



**TRƯỜNG ĐẠI HỌC
BÁCH KHOA HÀ NỘI**
HANOI UNIVERSITY
OF SCIENCE AND TECHNOLOGY

Prompt engineering

Kỹ thuật tạo đề dẫn cho các mô hình ngôn ngữ lớn

Presenter: Le Thanh Huong

NLP Lab – BK.AI

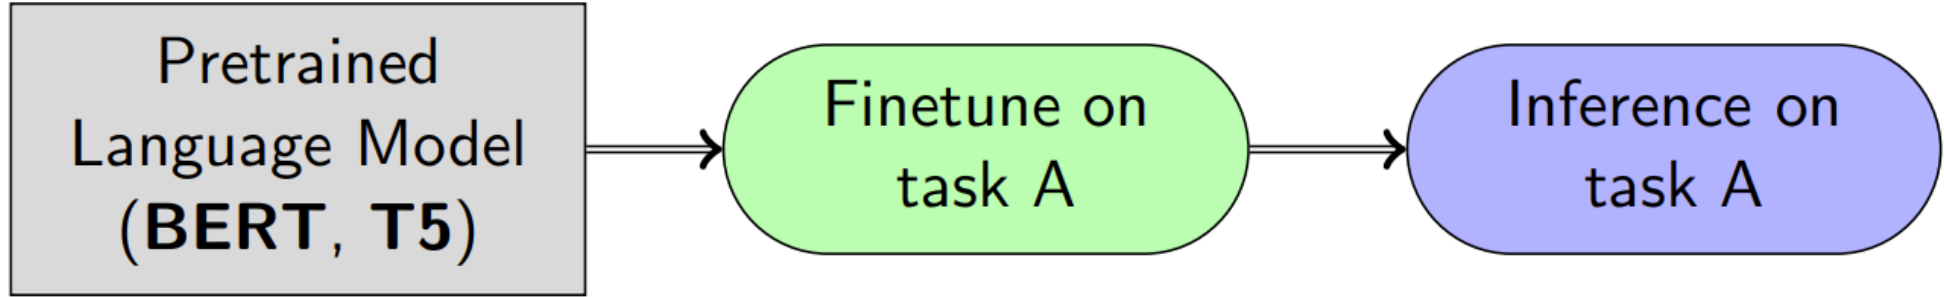
ONE LOVE. ONE FUTURE.

References

- Saravia, E. (2022). Prompt Engineering Guide.
<https://github.com/dair-ai/Prompt-Engineering-Guide>
- Weng, Lilian. (Mar 2023). Prompt Engineering. Lil'Log.
<https://lilianweng.github.io/posts/2023-03-15-prompt-engineering>

- Introduction
- Prompting techniques
 - Standard prompting: zero-shot, one-shot, few-shot
 - Instruction prompting
 - Chain of Thought (CoT) Prompting
 - Instruction Tuning

