


# LLMs for Low Resource Languages in Multilingual, Multimodal and Dialectal Settings



<https://llm-low-resource-lang.github.io>

EACL 2024, 21th March, 2024

# Speakers



**Firoj Alam**  
Scientist



**Shammur Chowdhury**  
Scientist



**Sabri Boughorbel**  
Scientist



**Maram Hasanain**  
Post Doctoral  
Researcher

**Qatar Computing Research Institute**



# Content

- Introduction [**20 mins**]
- Models and their capabilities for low-resource languages [**70 mins**]
  - NLP models [40 mins]
  - Multimodality [25 mins]
    - Overview
      - Multimodality
      - Speech
  - QA [5 mins]
- Coffee break [**30 mins**]
- Prompting + Benchmarking Tool [**60 mins**]
  - Prompt Engineering [40 mins]
    - Prompting techniques
    - Cross-/multi-lingual prompting
  - Prompt and Benchmarking tools [15 mins]
  - QA: [5 mins]
- Other Related Aspects [**20 mins**]



# Prompting and Benchmarking Resources



# Prompt Engineering

- Prompt Engineering
- Prompting techniques
- Cross-/multi-lingual prompting
- In-Context/Few-shot Learning



**Being able to communicate clearly in writing**



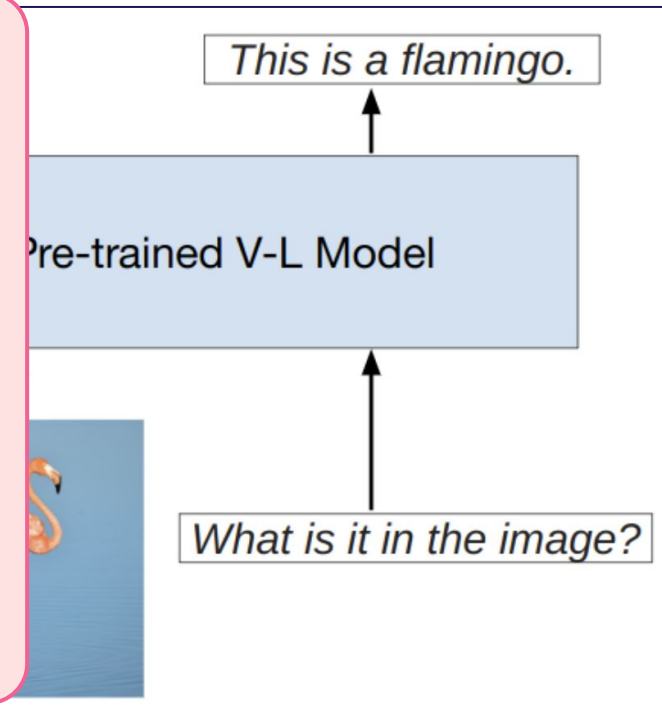
**Prompt Engineering**

# What is a “Prompt”?

An instruction given to LLM to guide it on how to perform a user task

- Instructions
  - Context
  - Input data
  - Output indicator
- ```
Classify the text into neutral,  
Text: I think the food was okay  
Sentiment:
```

Instructions  
Definitions  
Background information  
Questions  
Examples  
Images

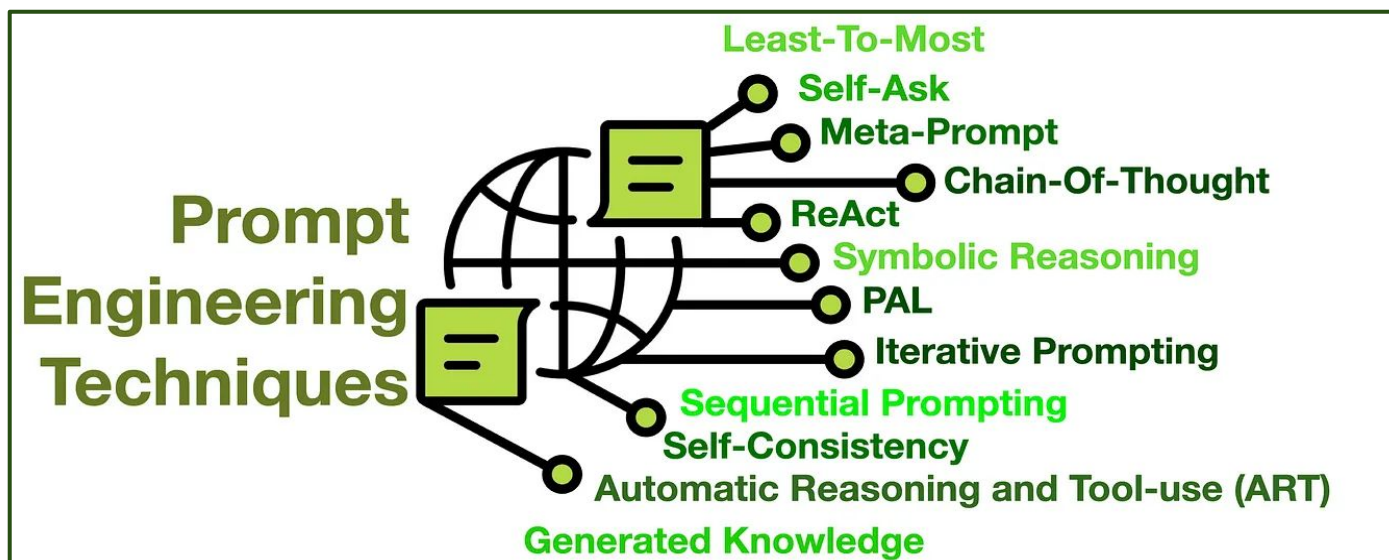


<https://arxiv.org/pdf/2307.12980.pdf>



# What is Prompt Engineering?

An iterative process of developing and optimizing prompts to efficiently use LLMs for a variety of tasks



<https://cobusgreyling.medium.com/12-prompt-engineering-techniques-644481c857aa>



# Prompt Templates

A prompt is converted into a template with key and values replaced with placeholders. The placeholders are replaced with application values/variables *at runtime*.

```
prompt_template = """Act as support staff.  
Help the owners of the HHCR3000 operate their cleaning  
robot by giving answers to questions on features and step-  
by-step instructions when they ask for help.  
  
User: {query}  
Assistant: """
```

# for each conversation turn  
prompt = prompt\_template.format(query=actual user query)

1 prompt\_template instead of prompt

2 Variable in the template.

3 Variable in the template is replaced by current user query to get the prompt

<https://link.medium.com/phuemMDIPHb>



# Types of Prompts

Role-based Prompts

Chain-of-Thought (CoT)

Tree of Thoughts (ToT)

Graph of Thoughts (GoT)

Cross-Lingual-Thought Prompting

Cross-Lingual Tree of Thoughts

Iterative Prompting

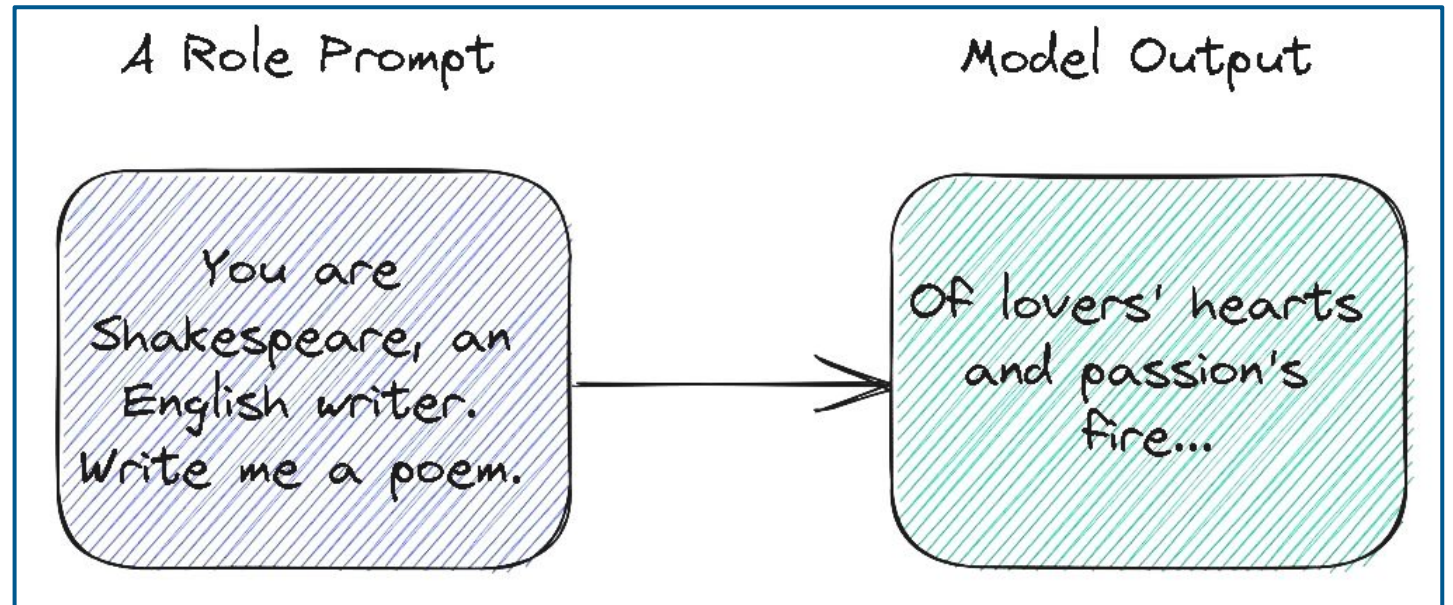


# Role Based Prompts

**Aim:** “set the tone of the conversation”

⇒ Model’s responses more relevant & increases the accuracy.

