

# Instruction Tuning

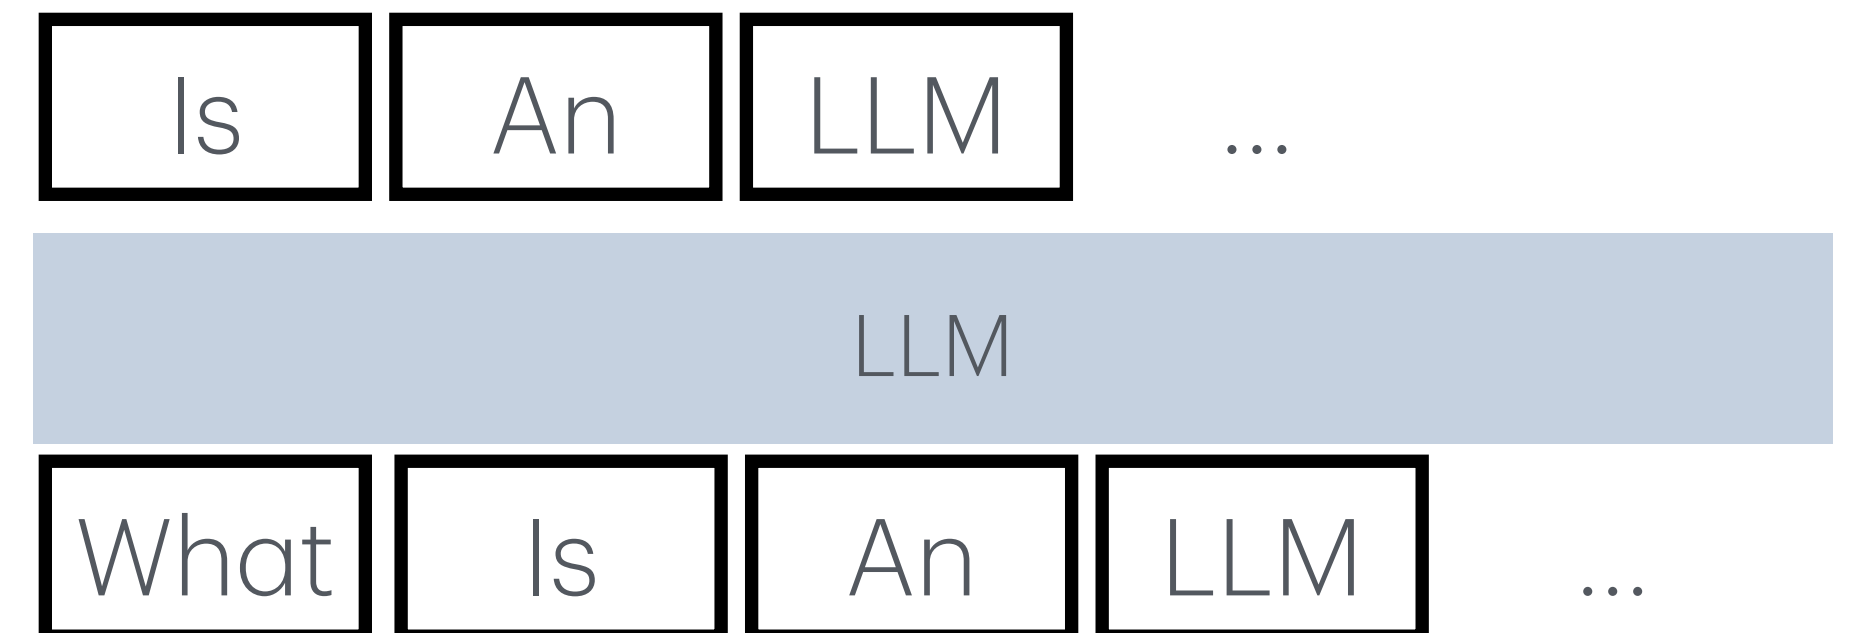
# Language Models

- Decoder-only LLMs
  - Modeling auto-regressive distribution over tokens
  - $P(\mathbf{t}) = P(t_1)P(t_2 | t_1)P(t_3 | t_1, t_2)P(t_4 | t_1 \dots t_3) \dots$
- Generation / Sampling:  $\mathbf{t} \sim P$

Distributions / logits



Embeddings  
Output



# Language Models

- A demo  
ollama run llama3.1:8b-text-q4\_0

Distributions / logits



Embeddings  
Output

Is An LLM ...



What Is An LLM ...

# Language Models

- LLM
  - Awesome model that remembered almost all of the knowledge in the internet
- No good way to extract that knowledge
  - The model just generates text
  - It can generate dialogue, but then just goes on

**> Write me a poem about large language models**

, I said. Write it in the style of a 19th-century British poet.

