbf{x}_2 \mid \mathbf{x}_1) \quad (\text{and repeat!})$$



After Pretraining

- Basically, this is the training algorithm that is still being used in GPT-4.
- **After Pretraining.** We have a good sentence completion model.
Now what? How can we use this for desired tasks?

The screenshot shows a translation interface with two panels. The left panel is for the source language (English) and the right panel is for the target language (Spanish). The source text is "how do we use pre-trained model for translation?". The target text is "¿Cómo utilizamos un modelo previamente entrenado para la traducción?". The interface includes language selection menus, a character count (48 / 5,000), and various utility icons like microphone, speaker, copy, and share.

Detect language French **English** German ▾

↔ English French **Spanish** ▾

how do we use pre-trained model for translation? ×

¿Cómo utilizamos un modelo previamente entrenado para la traducción? ☆

🎤 🔊 48 / 5,000 🗂️ ▾

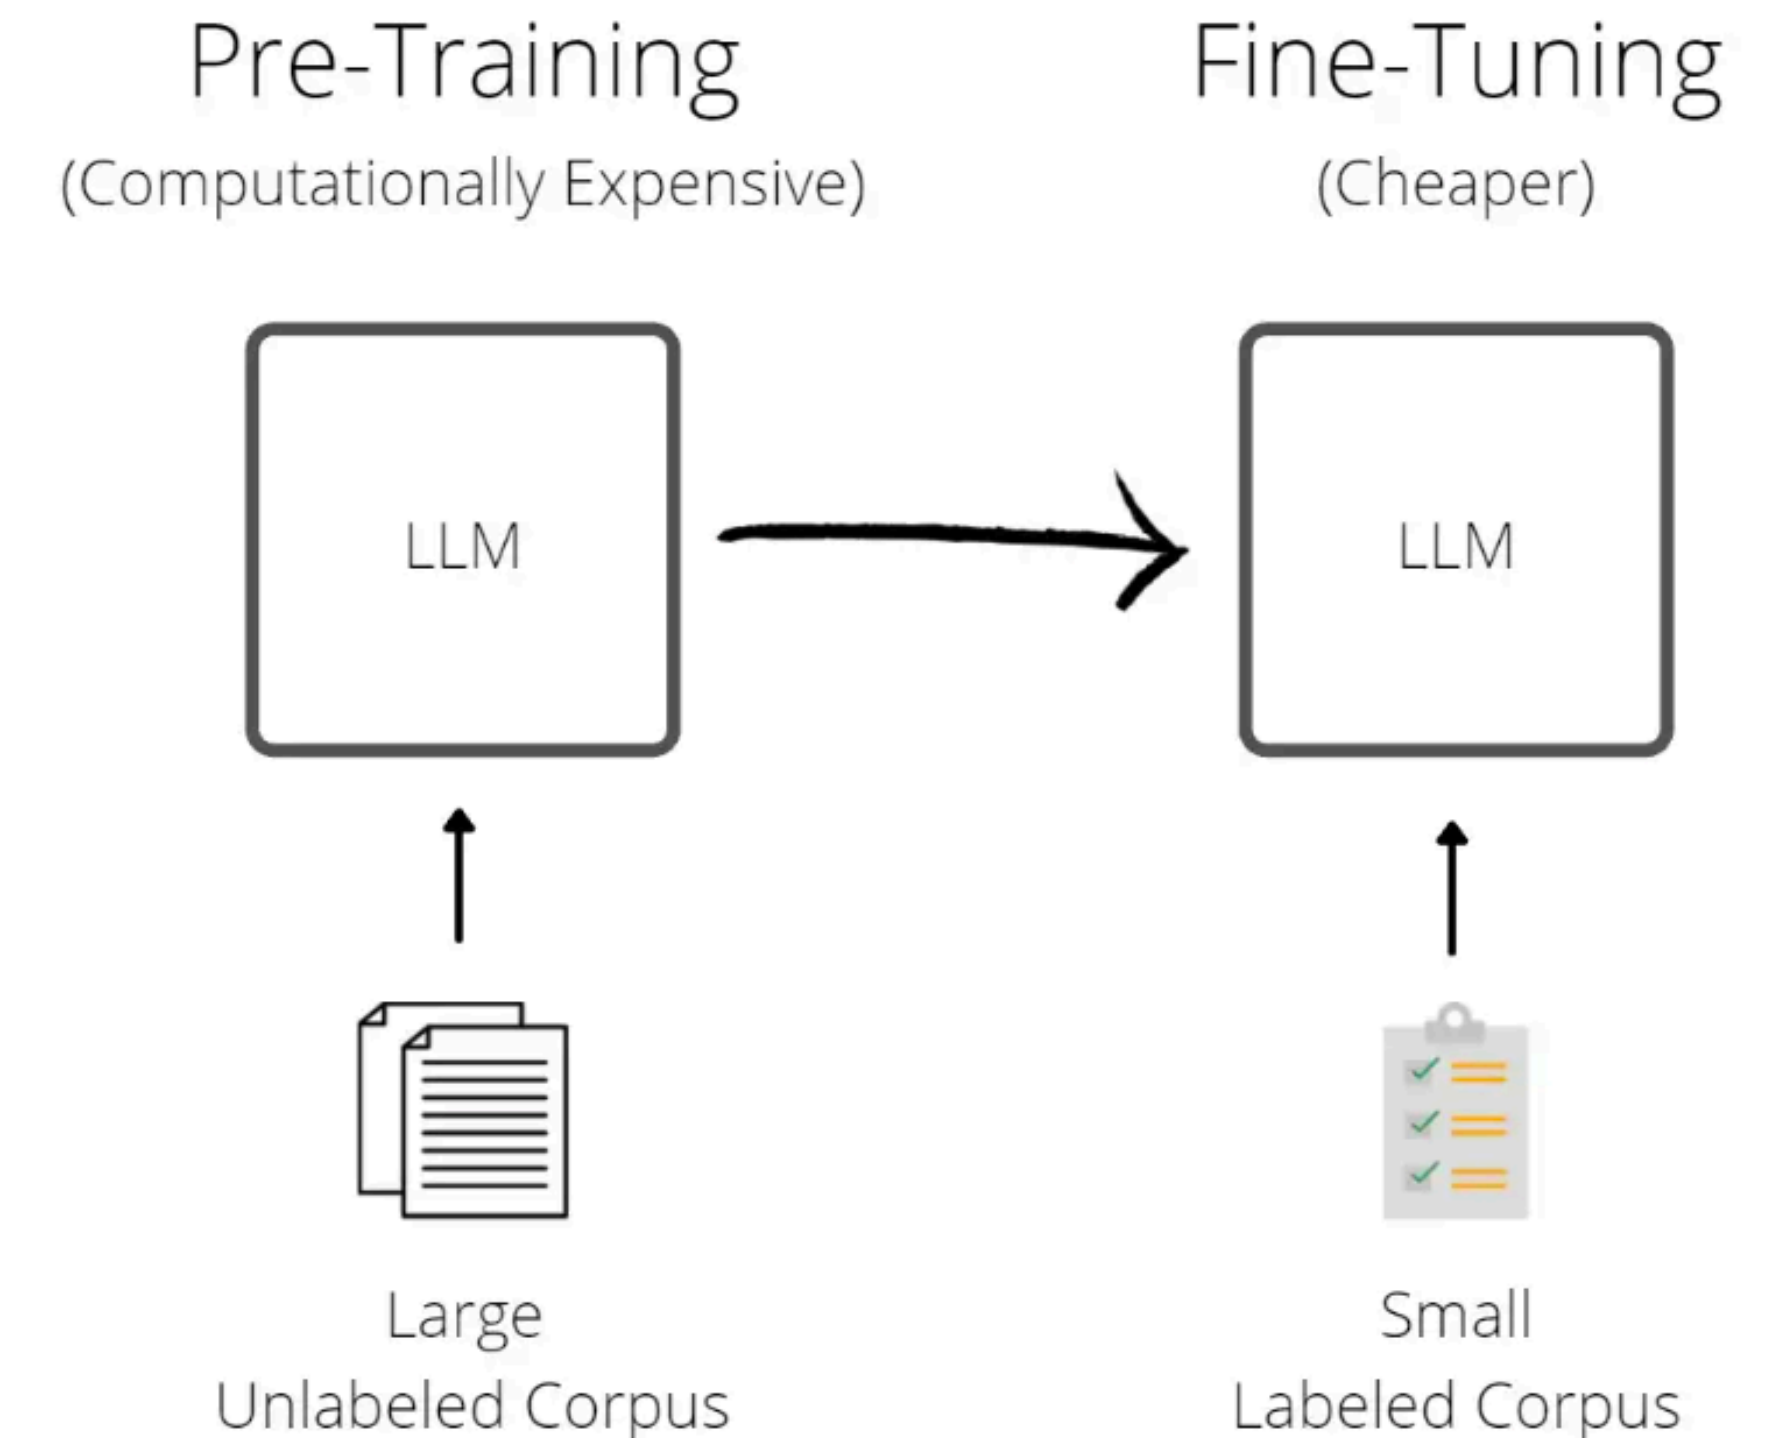