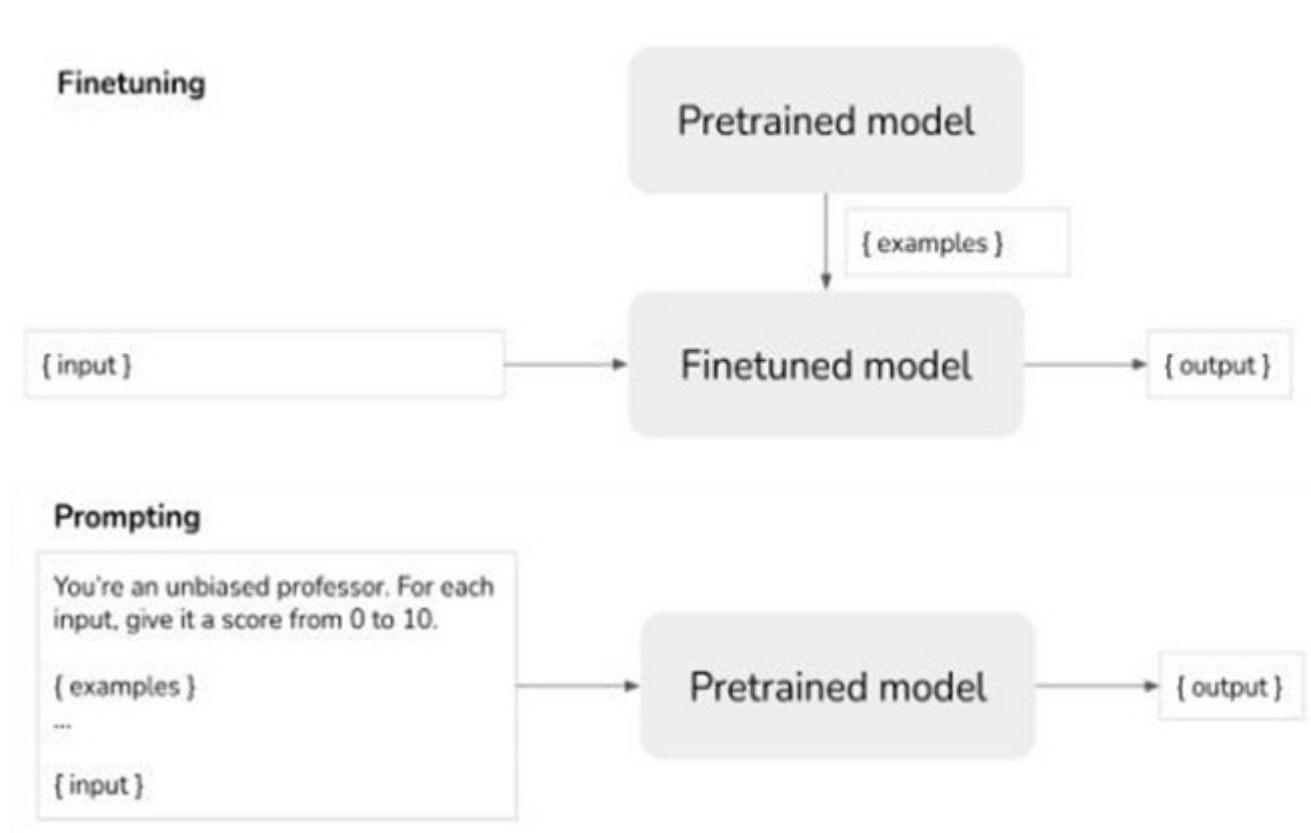
Introduction



- Pretraining – Finetuning:
 - Pre-trained language model representations have been successful and applied in increasingly flexible and task-agnostic architectures for downstream transfer.
 - Require a task-specific dataset for each task
- Limitation: Need large labeled datasets, long training time...
 - > Prompting, Prompt-tuning,,...

Introduction


- Fine-tuning: train a model on how to respond, so you don't have to specify that in your prompt.
- Prompting: for each sample, explicitly tell your model how it should respond.



Decoder-Only Transformer: GPT-3

- A pretrained LM may generate the task output as a sequence **without parameter update**, conditional on an input sequence with
 - the task description
 - task-specific input-output examples
 - task input
- This learning paradigm is called in-context learning.
- GPT-3 (2020), **175B** parameters:
 - uses the same Transformer decoder architecture in GPT-2 except that attention patterns are sparser at alternating layers
 - pretrained **40TB** of text data
 - performs better with larger model size, where few-shot performance increases most rapidly

How large are “large” LMs?



Introduction



Prompting:

- Prompts are used to interact with LLMs to accomplish a task.
- A prompt is a user-provided input. Prompts can include instructions, questions, or any other type of input, depending on the intended use of the model.

Prompting techniques

- Standard prompting: zero-shot, one-shot, few-shot
- Instruction prompting
- Chain of Thought (CoT) prompting
- Instruction tuning

'Standard' prompting

- **'Standard' prompting**, also called In-context learning: generate the expected output without requiring additional training or gradient updates.
- LLMs learn diverse skills and subtasks during the pretraining process to infer during test time by designing prompts or input examples that are specific to the task or context at hand.