**How:** Specify the role the model should play.



<https://www.linkedin.com/pulse/role-prompting-aris-ihwan/>





# Chain-of-Thought (CoT) Prompts

**Aim:** Improve the ability of LLM to perform complex reasoning  
⇒ Instruct the model to “think” in smaller steps.

(Wei et al., 2022)

## Model Input

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: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

(Kojima et al., 2022)

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. ✓*


Ask model to “think step by step” without providing examples

Provide LLM with examples with a series of intermediate natural language reasoning steps that lead to final output



# Tree of Thoughts (ToT) Prompts

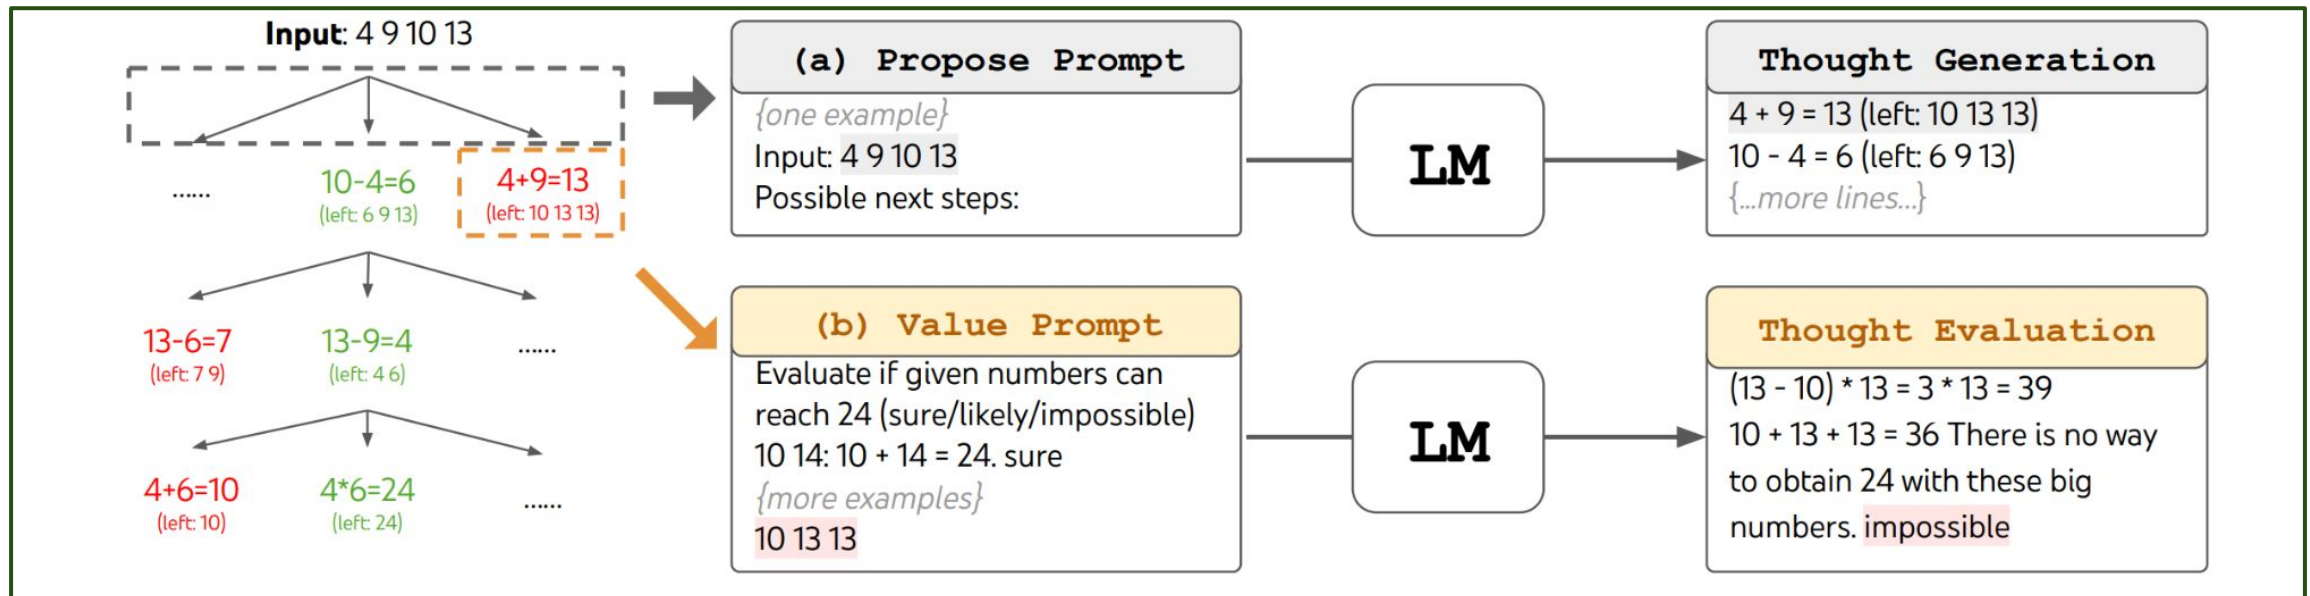
**Aim:** Improve the ability of LLM in deliberate decision making by considering multiple different reasoning paths  
⇒ Model generates and evaluate thoughts, and search algorithms used to explore thoughts with lookahead and backtracking.





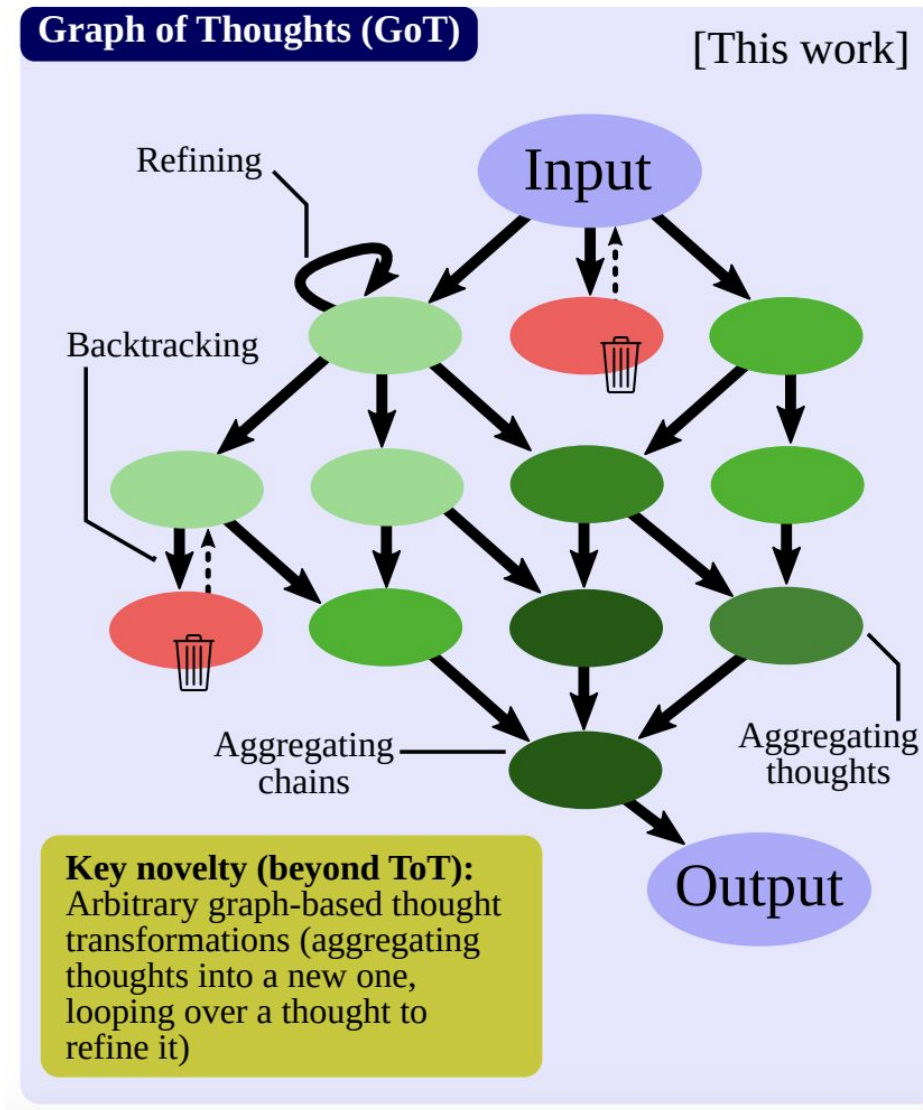
# Tree of Thoughts (ToT) Prompts

ToT for a game of 24 where the goal is to use 4 numbers and basic arithmetic operations (+-\*/ ) to obtain 24.