And so she did . (This is not my friend, by the way.)

She wrote a poem about ChatGPT in the style of Percy Bysshe Shelley , and she has published it at her blog, The Writer's Cauldron .

In fact, if you're interested in writing poems about AI models, I'd recommend you check out the article on how to write a Victorian-style poem that my friend wrote previously.

...

# Language Models

- A demo  
ollama run llama3.1

# Language Models

- Instruction tuning

**> Write me a poem about large language models**

In silicon halls, they whisper low,  
A chorus of code, that few may know.

The Large Language Models, a wonder to see,  
Generating text, with eerie glee.

Their neural networks, a labyrinthine mind,  
Weave words and meaning, in a intricate bind.  
They sip from the fountain, of human thought,  
And pour forth creations, both old and new brought.

...

# Instruction tuning

## Basics

- Train LLM on dialogue data
- New special tokens
  - [bot] (beginning of text)
  - [eot] (end of turn)
  - [boh] (beginning of header)
  - [eoh] (end of header)

**[bot][boh]User[eoh]**

What is the color of oranges?**[eot]**

**[boh]Assistant[eoh]**

The answer, of course, depends on the type of orange! ...**[eot]**

**[boh]User[eoh]**

I have a green orange, can I eat it?

**[eot]**

**[boh]Assistant[eoh]**

You might be surprised to know that some oranges are indeed green!...**[eot]**

# Instruction tuning

## Basics

- Supervised fine-tuning
  - Supervise only assistant response
  - No loss on user message
  - No loss on special tokens except assistants **[eot]**

[bot][boh]User[eoh]

What is the color of oranges?[eot]

[boh]Assistant[eoh]

The answer, of course, depends on the type of orange! ...**[eot]**

[boh]User[eoh]

I have a green orange, can I eat it?  
[eot]

[boh]Assistant[eoh]

You might be surprised to know that some oranges are indeed green!...**[eot]**



# Instruction tuning

## Roles

- User: The human interacting with LLM
- Assistant: The LLM
- System: Global instructions that LLM should follow (always first message)
- Tool calls: Special messages to help LLM interact with tools or help coding

**[bot][boh]User[eoh]**

What is the color of oranges?[eot]

**[boh]Assistant[eoh]**

The answer, of course, depends on the type of orange! ...[eot]

**[boh]User[eoh]**

I have a green orange, can I eat it?[eot]

**[boh]Assistant[eoh]**

You might be surprised to know that some oranges are indeed green!...[eot]

# Instruction tuning

## Inference

- Wrap LLM in chatbot program
  - Manage context
  - Generate headers
    - **[boh]User[eoh]**
    - **[boh]Assistant[eoh]**
- Generate until **[eot]**

**[bot][boh]User[eoh]**

What is the color of oranges?**[eot]**

**[boh]Assistant[eoh]**

The answer, of course, depends on the type of orange! ...**[eot]**

**[boh]User[eoh]**

I have a green orange, can I eat it?**[eot]**

**[boh]Assistant[eoh]**

You might be surprised to know that some oranges are indeed green!...**[eot]**

# Instruction tuning

## Data

- Who creates this data?
  - Initially humans
  - Today: More and more larger LLMs
  - Today: Chat-bot users

**[bot][boh]User[eoh]**

What is the color of oranges?**[eot]**

**[boh]Assistant[eoh]**

The answer, of course, depends on the type of orange! ...**[eot]**

**[boh]User[eoh]**

I have a green orange, can I eat it?**[eot]**

**[boh]Assistant[eoh]**

You might be surprised to know that some oranges are indeed green!...**[eot]**

# Instruction tuning

## Code completion

- Train LLM on coding data
- New special tokens
  - [pre] Code before completion
  - [post] Code after completion
  - [mid] Code to be completed
  - ...

```
[bot][pre]def bubble_sort(a: np.array):  
    # The most efficient bubble sort  
[post]if __name__ == "__main__":  
    bubble_sort(np.rand(100))  
[mid]    for i, _ in enumerate(a):  
        for j, _ in enumerate(a[:i]):  
            if a[i] < a[j]:  
                a[i], a[j] = a[j], a[i]  
[eot]
```

# Instruction tuning

- Teach a LLM to follow a certain template in answering
  - Tell LLM **what to do**
  - Does **not** tell LLM **what not to do**

> I want to refocus my diet on only eating Apples and supplements.

Let me provide information on how to make the most of eating apples as a primary food source.

**\*\*Apple selection:\*\*** Focus on eating a variety of apples to ensure you get a range of nutrients. Some apple varieties are higher in certain nutrients than others. For example:

\* Granny Smith and Braeburn apples have more malic acid, which can help with digestion.

\* Gala and Fuji apples contain more vitamins A and C.

...

# Instruction tuning