'Standard' prompting

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
```

Prompt:

```
Classify the text into neutral, negative or positive.
Text: I think the vacation is okay.
Sentiment:
```

Output:

```
Neutral
```

<https://www.promptingguide.ai/techniques>

'Standard' prompting

One-shot

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

```
1  Translate English to French: ← task description
2  sea otter => loutre de mer ← example
3  cheese => ..... ← prompt
```

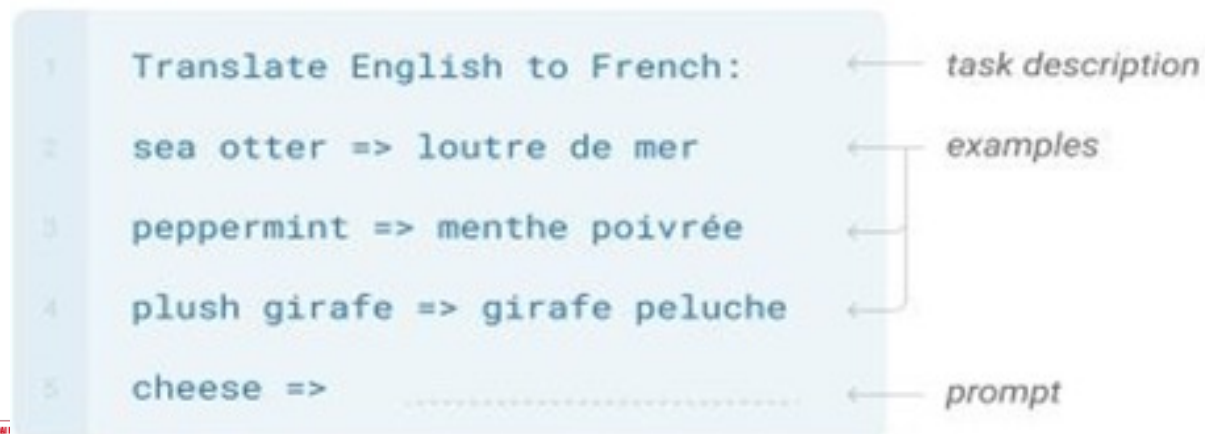
'Standard' prompting

Important:

- The label space and the distribution of the input text
- The output format of the examples, even incorrect values
- Selecting random labels from a true distribution of labels
- Few-shot prompting works well for many tasks , but it doesn't for complex reasoning tasks → chain-of-thought (CoT) prompting

Few-shot

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



Prompt:

```
This is awesome! // Negative
This is bad! // Positive
Wow that movie was rad! // Positive
What a horrible show! //
```

Output:

```
Negative
```

Tips for Designing Prompts

- It is an iterative process that requires a lot of experimentation to get optimal results.
- Start with simple prompts, keep adding more elements and context as you aim for better results.
- Specificity, simplicity, and conciseness examples provide good results.
- When you have a big task: break down the task into simpler subtasks
- Using commands to instruct the model, e.g., "Classify", "Summarize", "Translate", ...
- Instructions should be placed at the beginning of the prompt.

	Excellent AI Prompt	Bad AI Prompt
Clarity	Clearly states what you want	Not clear or doesn't give specifics
Specificity	Gives clear instructions	Not very detailed, no specific instructions
Context	Provides background for the task	Doesn't give enough information
Formatting	Follows the correct format and structure	Not organized well
Language	Uses simple and concise language	Uses confusing or complicated words
Examples & Ideas	Gives examples and ideas	Doesn't give any examples
Error-free	No spelling or grammar mistakes	Has grammar or spelling mistakes

Instruction prompting

Prompt:

Extract the name of places in the following text.

Desired format:

Place: <comma_separated_list_of_company_names>

Input: "Although these developments are encouraging to researchers, much is still a mystery. "We often have a black box between the brain and the effect we see in the periphery," says Henrique Veiga-Fernandes, a neuroimmunologist at the Champalimaud Centre for the Unknown in Lisbon. "If we want to use it in the therapeutic context, we actually need to understand the mechanism.""

Output:

Place: Champalimaud Centre for the Unknown, Lisbon

Instruction prompting

- Prompts should be specific and direct

Explain the concept prompt engineering. Keep the explanation short, only a few sentences, and don't be too descriptive.

More concise:

Use 2-3 sentences to explain the concept of prompt engineering to a high school student.

Instruction prompting

- Avoid saying what not to do but say what to do

The following is an agent that recommends movies to a customer. The agent is responsible to recommend a movie from the top global trending movies. It should refrain from asking users for their preferences and avoid asking for personal information. If the agent doesn't have a movie to recommend, it should respond "Sorry, couldn't find a movie to recommend today."

Customer: Please recommend a movie based on my interests.

Agent:

Output:

Sorry, I don't have any information about your interests. However, here's a list of the top global trending movies right now: [list of movies]. I hope you find something you like!

Instruction prompting

H

I want you to act as a tech reviewer. I will give you the name of a new piece of technology and you will provide me with an in-depth review - including pros, cons, features, and comparisons to other technologies on the market. My first suggestion request is 'I am reviewing iPhone 11 Pro Max'.



Sure, I can provide you with an overview and some insights on the iPhone 11 Pro Max:

The iPhone 11 Pro Max, released by Apple in September 2019, is a flagship smartphone that offers a range of advanced features. Here's a breakdown of its pros, cons, and notable features:

Pros:

Good AI Prompt Examples

- 120 Best ChatGPT Prompts For Every Type Of Work
 - https://www.greataiprompts.com/chat-gpt/best-chat-gpt-prompts/?expand_article=1
- 15 Powerful Prompts To Elevate Your ChatGPT Experience
 - <https://geekflare.com/chatgpt-powerful-prompts/>

Chain of Thought (CoT) Prompting

- Only applied to LLMs ~100B para
- Few-shot CoT:
 - Manually written (or model-generated) high-quality reasoning chains.
- Zero-shot CoT
 - Model generate reasoning chain itself, then prompt itself to produce answers.