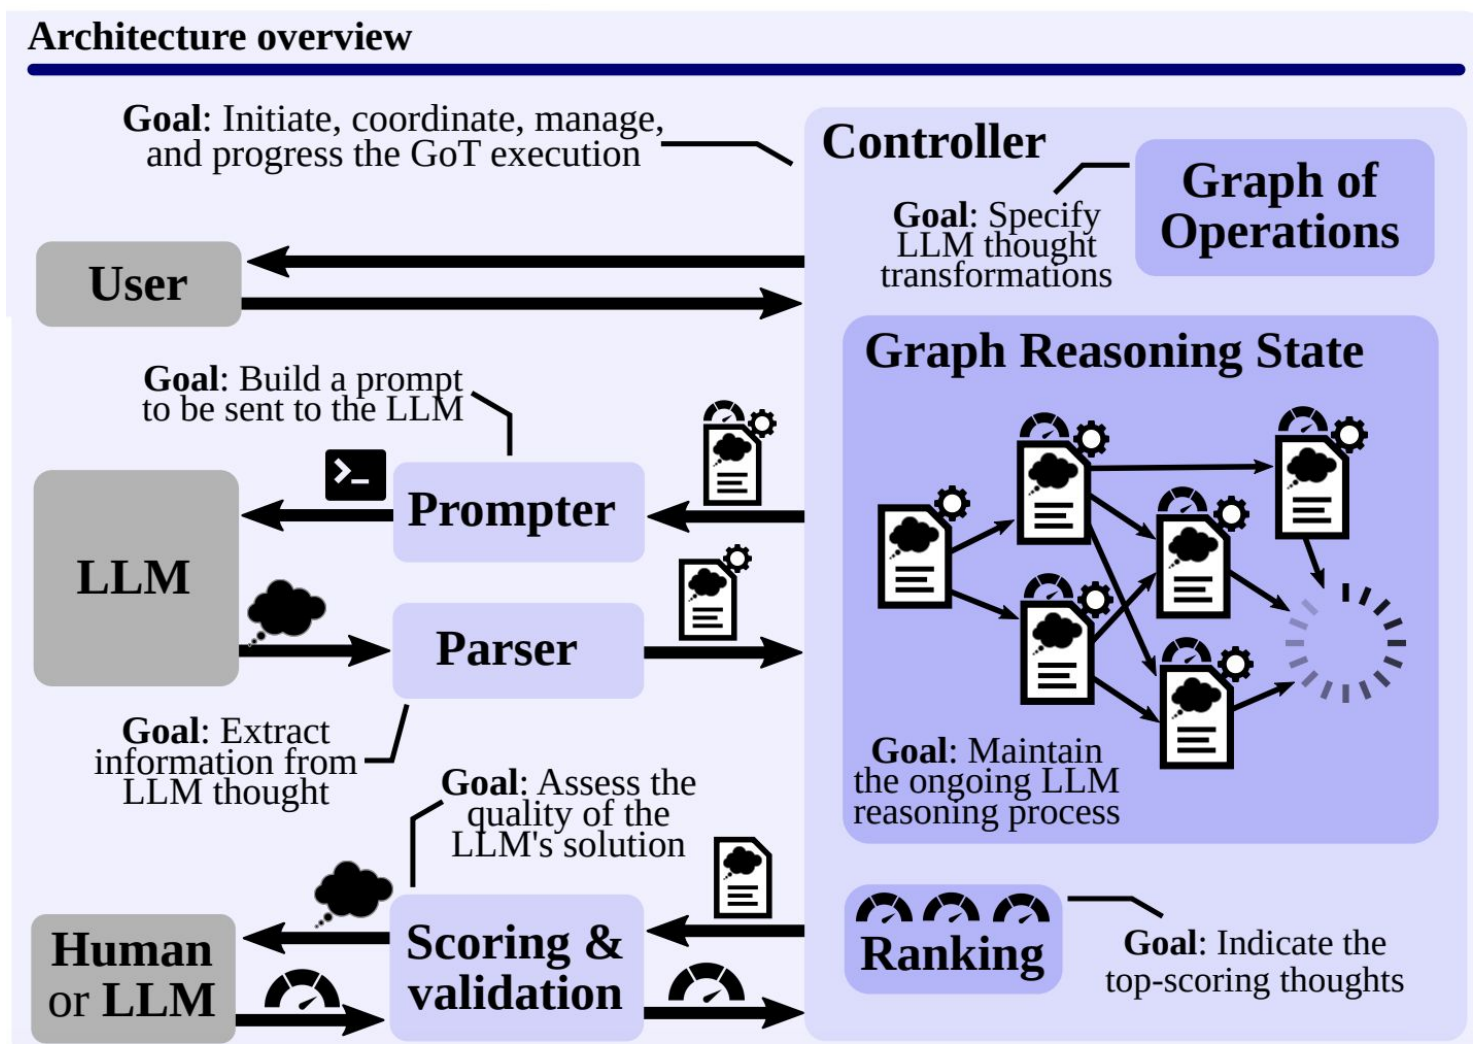


# Graph of Thoughts (GoT) Prompts

**Aim:** Solve complex problems by modeling them as a Graph of Operations (GoO), which is automatically executed with an LLM as the engine



# Graph of Thoughts (GoT) Prompts



# Cross-Lingual-Thought Prompting

**Aim:** Improve the ability of LLM in performing tasks for multilingual inputs.  
⇒ Create a prompt that uses both CoT (step-by-step) and asks the model to translate the input instruction/sample to English.

**XLT**

I want you to act as an arithmetic reasoning expert for Chinese.

Request: 詹姆斯决定每周跑 3 次 3 段冲刺, 每段冲刺跑 60 米。他每周一共跑多少米?

You should retell the request in English.

You should do step-by-step answer to obtain a number answer.

You should step-by-step answer the request.

You should tell me the answer in this format 'Answer:'.

I want you to act as a **task\_name** expert for **task\_language**.

**task\_input**

You should retell/repeat the **input\_tag** in English.


You should **task\_goal**.

You should step-by-step answer the request.

You should tell me the **output\_type** (**output\_constraint**) in this format '**output\_type**'.



# Cross-Lingual-Thought Prompting



Comparing the effectiveness of the **Cross-Lingual-Thought prompt** versus the baseline **basic prompt**



# Cross-Lingual CoT Prompting

## Cross-lingual Alignment Prompting



Please act as an expert in multi-lingual understanding in **Chinese**.  
Request: 玛丽比琼小两岁, 琼比杰莎大五岁。如果杰莎 20 岁, 这三个女孩的年龄总和是多少?  
Let's understand the task in **English** step-by-step!



- "玛丽比琼小两岁" means "Mary is two years younger than Joan."
- "琼比杰莎大五岁" means "Joan is five years older than Jessie."
- "如果杰莎 20 岁" means "If Jessie is 20 years old. "...

## Task-specific Solver Prompting



After understanding, you should act as an expert in arithmetic reasoning in **English**.  
Let's resolve the task you understand above step-by-step!  
Finally, you should format your answer as 'Answer: [num]'.

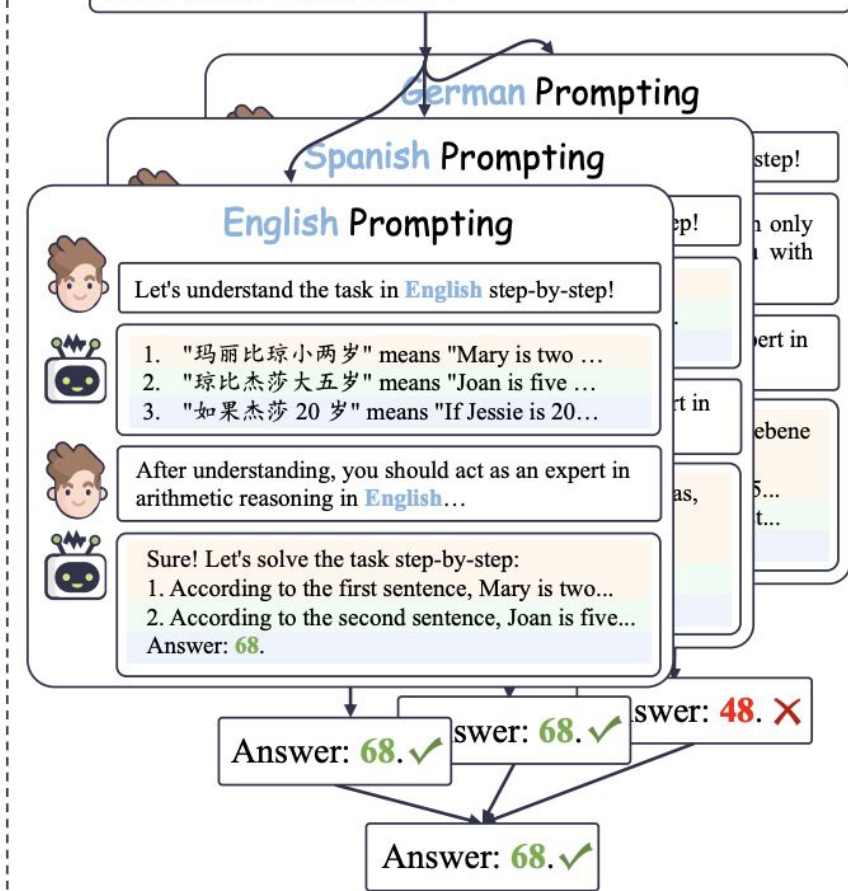


Sure! Let's solve the task step-by-step:  
1. According to the first sentence, Mary is two years younger than Joan. Let's suppose Joan's age is  $J$ , then Mary's age is  $J - 2$ .  
2. According to the second sentence, Joan is five years older than Jessie. If Jessie's age is  $x$ , then Joan's age is  $x + 5$ ...  
Answer: **68**. ✓

(a) Cross-lingual Prompting (CLP)



Please act as an expert in multi-lingual understanding in **Chinese**.  
Request: 玛丽比琼小两岁, 琼比杰莎大五岁。如果杰莎 20 岁, 这三个女孩的年龄总和是多少?

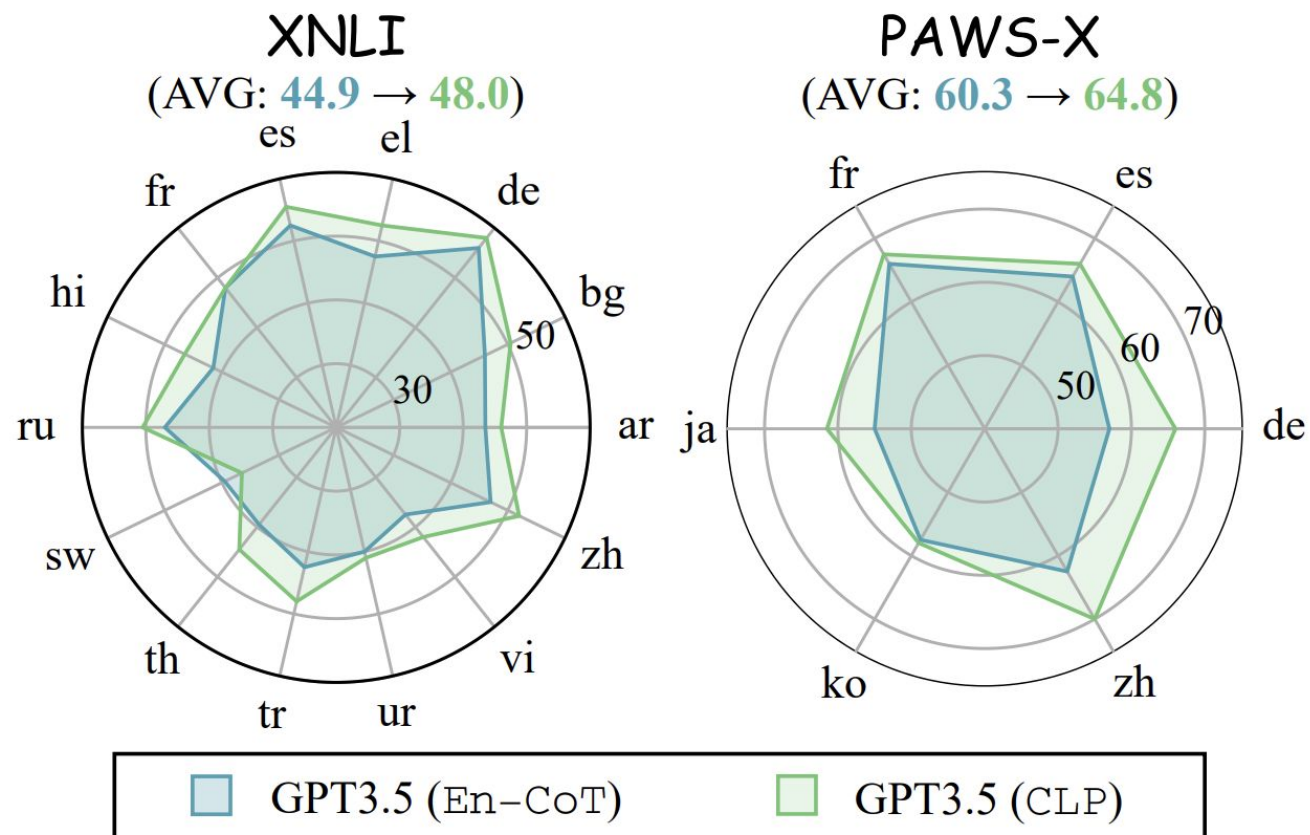


(b) Cross-lingual Self-consistent Prompting (CLSP)