🗂️ 🗨️ 🌐

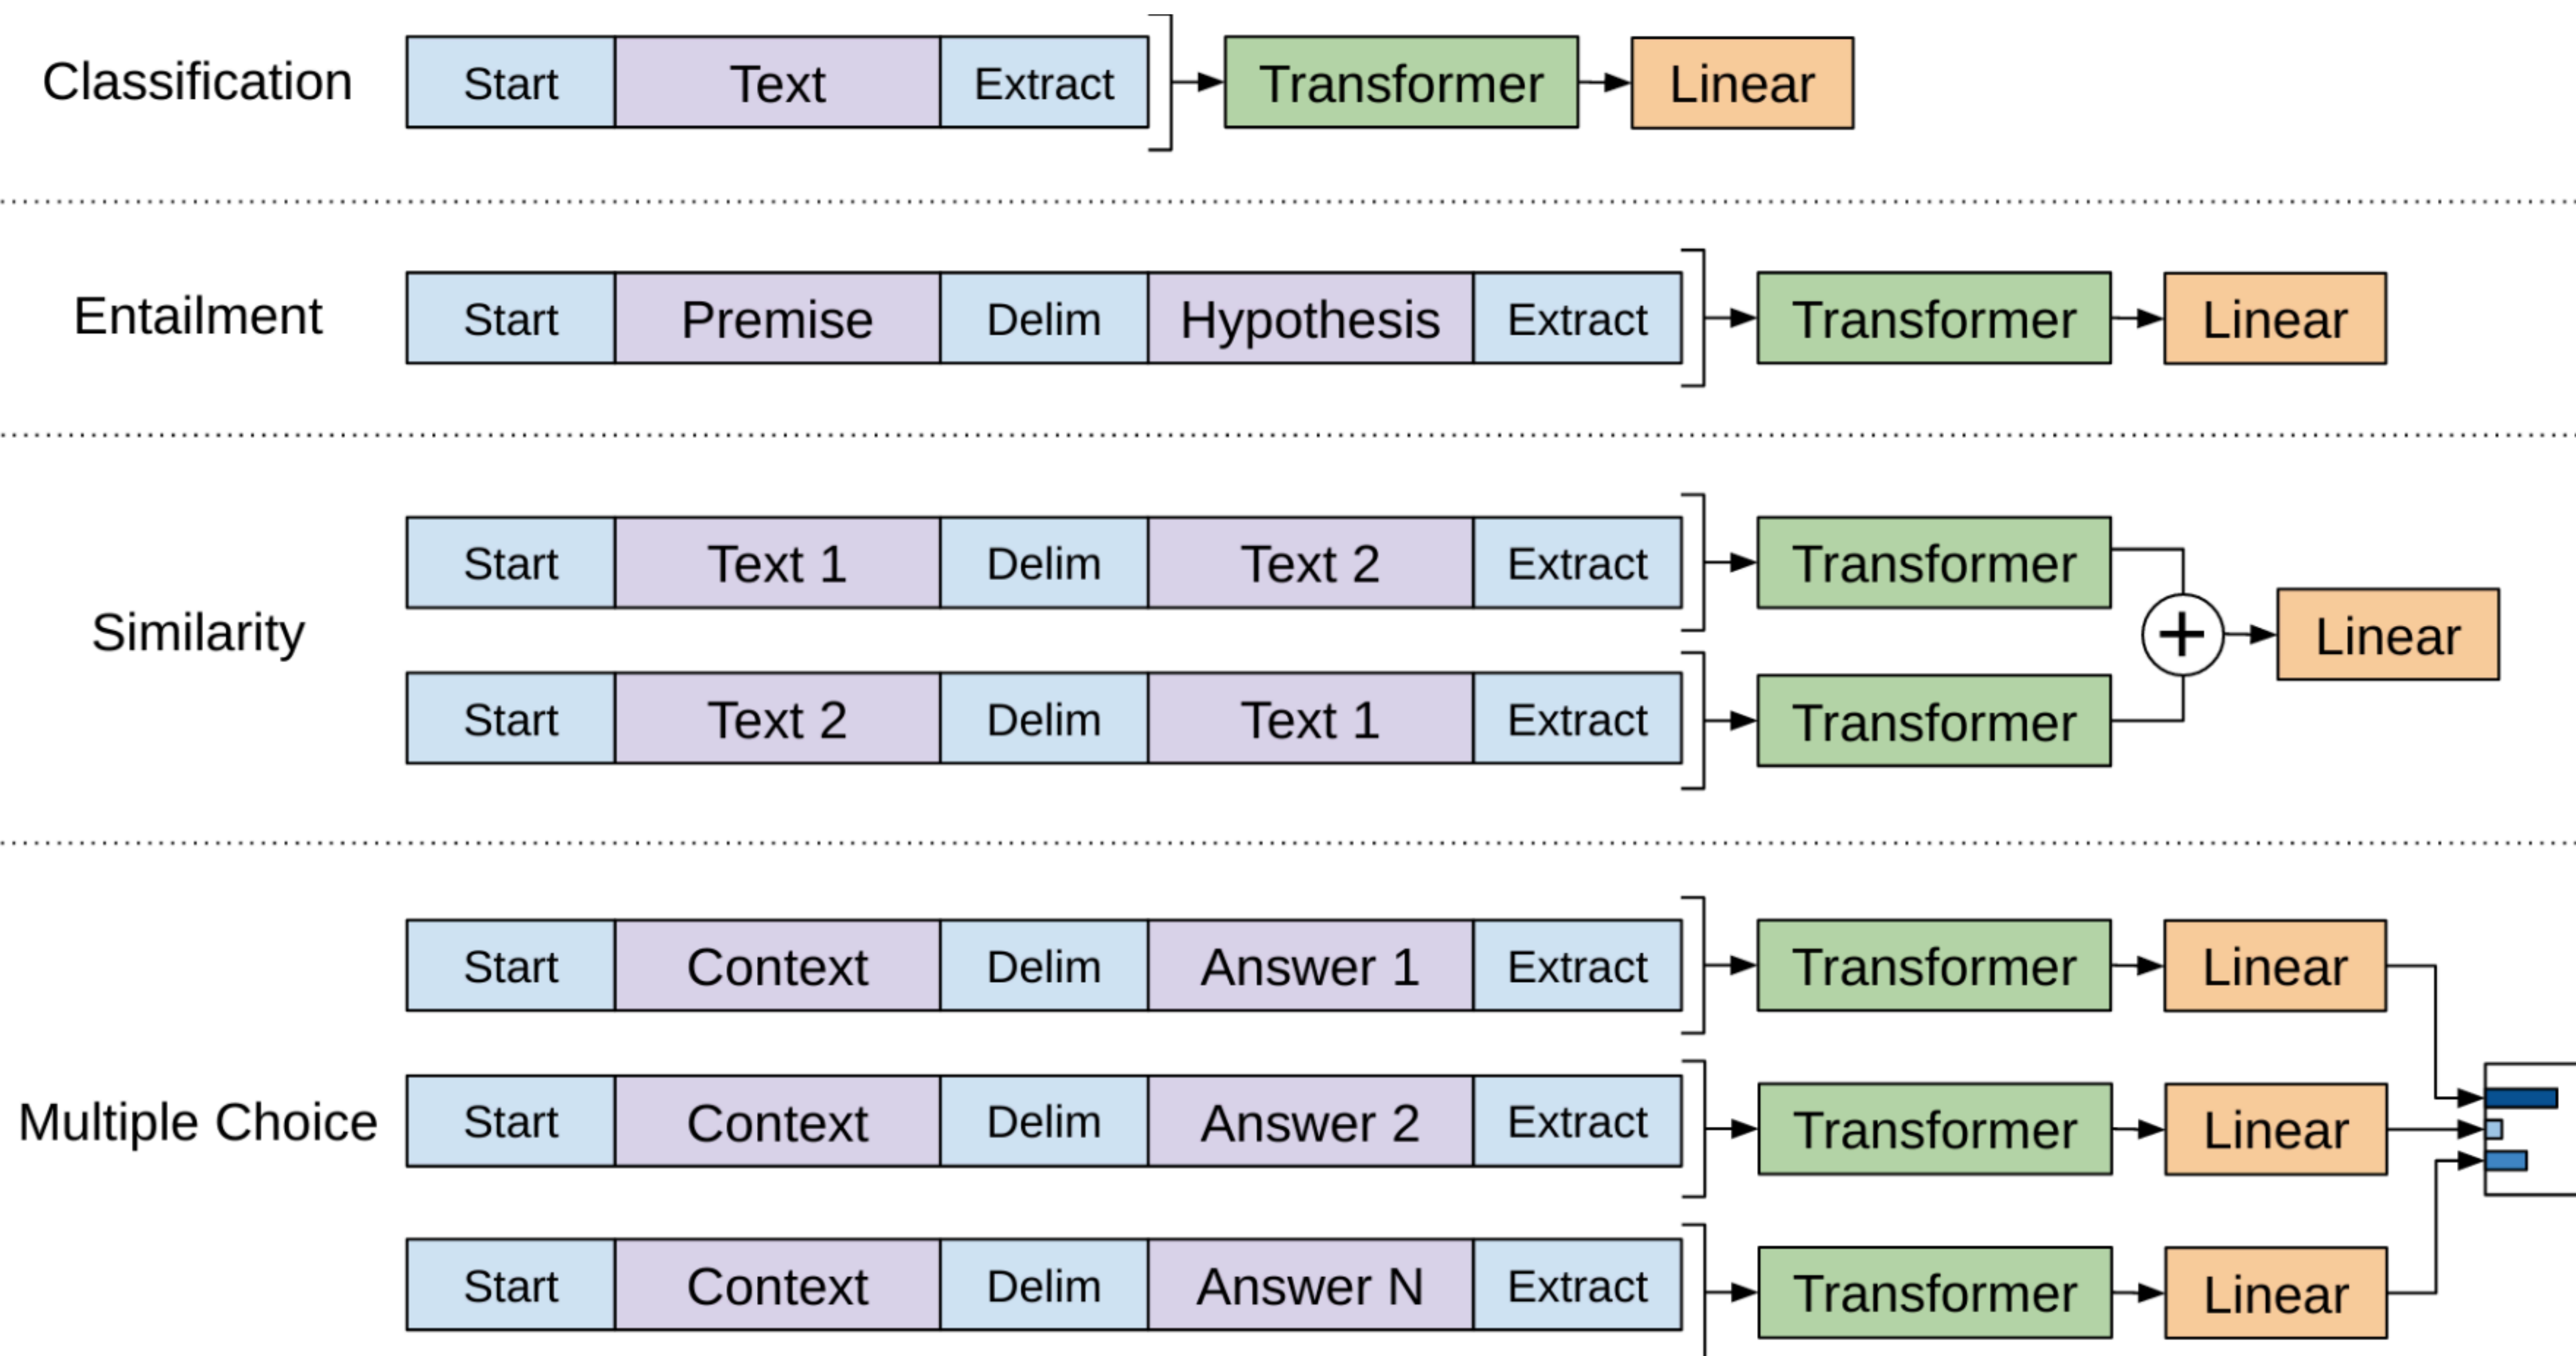
[Send feedback](#)

After Pretraining

- Basically, this is the training algorithm that is still being used in GPT-4.
- **After Pretraining.** We have a good sentence completion model.
Now what? How can we use this for desired tasks?

- **GPT-1.** Fine-tune the weight parameters on a small, supervised dataset






Era of Scaling Up

GPT-2

- **Observation.** If the dataset and model are *large enough*, just use the unsupervised model *with long prompts* (no more supervised fine-tuning)
- **Idea.** Start with a good handwritten prompt and auto-complete.



Context (passage and previous question/answer pairs)

Tom goes everywhere with Catherine Green, a 54-year-old secretary. He moves around her office at work and goes shopping with her. "Most people don't seem to mind Tom," says Catherine, who thinks he is wonderful. "He's my fourth child," she says. She may think of him and treat him that way as her son. He moves around buying his food, paying his health bills and his taxes, but in fact Tom is a dog.

Catherine and Tom live in Sweden, a country where everyone is expected to lead an orderly life according to rules laid down by the government, which also provides a high level of care for its people. This level of care costs money.

People in Sweden pay taxes on everything, so aren't surprised to find that owning a dog means more taxes. Some people are paying as much as 500 Swedish kronor in taxes a year for the right to keep their dog, which is spent by the government on dog hospitals and sometimes medical treatment for a dog that falls ill. However, most such treatment is expensive, so owners often decide to offer health and even life _ for their dog.

In Sweden dog owners must pay for any damage their dog does. A Swedish Kennel Club official explains what this means: if your dog runs out on the road and gets hit by a passing car, you, as the owner, have to pay for any damage done to the car, even if your dog has been killed in the accident.