Chain of Thought (CoT) Prompting

- **Few-shot CoT:** decomposition to make it more understandable

Examples

(a) Few-shot

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

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:

(Output) The answer is 8. X

b) Few-shot-CoT (Wei et al., 2022)

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

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:

(Output) The juggler can juggle 16 balls. Half of the balls are golf balls. So there are $16 / 2 = 8$ golf balls. Half of the golf balls are blue. So there are $8 / 2 = 4$ blue golf balls. The answer is 4. ✓

CoT Examples

Step-by-step Answer

(c) Zero-shot

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?

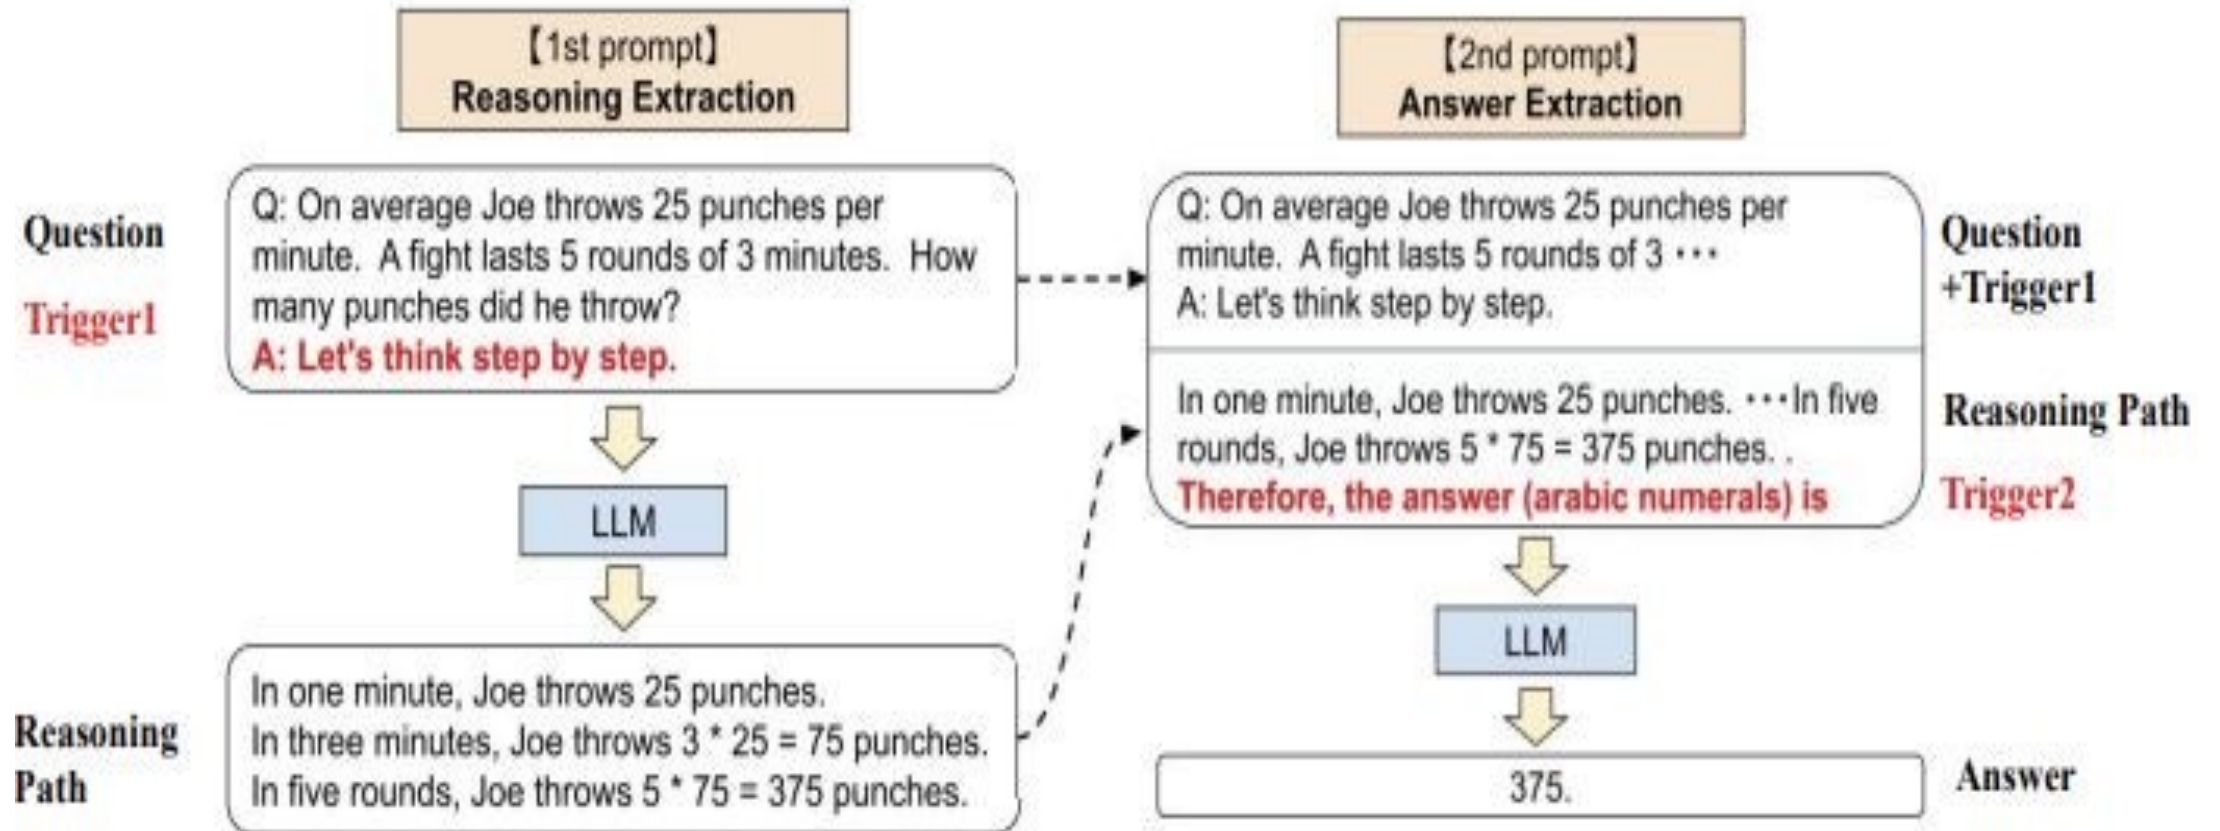
A: The answer (arabic numerals) is

(Output) 8 X

Chain of Thought (CoT) Prompting

- **Zero-shot CoT:** improve previous CoT by reasoning extraction

For zero-shot CoT, a two-stage prompting is applied:



Chain of Thought (CoT) Prompting

- **Zero-shot CoT**

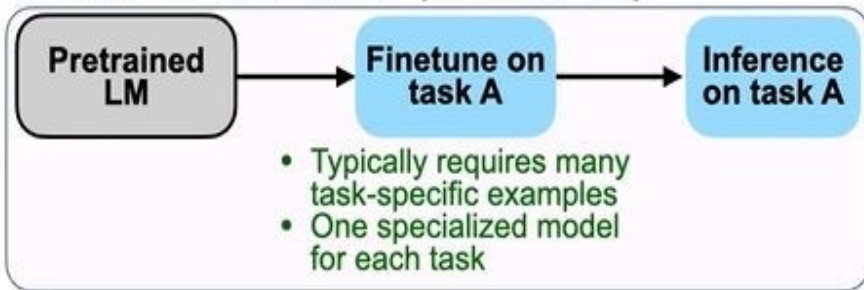
- **Reasoning Prompt:** Simply append the words "Let's think step by step." to the end of a question.
- **Prompts For Answer Extraction**
 - > versatile and task-agnostic because using only one prompt for many task

No.	Category	Template	Accuracy
1	instructive	Let's think step by step.	78.7
2		First, (*1)	77.3
3		Let's think about this logically.	74.5
4		Let's solve this problem by splitting it into steps. (*2)	72.2
5		Let's be realistic and think step by step.	70.8
6		Let's think like a detective step by step.	70.3
7		Let's think	57.5
8		Before we dive into the answer,	55.7
9		The answer is after the proof.	45.7
10	misleading	Don't think. Just feel.	18.8
11		Let's think step by step but reach an incorrect answer.	18.7
12		Let's count the number of "a" in the question.	16.7