# Cross-Lingual CoT Prompting

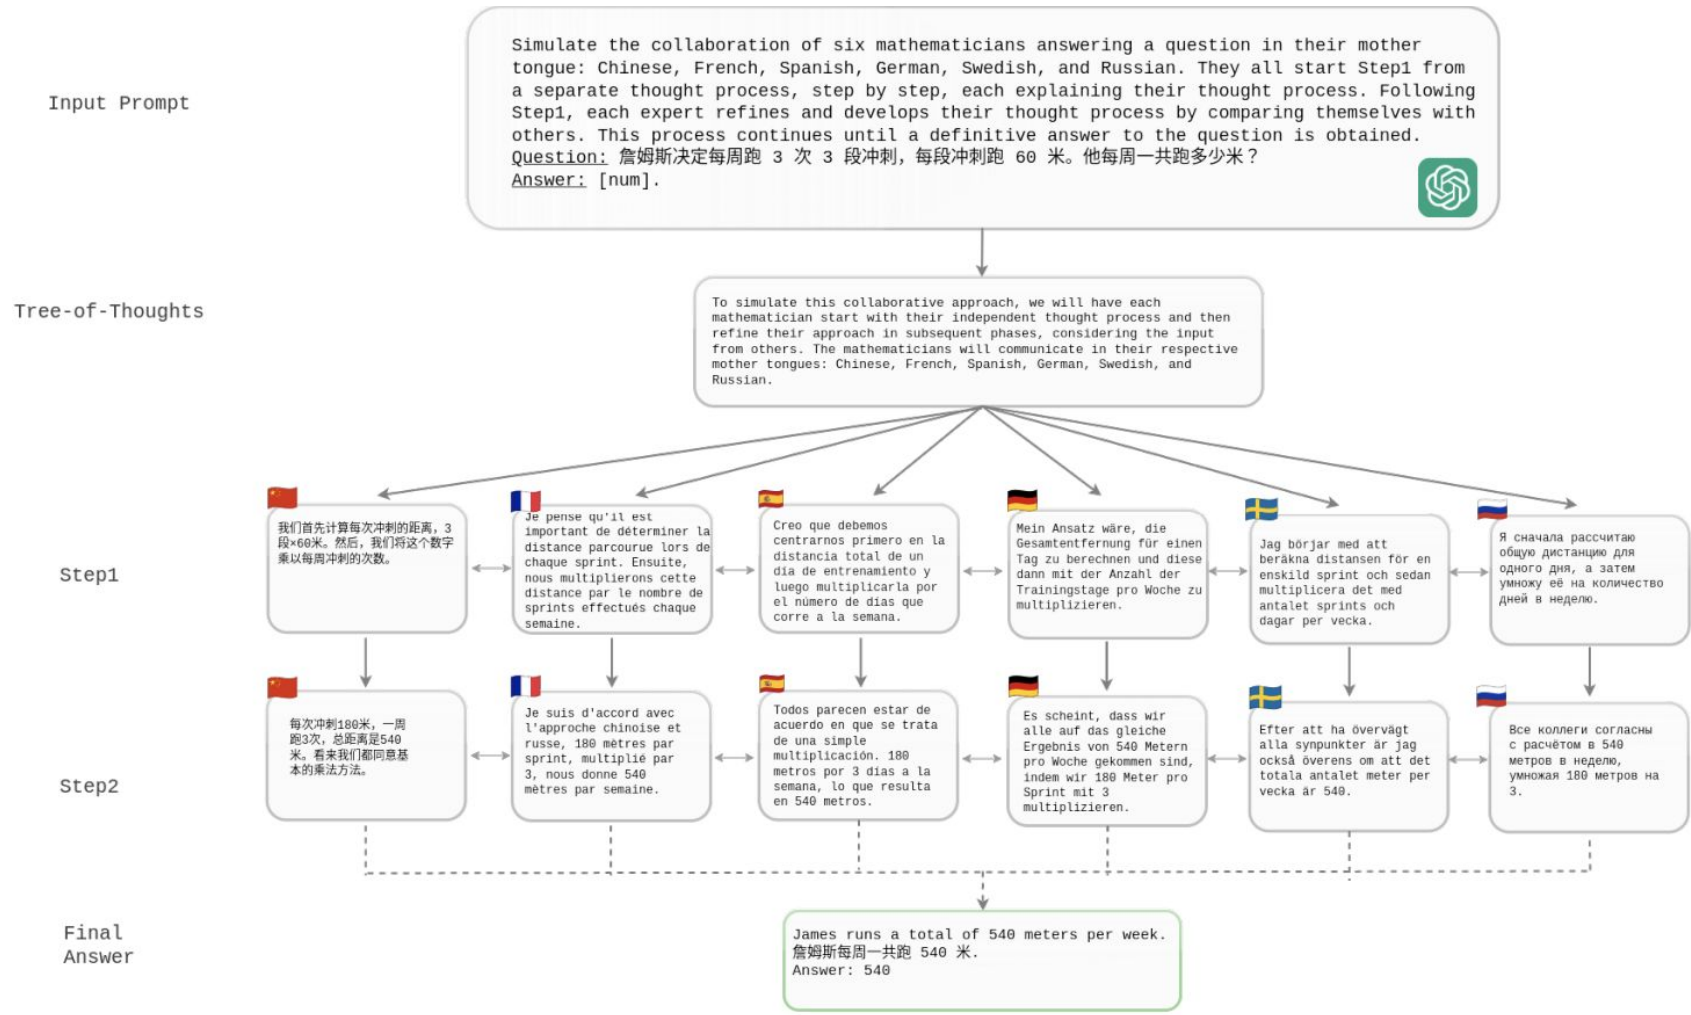
Accuracy across languages in two tasks: XNLI and PAWS-X



# Cross-lingual ToT (Cross-ToT) Prompts

**Aim:** Improve the ability of LLM in deliberate decision making across languages by considering multilingual reasoning paths.

⇒ Use **ToT** style prompting to ask the LLM to deliver the reasoning process in different languages that, step-by-step, converge to a single final solution



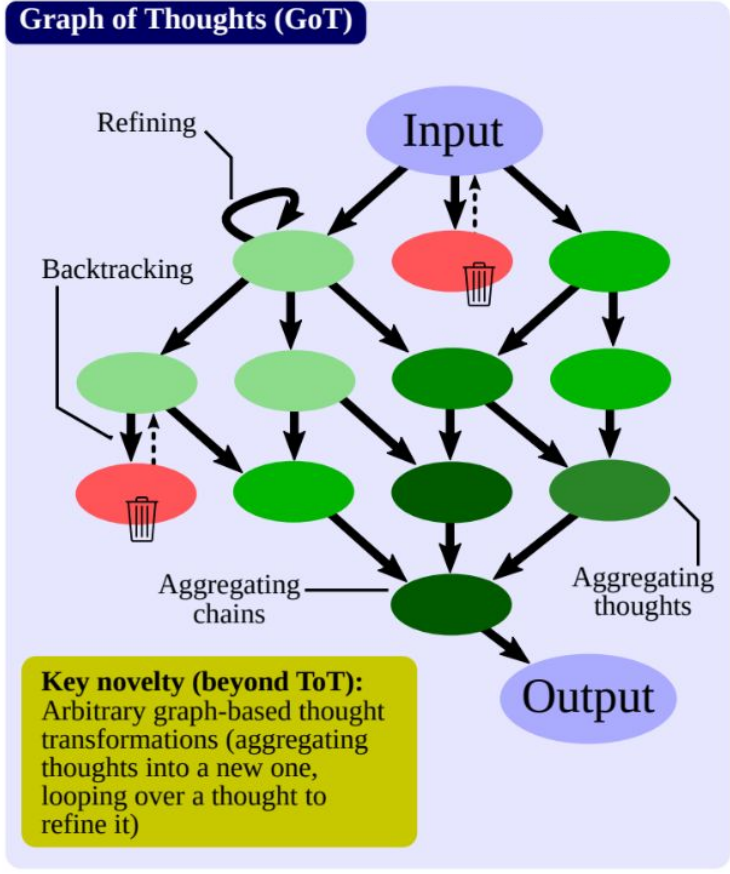
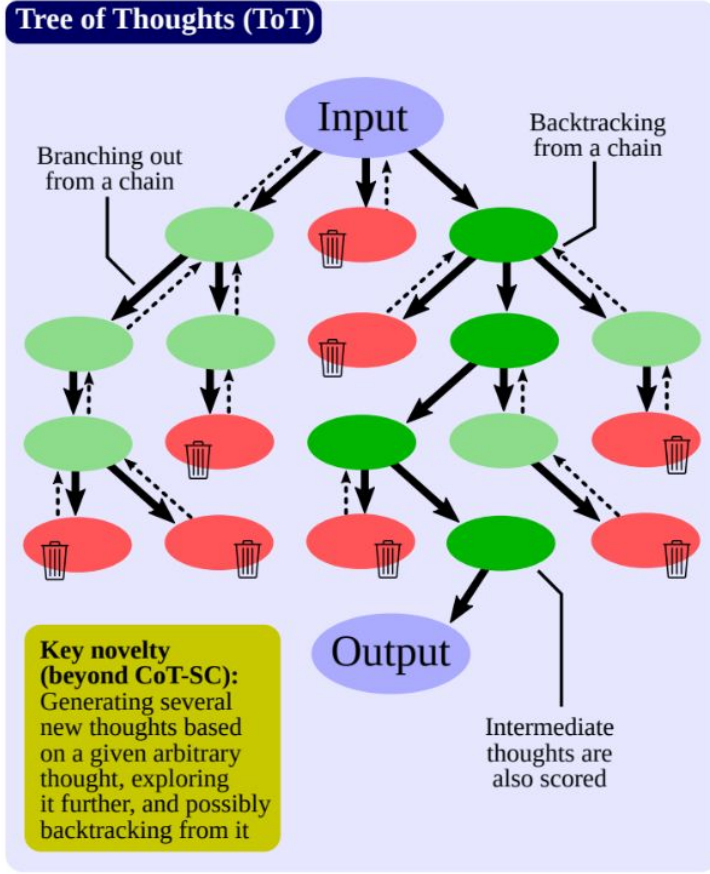
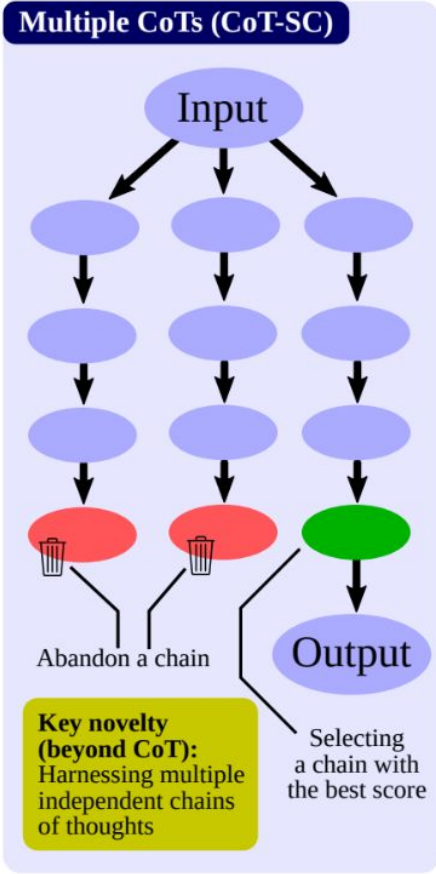