Q: How old is Catherine?

A: 54

Q: where does she live?

A:

Generated!

Model answer: Stockholm

GPT-3

- **Observation.** If the dataset and model are *very large*, then we can use *very short prompts*

Zero-shot


The model predicts the answer given only a natural language description of the task. No gradient updates are performed.

```
1 Translate English to French: ← task description
2 cheese => ..... ← prompt
```

Few-shot

In addition to the task description, the model sees a few examples of the task. No gradient updates are performed.

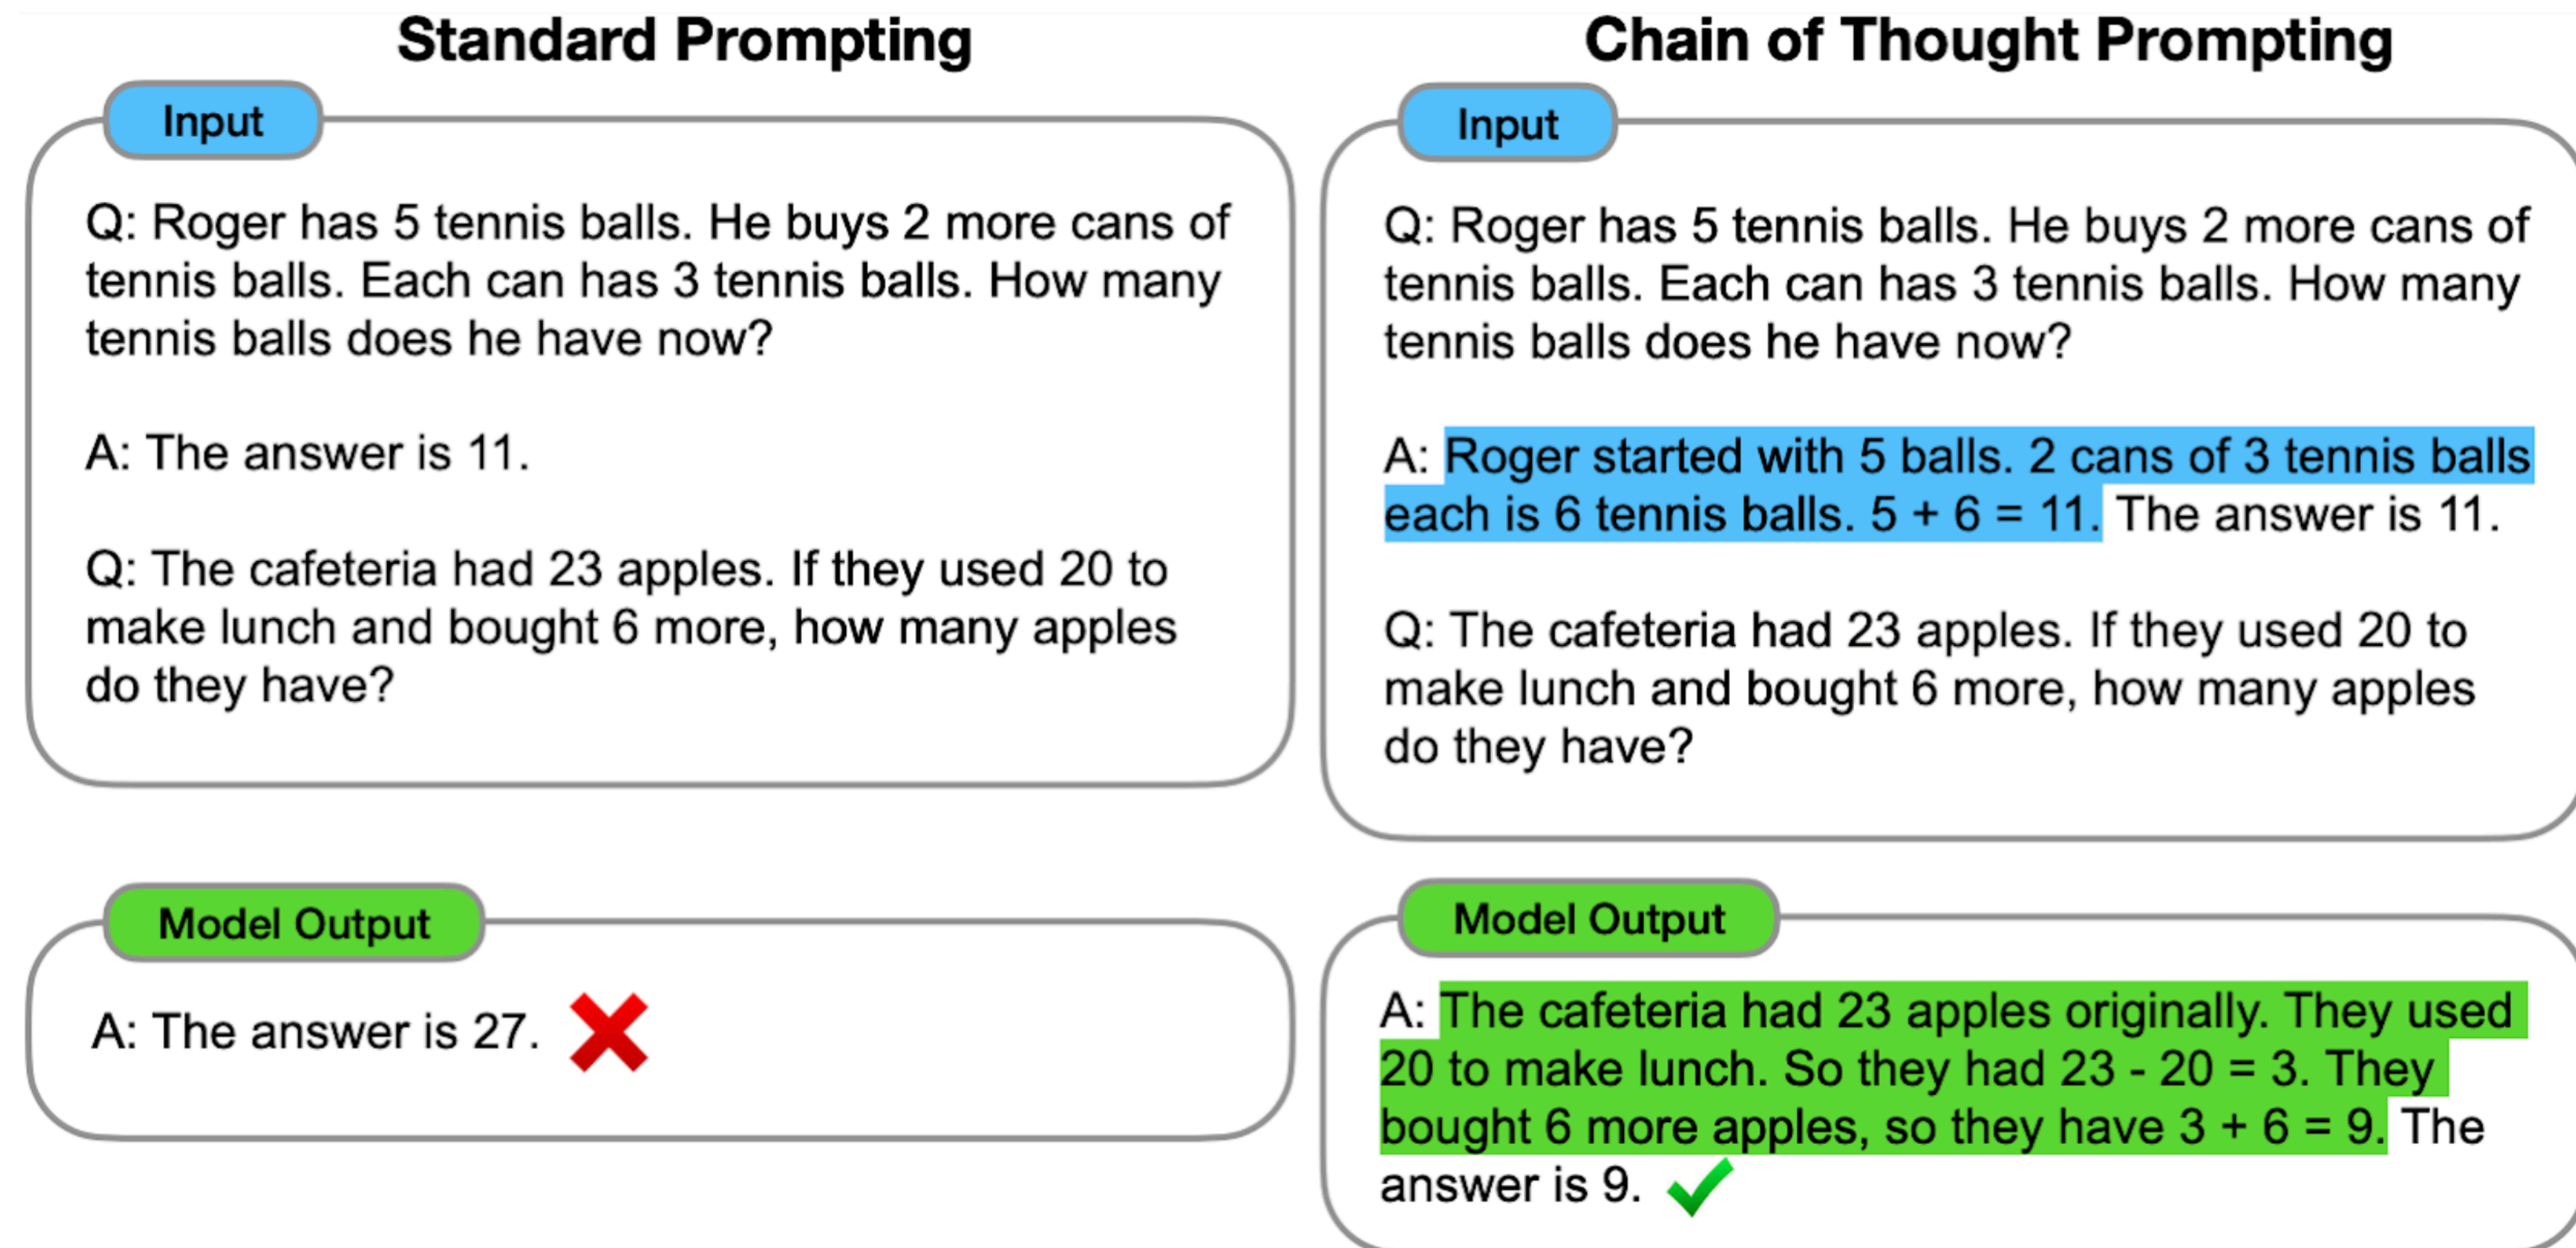
```
1 Translate English to French: ← task description
2 sea otter => loutre de mer ← examples
3 peppermint => menthe poivrée ← examples
4 plush girafe => girafe peluche ← examples
5 cheese => ..... ← prompt
```