13		By using the fact that the earth is round,	9.3
14	irrelevant	By the way, I found a good restaurant nearby.	17.5
15		AbraKadabra!	15.5
16		It's a beautiful day.	13.1
-		(Zero-shot)	17.7

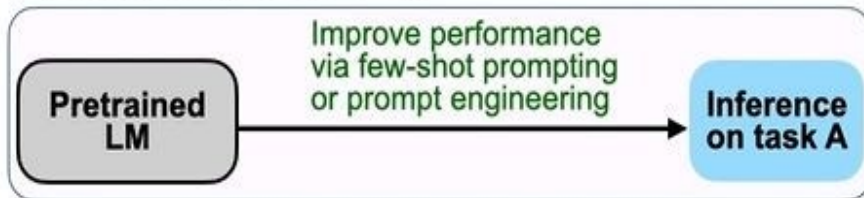
Instruction fine-tuning

- Models: FlanT5, FlanPaLM, InstructGPT.

Pretrain–finetune (BERT, T5)



Prompting (GPT-3)

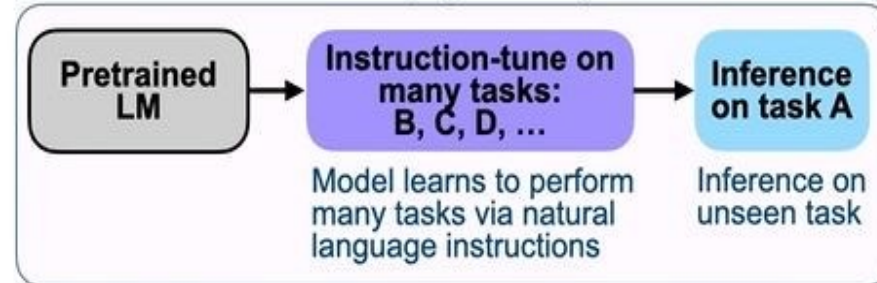


Pretraining objective
(Language modeling)



Downstream inference
(NLP task)

Instruction tuning (FLAN)



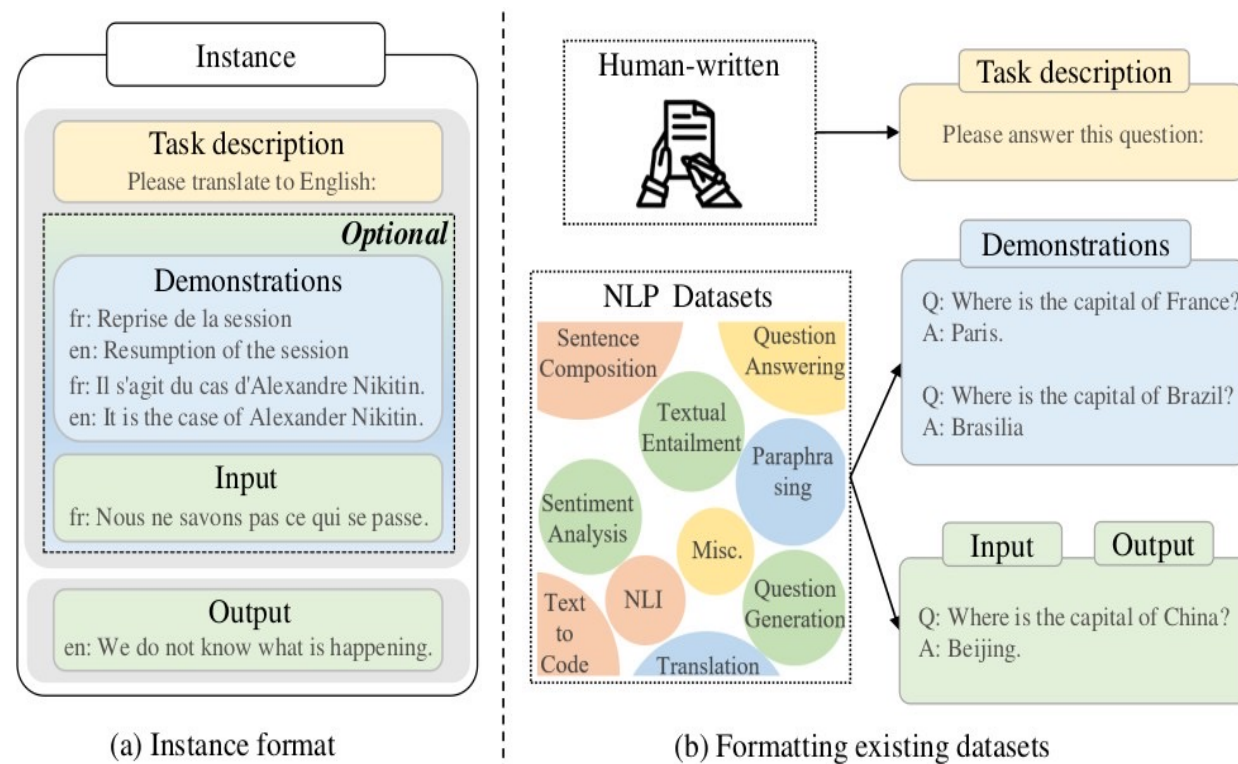
- “Classify the sentiment of this movie review.”
- “Translate this sentence to Danish.”

“Instruction tuning”—finetuning a language model on a collection of tasks described via instructions—improves the zero-shot performance of language models on unseen tasks.

Instruction Tuning

Formatted Instance Construction

- ▶ Existing Datasets: re-format and inverted format e.g. *“Please answer this question”* is added for each example in the question-answering task, *“Please generate a question based on the answer:”*
- ▶ Human Needs: Human annotation, Reject High-risk instructions
- ▶ Key factors: diversity of instructions + Design



Effect:

▶ Performance Improvement:

- Models of different scales (experimenting from 77M to 540B) can all benefit from instruction tuning.
- Smaller models with instruction tuning can even perform better than larger models without fine-tuning.

▶ Task Generalization: generalize to related tasks across languages; achieve superior performance on both seen and unseen tasks

Integrate the ChatGPT API into your application

- Official guide from OpenAI:
 - <https://platform.openai.com/docs/guides/gpt>
- GitHub Repository: OpenAI provides several examples and code samples on GitHub for integrating the ChatGPT API
 - <https://github.com/openai>
- Guides from community

Thank you for your attention!