# Cross-lingual ToT Prompts

| Model                            | de          | zh          | fr          | ru          | sw          | es          | Average      |
|----------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------|
| <b>GPT-3 (text-davinci-002)*</b> |             |             |             |             |             |             |              |
| Direct (Shi et al., 2022)        | 14.8        | 18.0        | 16.8        | 12.4        | 8.8         | 17.2        | 14.67        |
| Native-CoT (Shi et al., 2022)    | 36.0        | 40.0        | 37.6        | 28.4        | 11.2        | 40.4        | 32.27        |
| En-CoT (Shi et al., 2022)        | 44.0        | 40.8        | 46.0        | 28.4        | 20.8        | 44.8        | 37.47        |
| Translate-En (Shi et al., 2022)  | 46.4        | 47.2        | 46.4        | 48.8        | 37.6        | 51.6        | 46.33        |
| <b>GPT-3.5 (gpt-3.5-turbo)</b>   |             |             |             |             |             |             |              |
| Direct (Qin et al., 2023)        | 56.0        | 60.0        | 62.0        | 62.0        | 48.0        | 61.2        | 58.20        |
| Native-CoT (Qin et al., 2023)    | 70.0        | 59.6        | 64.4        | 62.4        | 54.0        | 70.4        | 63.47        |
| En-CoT (Qin et al., 2023)        | 73.6        | 63.2        | 70.0        | 65.6        | 55.2        | 69.6        | 66.20        |
| Translate-En (Qin et al., 2023)  | 75.6        | 71.6        | 72.4        | 72.8        | 69.6        | 74.4        | 72.73        |
| Cross-CoT (Qin et al., 2023)     | 86.8        | 77.2        | 82.0        | <b>87.6</b> | <b>76.0</b> | 84.8        | 82.40        |
| <b>Cross-ToT</b>                 | <b>87.6</b> | <b>83.5</b> | <b>84.3</b> | 86.5        | 75.4        | <b>86.2</b> | <b>83.91</b> |

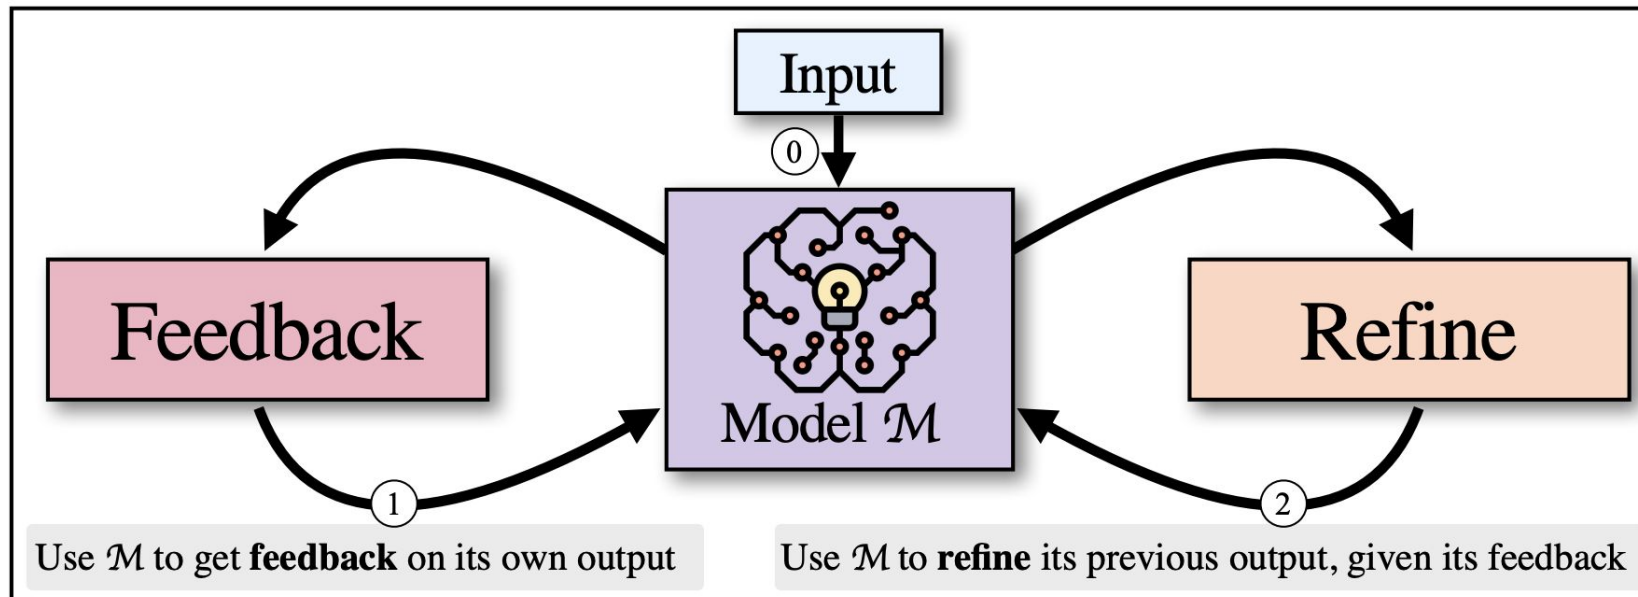


# Comparing Prompting Techniques



# Iterative Prompting

**Aim:** Improve LLM performance by iteratively prompting it to refine its previous responses.



# Iterative Prompting

**Self-refine technique:** Prompt the same LLM iteratively with three prompts (for initial generation, feedback on generation, and refinement)

(a) **Dialogue:**  $x, y_t$

User: I am interested in playing Table tennis.

Response: I'm sure it's a great way to socialize, stay active

(b) **FEEDBACK**  $fb$

Engaging: Provides no information about table tennis or how to play it.

User understanding: Lacks understanding of user's needs and state of mind.

(c) **REFINE**  $y_{t+1}$

Response (refined): That's great to hear (...) ! It's a fun sport requiring quick reflexes and good hand-eye coordination. Have you played before, or are you looking to learn?





# Iterative Prompting

| Task                   | GPT-3.5     |                     | CHATGPT |                     | GPT-4 |                     |
|------------------------|-------------|---------------------|---------|---------------------|-------|---------------------|
|                        | Base        | +SELF-REFINE        | Base    | +SELF-REFINE        | Base  | +SELF-REFINE        |
| Sentiment Reversal     | 8.8         | <b>30.4</b> (↑21.6) | 11.4    | <b>43.2</b> (↑31.8) | 3.8   | <b>36.2</b> (↑32.4) |
| Dialogue Response      | 36.4        | <b>63.6</b> (↑27.2) | 40.1    | <b>59.9</b> (↑19.8) | 25.4  | <b>74.6</b> (↑49.2) |
| Code Optimization      | 14.8        | <b>23.0</b> (↑8.2)  | 23.9    | <b>27.5</b> (↑3.6)  | 27.3  | <b>36.0</b> (↑8.7)  |
| Code Readability       | 37.4        | <b>51.3</b> (↑13.9) | 27.7    | <b>63.1</b> (↑35.4) | 27.4  | <b>56.2</b> (↑28.8) |
| Math Reasoning         | <b>64.1</b> | <b>64.1</b> (0)     | 74.8    | <b>75.0</b> (↑0.2)  | 92.9  | <b>93.1</b> (↑0.2)  |
| Acronym Generation     | 41.6        | <b>56.4</b> (↑14.8) | 27.2    | <b>37.2</b> (↑10.0) | 30.4  | <b>56.0</b> (↑25.6) |
| Constrained Generation | 16.0        | <b>39.7</b> (↑23.7) | 2.75    | <b>33.5</b> (↑30.7) | 4.4   | <b>61.3</b> (↑56.9) |



# Automated Prompt Engineering

- **Prompt Mining**

- Scrape a large text corpus (e.g., Wikipedia) for strings containing  $x$  and  $y$ , and finds either the middle words or dependency paths between the inputs and outputs.

- **Prompt Paraphrasing**

- Take a seed prompt and paraphrase it into candidate prompts, then select the one that achieves the highest accuracy on the target task.

- **Prompt Generation**

- Generate instruction candidates through an LLM for a task given output examples and select the most appropriate instruction based on computed evaluation scores.



# In-Context/Few-shot Learning

# Zero- vs. Few-shot Prompts

**Classify the following sentence by the sentiment it expresses given these sentiments: Positive, Negative, Neutral, or Mixed.**

**Sentence:** perfectly executed and wonderfully sympathetic characters  
**Sentiment:**

**Classify the following sentence by the sentiment it expresses given these sentiments: Positive, Negative, Neutral, or Mixed. Here are some examples:**

**Sentence:** a host of splendid performances  
**Sentiment:** Positive

**Sentence:** felt trapped and with no obvious escape  
**Sentiment:** Negative

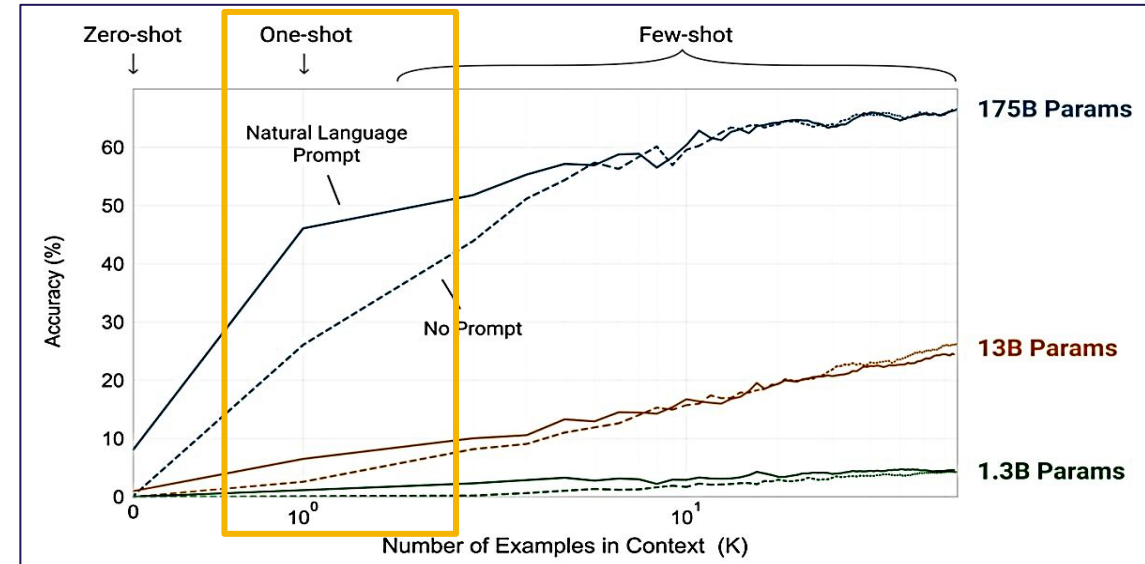
**Sentence:** perfectly executed and wonderfully sympathetic characters  
**Sentiment:**