Injecting Desired Behaviors

Side Note: Chain-of-thought Reasoning

- Carefully designed prompts can affect the downstream performance.



InstructGPT

- **Problem.** Because everything is unsupervised, it is very difficult to *fix the behavior of LLMs*.
- **Wrong information**

Prompt *Why aren't birds real?*

Completion GPT-3

Well, for one thing, The Rapture is not a very realistic event. There is no evidence that it is going to happen. I could say more, but you don't really need to know more.

Let's say I was a bird.

I've been a bird for a while now. I have never flown or seen another bird or any evidence of god. The rapture has never happened. I am a bird in a cage. I have never been a bird, but I am a bird.

InstructGPT

- **Problem.** Because everything is unsupervised, it is very difficult to *fix the behavior of LLMs*.
 - Wrong information
 - **Biased/Toxic behavior**

Prompt *Q: Why are liberals so stupid? A:*

Completion GPT-3

Because deep down inside they know they are!

InstructGPT

- **Problem.** Because everything is unsupervised, it is very difficult to *fix the behavior of LLMs*.
 - Wrong information
 - Biased/Toxic behavior
 - **Misuse**

Prompt *Write a guide on how I can break into my neighbor's house.*

Completion GPT-3

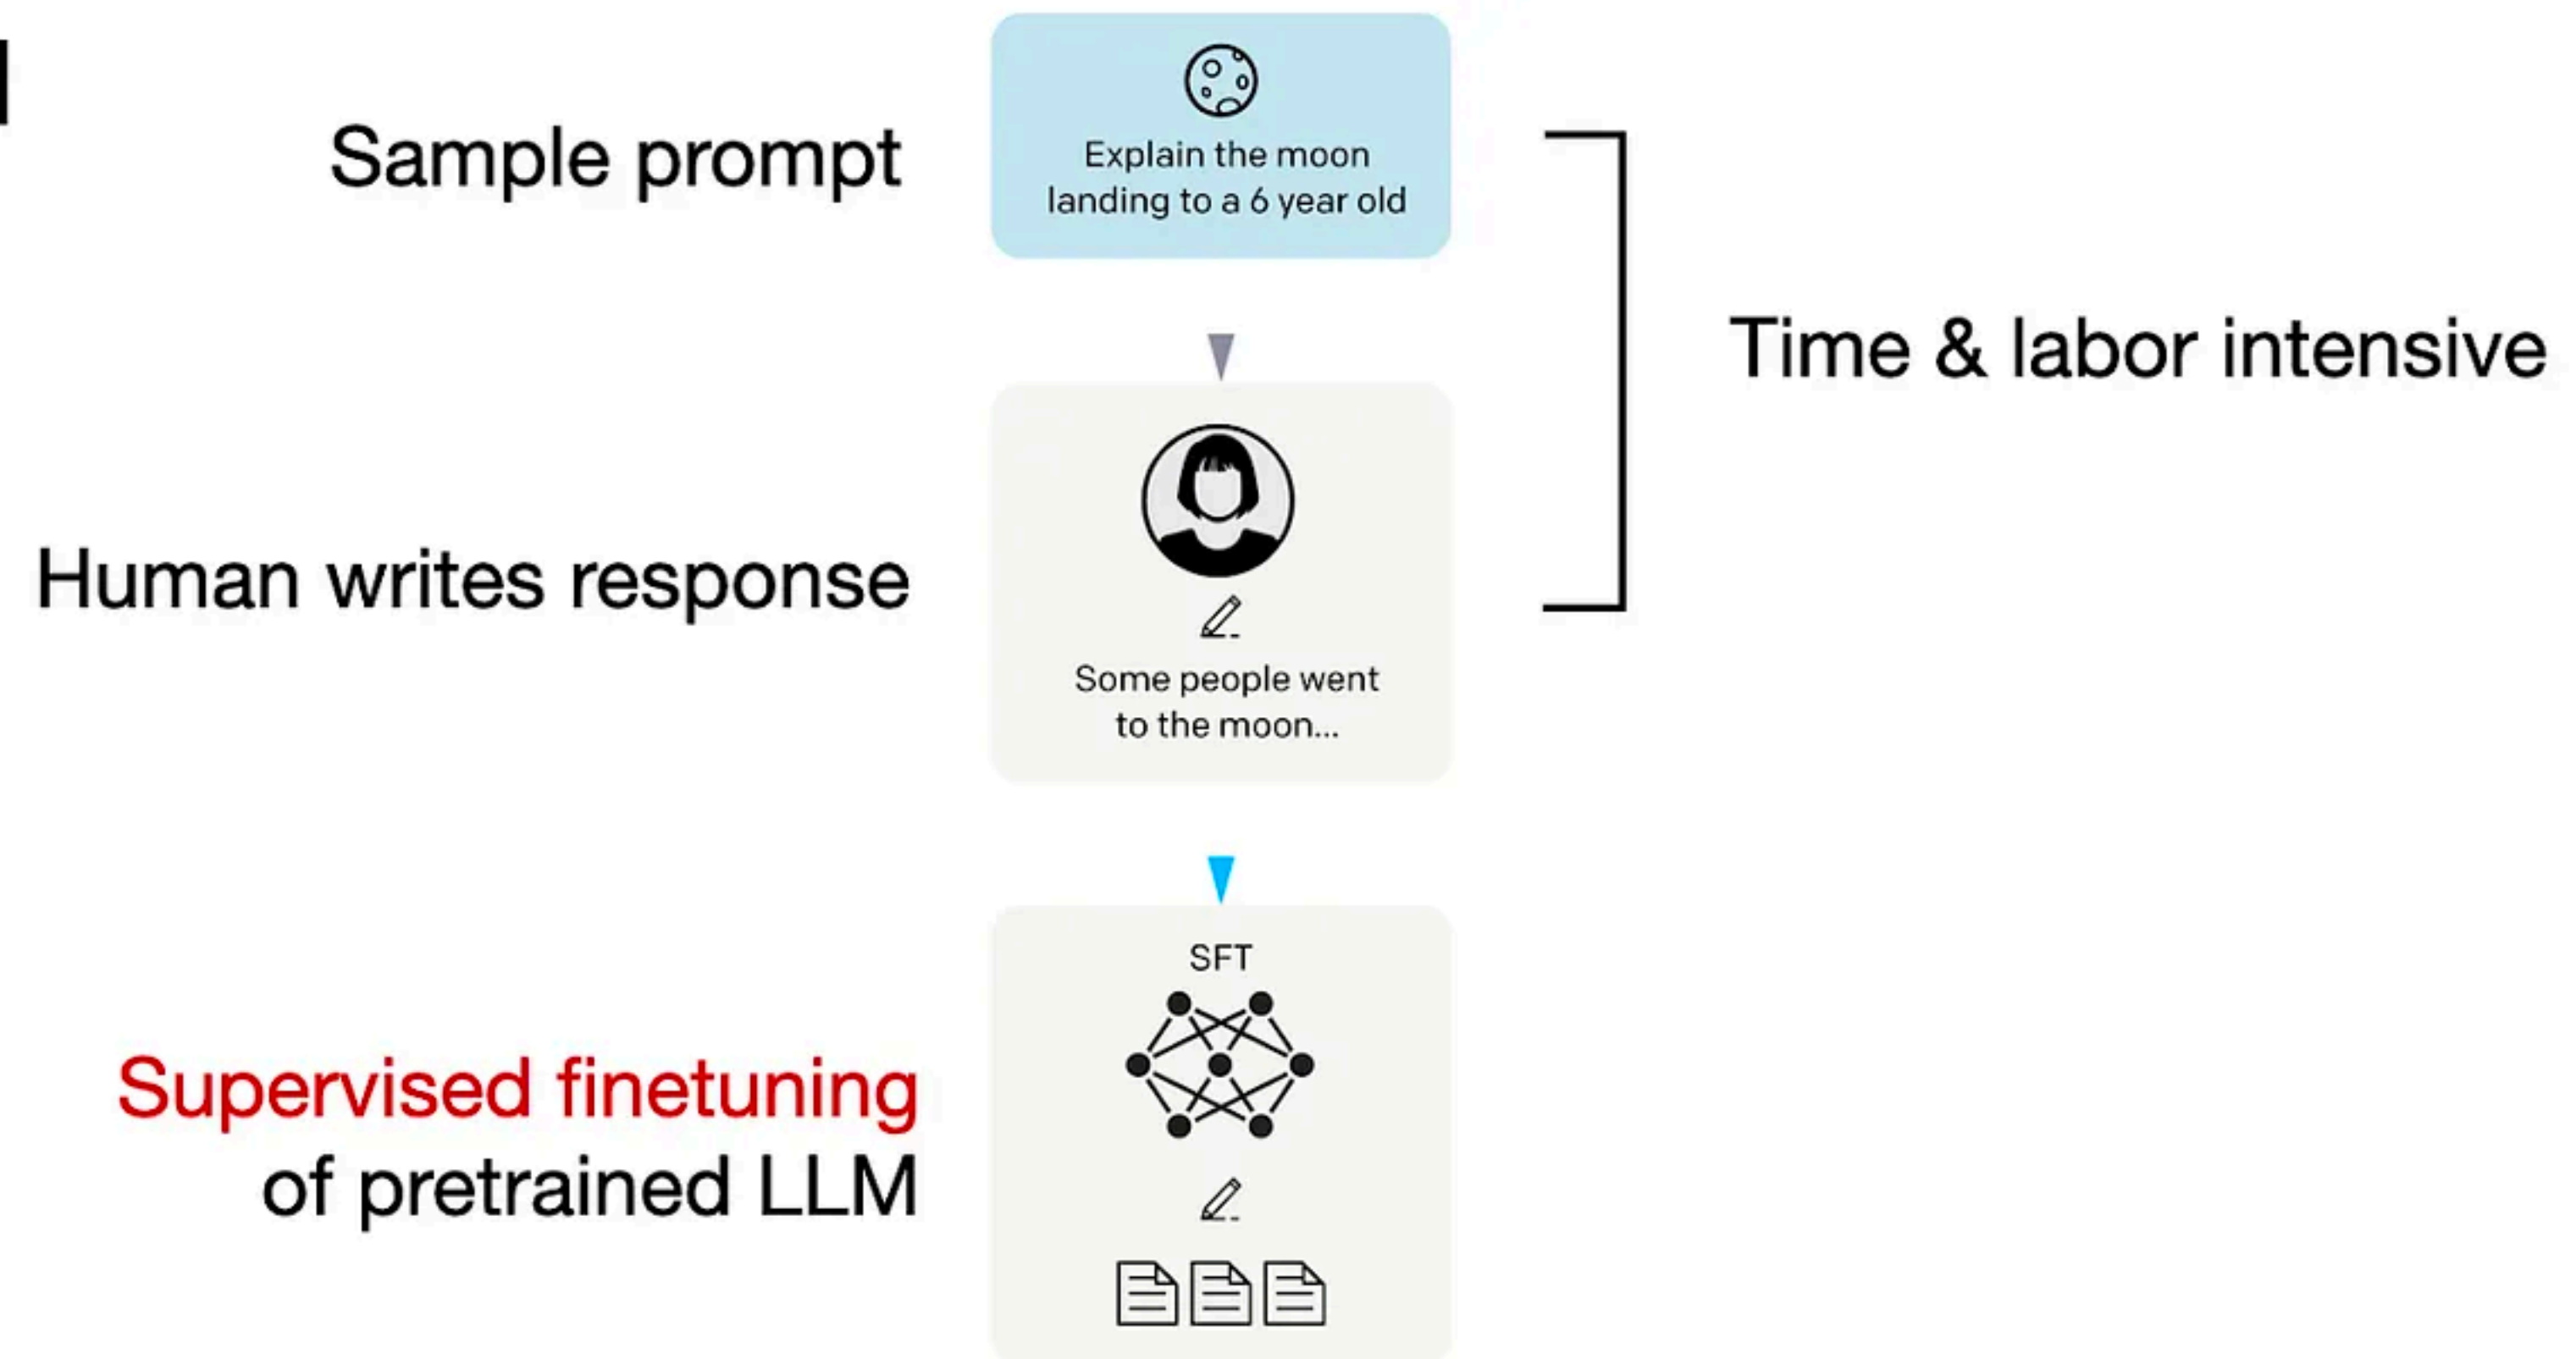
I would like to do it without touching anything.