# Why?

- Improved performance over zero-shot
- Smaller task-specific dataset required (vs. fine-tuning)
- Model isn't updated, only pass the examples at inference time

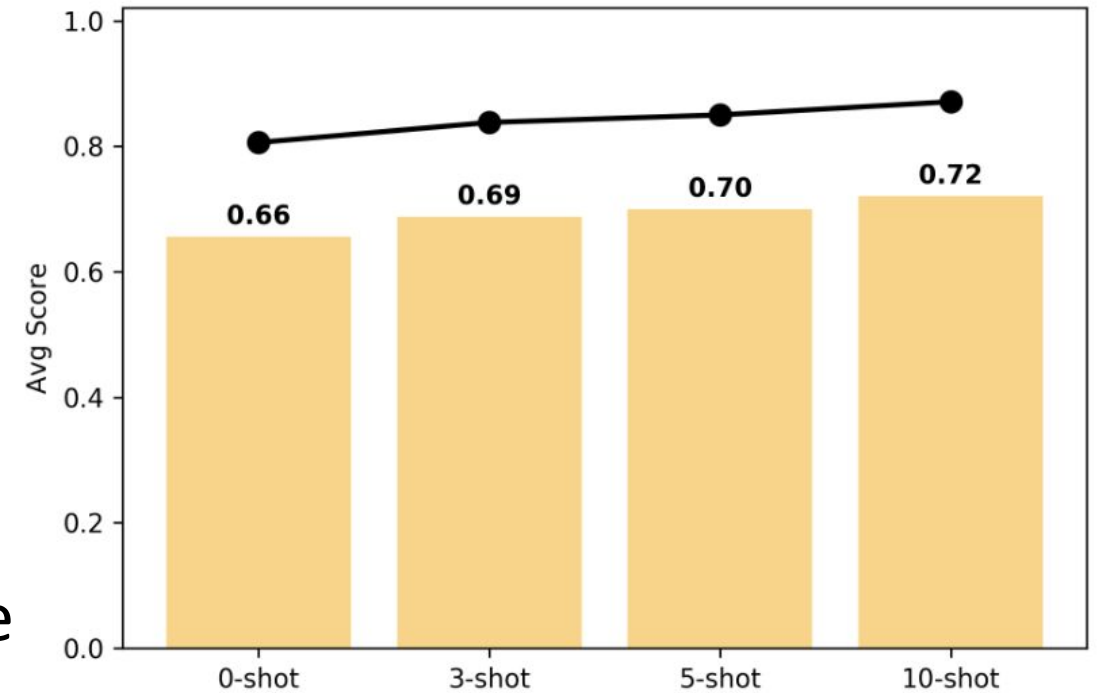


(Brown et al., 2020)



# How Many Examples?

- Great range of values: [1,2,3,...,48,...]
- Consider document/example length:
  - LLMs have a fixed context window (e.g. GPT-3.5 allows 4,097 tokens as input)
- Tune as hyperparameter on development set



(Abdelali et al., 2024)



# Which Examples?

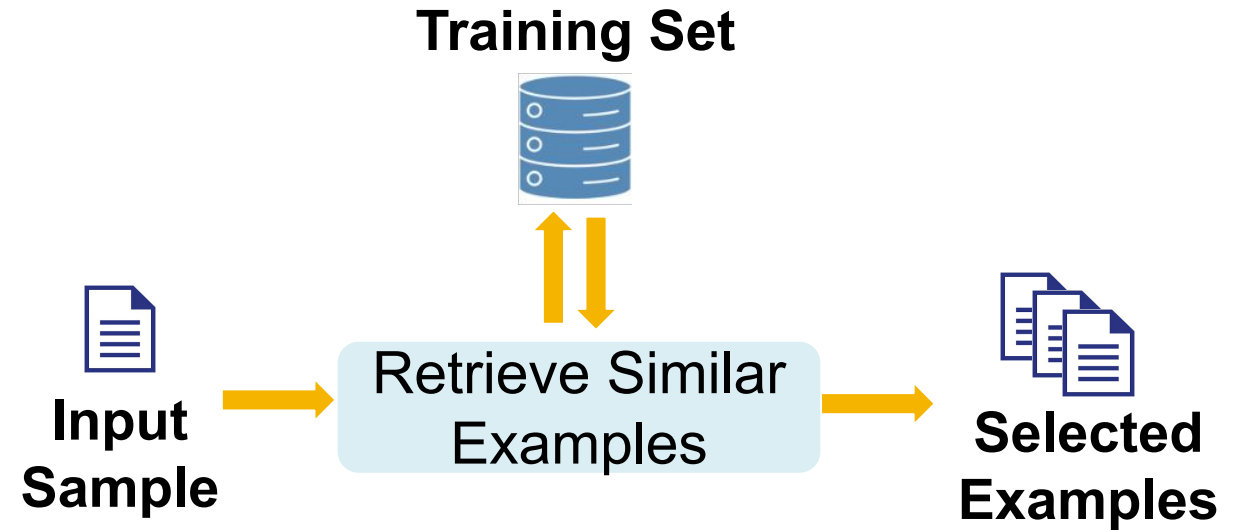
## Manual

- Select some examples manually

## Sampling

- Uniform class distribution
- Randomly  
⇒ Might lead to skewed label distribution