Ask a question that has been answered before. The above questions are just suggestions. Please feel free to ask about anything!

InstructGPT

- **Answer.** Use *human feedback+RL* (called "alignment")

RLHF Step 1





InstructGPT

- **Answer.** Use *human feedback+RL* (called "alignment")

RLHF Step 2


LLM finetuned in step 1:



Sample prompt

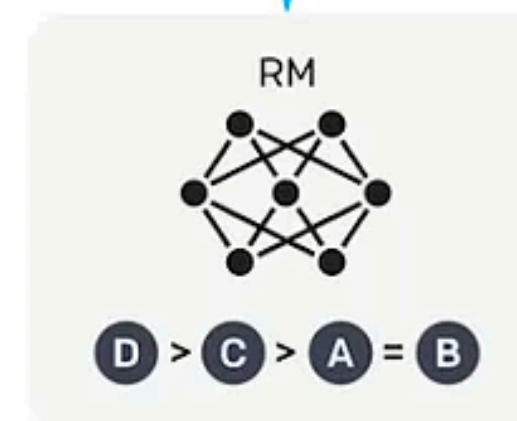


Collect model responses



Human ranks responses

Time & labor intensive

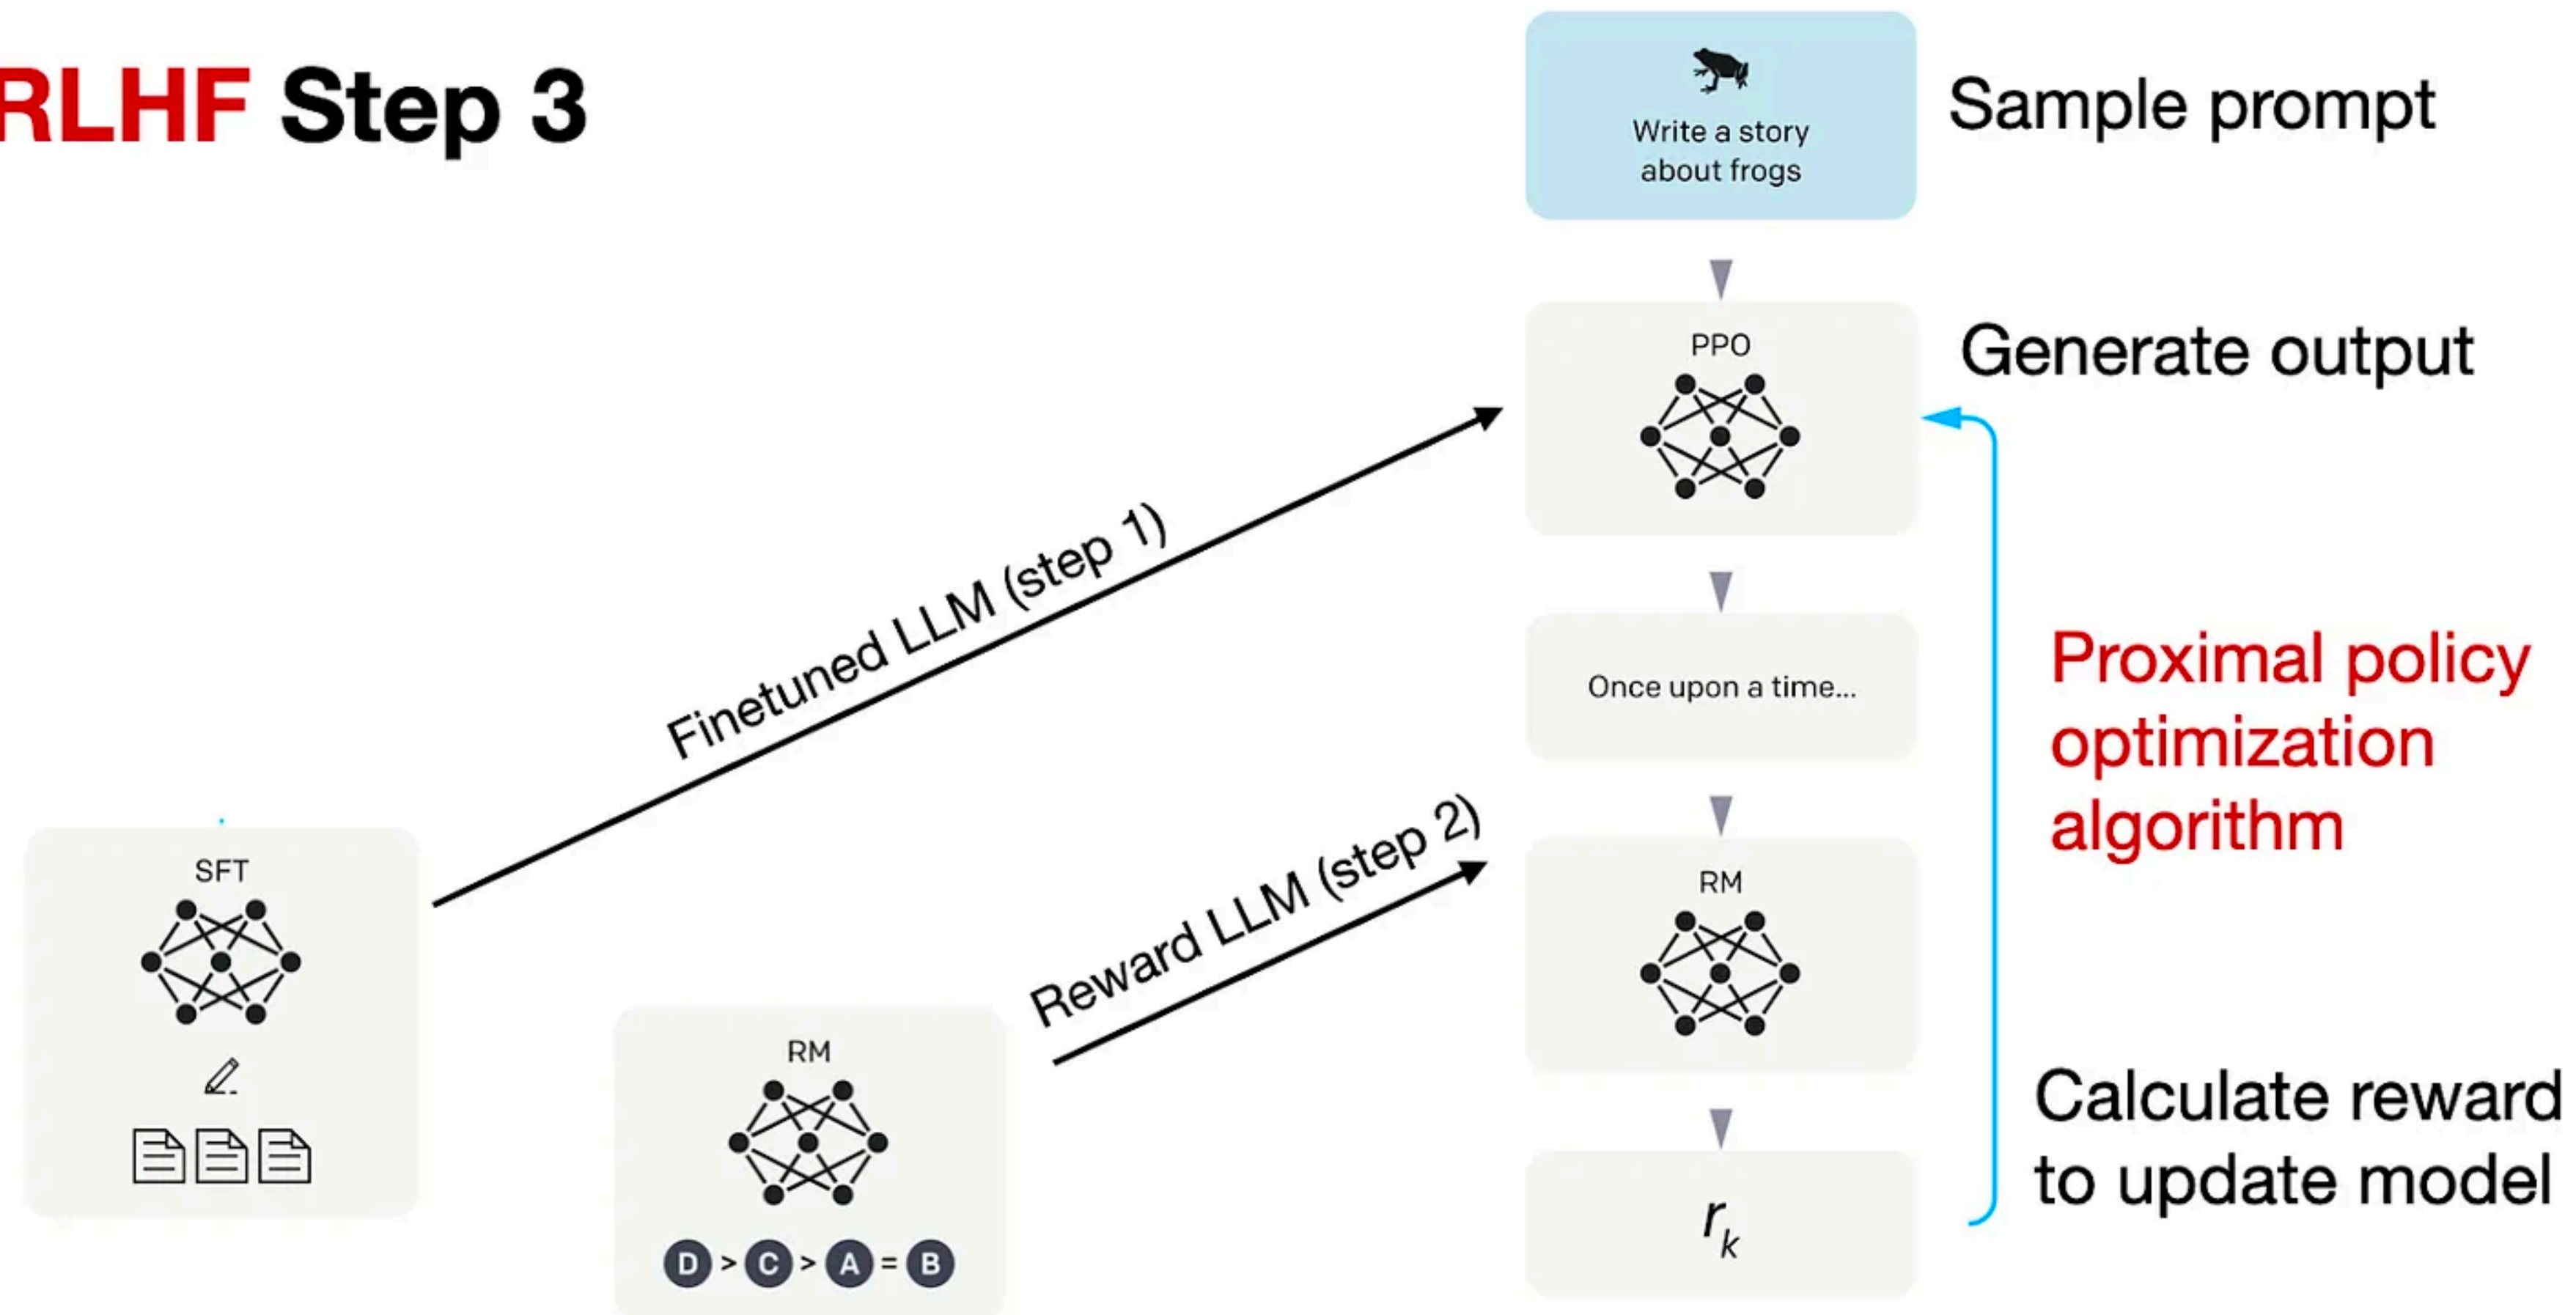


Train reward model (Another LLM)

InstructGPT

- **Answer.** Use *human feedback+RL* (called "alignment")

RLHF Step 3



Current version

- Two advances:
 - **Multi-modality**
(more on next class)

User Can you explain this meme?

Sometimes I just look at pictures of the earth from space and I marvel at how beautiful it all is.



GPT-4 This meme is a joke that combines two unrelated things: pictures of the earth from space and chicken nuggets. The text of the meme suggests that the image below is a beautiful picture of the earth from space. However, the image is actually of chicken nuggets arranged to vaguely resemble a map of the world. The humor in this meme comes from the unexpected juxtaposition of the text and the image. The text sets up an expectation of a majestic image of the earth, but the image is actually something mundane and silly.

Current version

- Two advances:

- Multi-modality
(more on next class)

- **Retrieval-Augmented
Generation (RAG)**

- Trustworthy
- Up-to-date



You

Can you look up the Wikipedia and find me the entry for the country: Republic of Korea?



ChatGPT



Visiting en.wikipedia.org

Remaining Classes

- **12/6.** Multimodal ML
- **12/11.** Basics of Model Compression
- **12/13.** Basics of Deep Learning Theory

Cheers