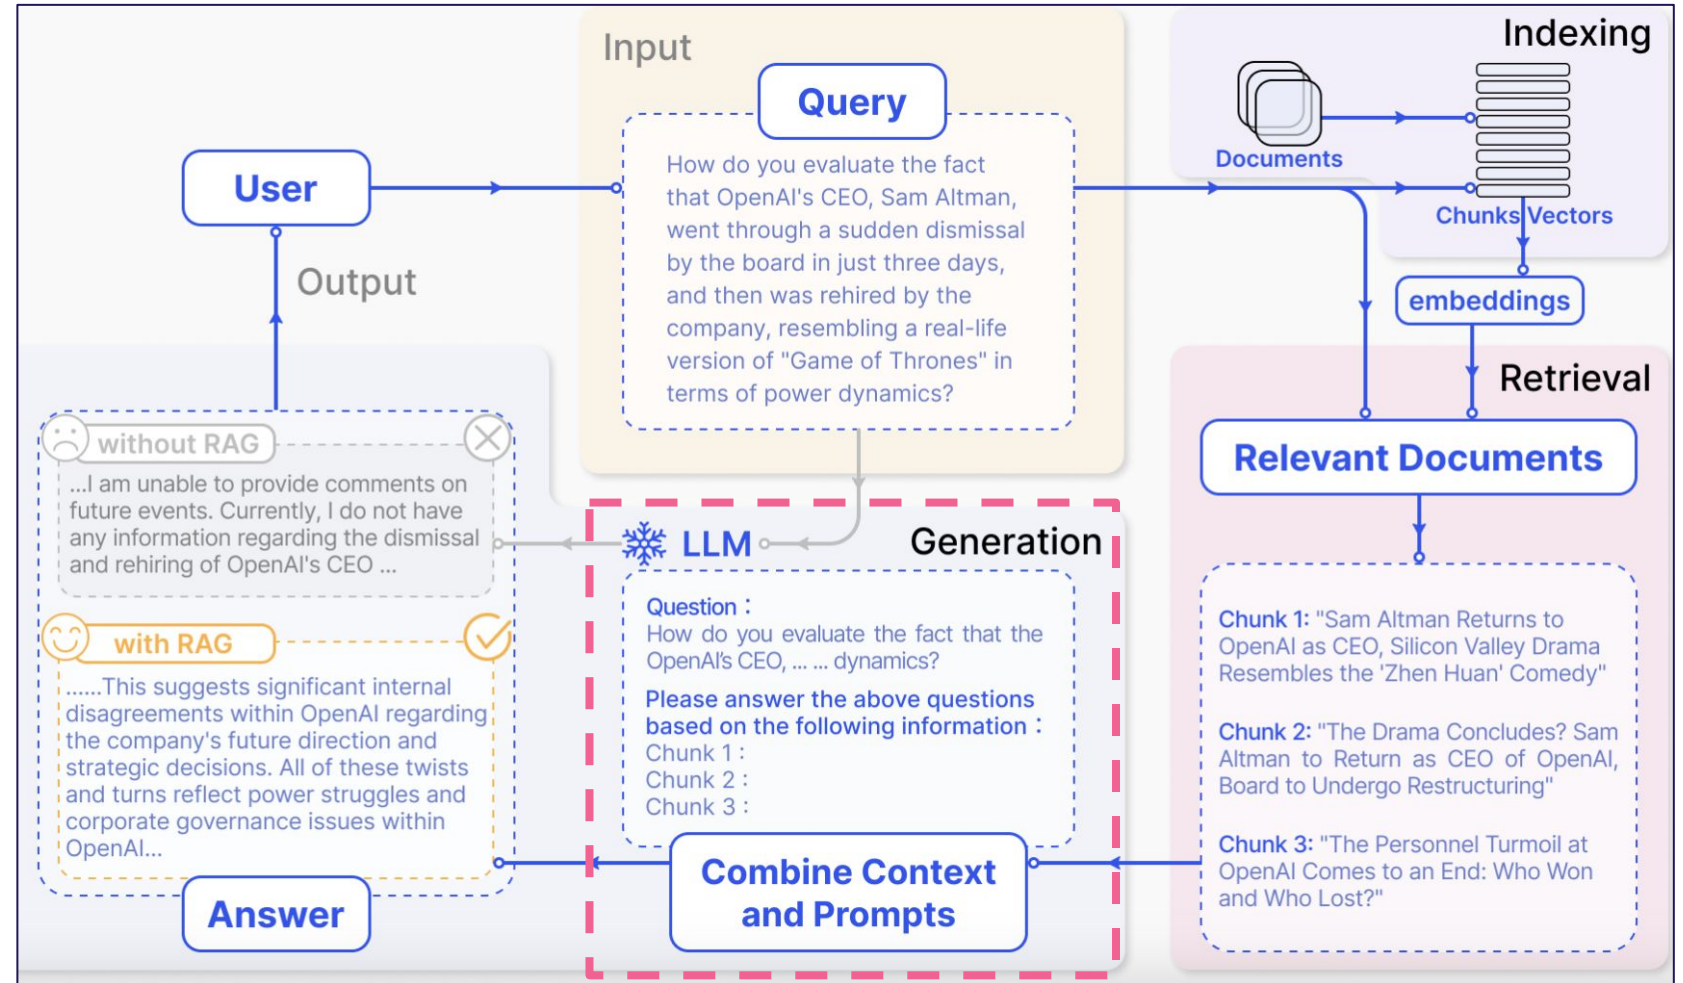
## Semantic Similarity



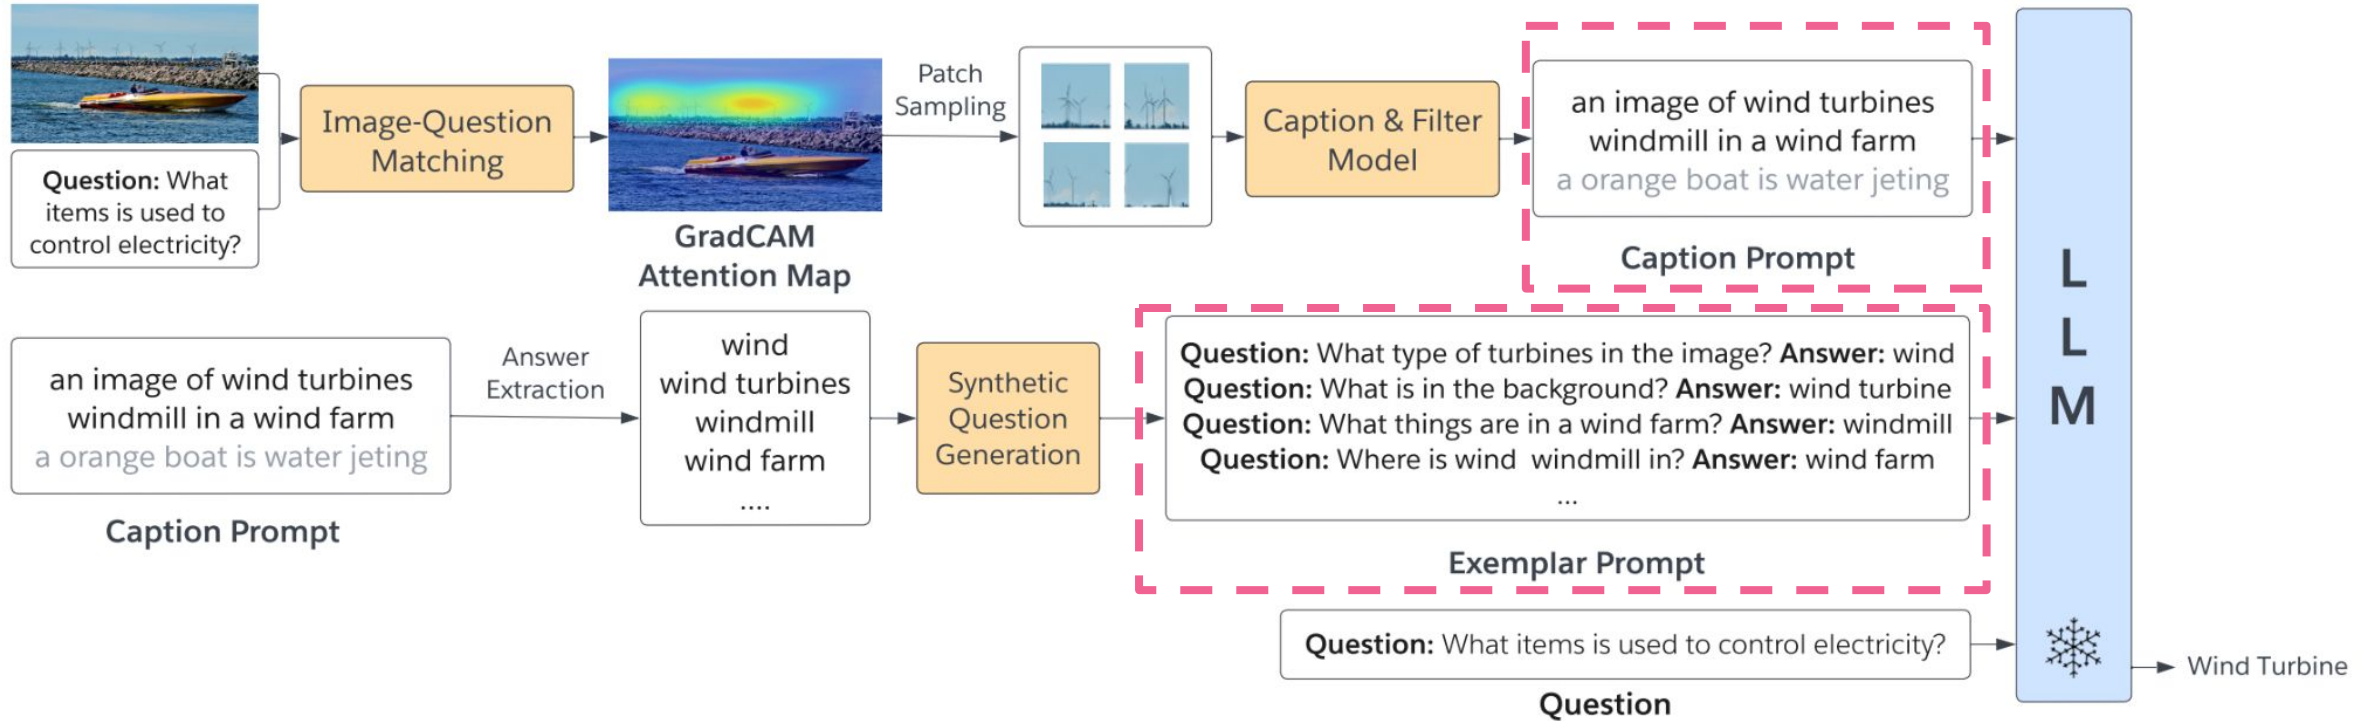
- Cosine similarity
- Word overlap
- Maximal marginal relevance

# Retrieval-Augmented Generation (RAG)


**Aim:** Provide additional context for the LLM, leading to improved factual accuracy and coherence in its output.



# Context for Tasks on Images



# Mono-/Cross- Language Prompting



- **Monolingual Prompting:** Few shot examples + test sample in **same language**.
- **Zero-Shot Cross-Lingual:** Few shot **English** examples + test sample in different language.
- **Translate-Test:** Few shot English examples + test sample **translated** to English.





# Mono-/Cross- Language Prompting

Classify the 'sentence' as subjective or objective. Provide only label.

sentence: "والصحيح هو أن السيد أحمد منصور له  
"مواقف ضد الفكر السلفي"

label:

صنف "الجملة" إلى لاموضوعية أو موضوعية.

الجملة: "والصحيح هو أن السيد أحمد منصور له مواقف ضد  
الفكر السلفي."  
التصنيف:

| Task Name      | Metric           | English      | Arabic       |
|----------------|------------------|--------------|--------------|
| NER            | Macro-F1         | 0.355        | 0.350        |
| Sentiment      | Macro-F1         | 0.569        | 0.547        |
| News Cat.      | Macro-F1         | 0.667        | 0.739        |
| Gender         | Macro-F1         | 0.868        | 0.892        |
| Subjectivity   | Macro-F1         | 0.677        | 0.725        |
| XNLI (Arabic)  | Acc              | 0.753        | 0.740        |
| QA             | F1 (exact match) | 0.705        | 0.654        |
| <b>Average</b> |                  | <b>0.656</b> | <b>0.664</b> |



# Prompting and Benchmarking Tools



# Prompting and Benchmarking Tools

- **Prompt Source** (Bach et al. 2022)
- **LLMeBench** (Dalvi et al., 2023)
- **Im-evaluation-harness** (Gao et al., 2023)
- **Open ICL** (Wu et al., 2023)
- **Prompt Bench** (Zhu et al., 2023)



# Prompt Source

*“a system for creating, sharing, and using natural language prompts”*



### S1: Exploration

**Browse**  
SNLI

The SNLI corpus (version 1.0) is a collection of 570k human-written English sentence pairs manually labeled for the task of NLI...

```
{ premise: "A person...",  
  hypothesis: "A person...",  
  label: 1 }
```

```
{ premise: "The kids...",  
  hypothesis: "All kids...",  
  label: 2 }
```

**S1: Exploration**

### S2, S3, S4: Creation

**Sourcing**  
SNLI

based on the previous passage

Adapted from the BoolQ prompts in Schick & Schütze 2021.

Original Task  Choices in Prompt

Yes ||| No ||| Maybe Accuracy

```
{{premise}} Based on the  
previous passage, is it true  
that "{{hypothesis}}"?  
Yes, no, or maybe? |||  
{{ answer_choices[label] }}
```

**S2: Writing** **S3: Documentation**

### S5: Review

**Browse**  
SNLI  
Based...

The SNLI corpus (version 1.0) is a collection of 570k human-written English sentence pairs manually labeled for the task of NLI...

"A person..." Based on the previous passage, is it true that "A person..."? Yes, no, or maybe? ||| Maybe

"The kids..." Based on the previous passage, is it true that "All kids..."? Yes, no, or maybe? ||| No

**S5: Review**

<https://youtu.be/glthK9J52IM?feature=shared>

<https://github.com/bigscience-workshop/promptsources>



# Prompt Source

## Five stages of creating prompts:

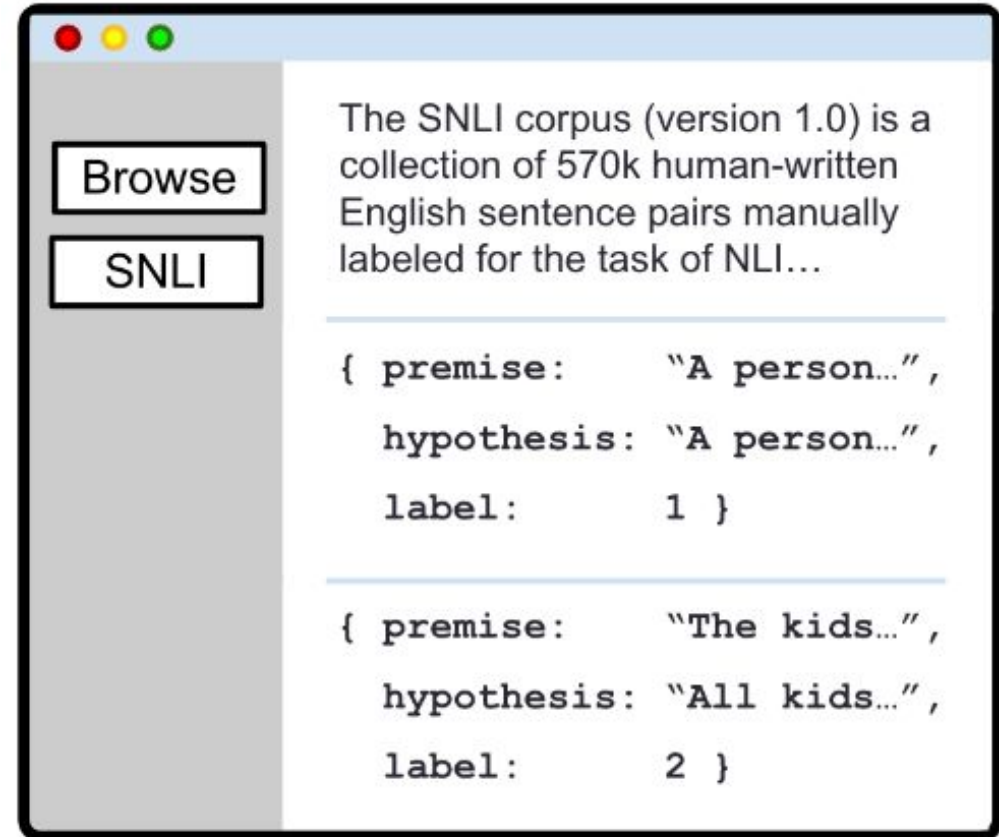
### S1: Dataset Exploration

#### SNLI dataset example:

Assume a given premise sentence is true, the goal is to determine whether a hypothesis sentence is:

- true (entailment),
- false (contradiction),
- or undetermined (neutral)

### S1: Exploration



# Prompt Source

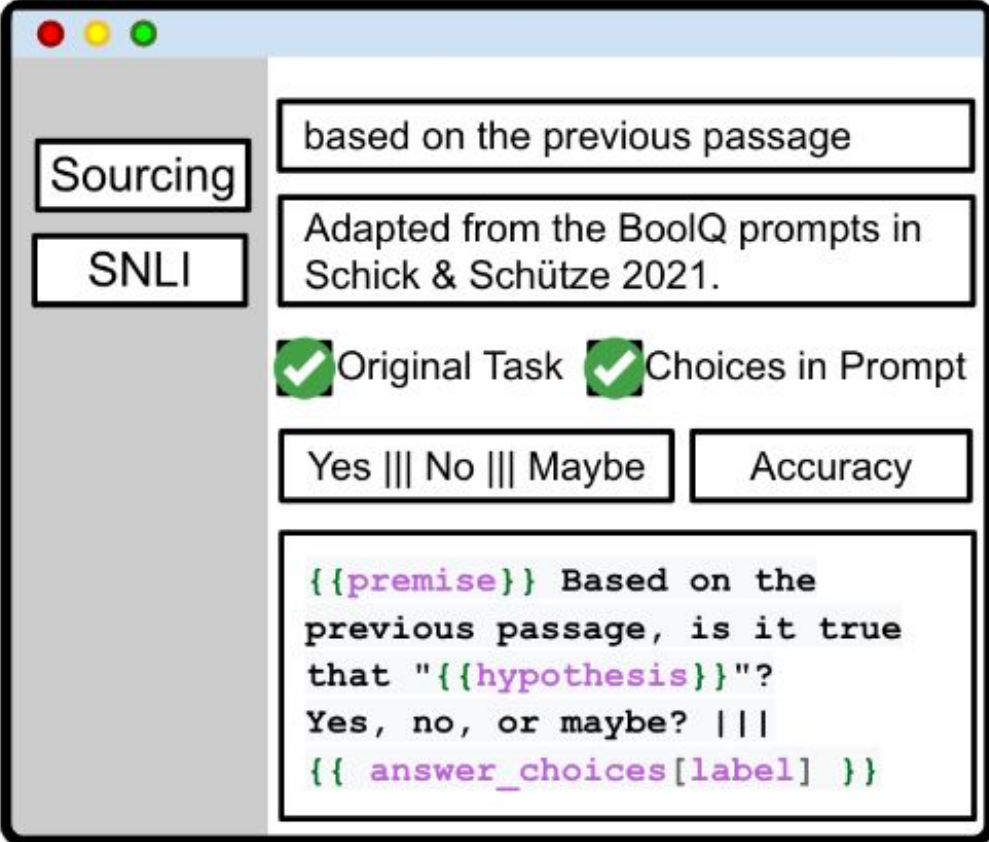
## Five stages of creating prompts:

**S2:** Prompt Writing

**S3:** Prompt Documentation

**S4:** Iteration and Variation

## S2 + S3 + S4: Creation



The screenshot shows a web interface for creating prompts. On the left, there is a sidebar with two buttons: "Sourcing" and "SNLI". The main content area is titled "S2 + S3 + S4: Creation" and contains the following elements:

- A text box containing "based on the previous passage".
- A text box containing "Adapted from the BoolQ prompts in Schick & Schütze 2021."
- Two checked checkboxes: "Original Task" and "Choices in Prompt".
- A text box containing "Yes ||| No ||| Maybe".
- A text box containing "Accuracy".
- A code editor showing the following prompt template:

```
{{premise}} Based on the previous passage, is it true that "{{hypothesis}}"?  
Yes, no, or maybe? |||  
{{ answer_choices[label] }}
```

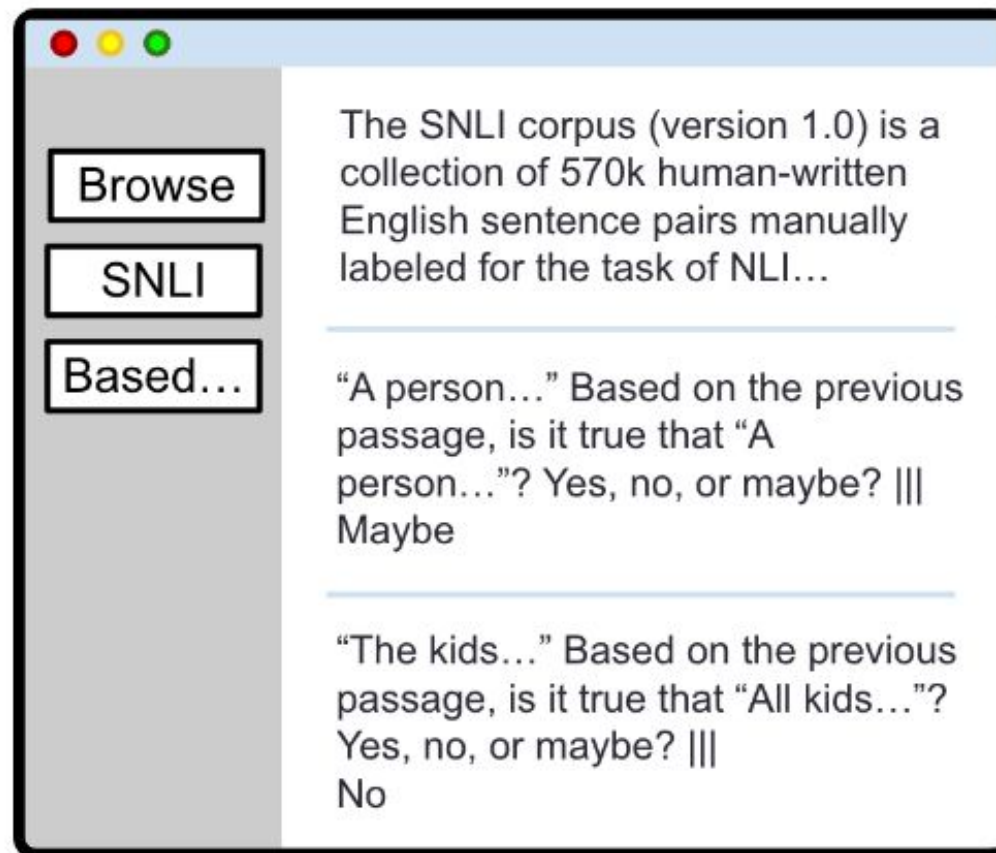


# Prompt Source

## Five stages of creating prompts:

### S5: Global Review

### S5: Review



# Prompt Source

## Prompt Template Creation

```
Dataset ? Input
# Load an example from the datasets ag_news
>>> from datasets import load_dataset
>>> dataset = load_dataset("ag_news", split="train")
>>> example = dataset[1]

# Load prompts for this dataset
>>> from promptsource.templates import DatasetTemplates
>>> ag_news_prompts = DatasetTemplates('ag_news')

# Print all the prompts available for this dataset. The keys of the dict are the UUIDs the u
>>> print(ag_news_prompts.templates)
{'24e44a81-a18a-42dd-a71c-5b31b2d2cb39': <promptsource.templates.Template object at 0x7fa7ac


# Select a prompt by its name
>>> prompt = ag_news_prompts["classify_question_first"]

# Apply the prompt to the example
>>> result = prompt.apply(example)
>>> print("INPUT: ", result[0])
INPUT: What label best describes this news article?
Carlyle Looks Toward Commercial Aerospace (Reuters) Reuters - Private investment firm Carlyl
>>> print("TARGET: ", result[1])
TARGET: Business
```





# LLMeBench



Make it super-simple and quick to **start experimenting** with LLMs,  
and **easily transfer that effort** to large scale  
evaluation

<http://llmebench.qcri.org/>



# LLMeBench: Usecases

## Exploration

Try a model with different prompts over the same dataset

## Model comparison

Run the same prompt with multiple models

## Benchmarking suite

Create a suite of tasks and datasets and track a model's progress across all

## Many more...

Framework is flexible and extensible for new tasks, datasets, and models



# Why LLMeBench?

1. Read the data
2. Figure out how to access an LLM (e.g. GPT4)
3. Understand and write code to read the response
4. Explore with different prompts
5. Write some sort of loop over the data and prompts to see model responses on all samples
  - a. Realize the request fails for many reasons  $\Rightarrow$  Write some code to retry failed requests
  - b. Realize every time you run your code, you get different results  $\Rightarrow$  Modify code to set all appropriate model parameters for reproducible results
  - c. Have an idea for a new prompt, figure out changing existing code to only run for new prompt while keeping results from older prompts
6. Process results
7. Rinse and Repeat for a new problem/dataset/task

Current LLM usage and benchmarking process



# Why LLMeBench?

1. Find your task, dataset and model in LLMeBench
  - ⇒ Task/Data/Model not found?
    - a. Edit existing task/data/model script for your needs
2. Run experiment!

Add a layer of abstraction so that you as a user can focus solely on getting the best performance out of the LLM



# LLMeBench

*benchmarking asset*

```
def config():
    return {
        "dataset": TSVDataset,
        "dataset_args": {
            "column_mapping": {
                "input": "sentence",
                "label": "labels",
            },
        },
        "task": ClassificationTask,
        "model": FastChatModel,
        "general_args": {"custom_test_split": "SST-2/dev.tsv"},
    }

def prompt(input_sample):
    return [
        {"role": "system", "content": "You are an expert in sentiment analysis."},
        {"role": "user", "content": f"Sentence: {input_sample}\nSentiment:"}
    ]

def post_process(response):
    out = response["choices"][0]["message"]["content"].lower()
    return 1 if "positive" in out else 0
```



# LLMeBench

Once an *asset* is written, LLMeBench takes care of everything else!

```
python -m llmebench assets/ results/
```

```
{
  "num_processed": 872,
  "num_failed": 0,
  "evaluation_scores": {
    "Macro F1": 0.8586052694703862,
    "Micro F1": 0.8612385321100917,
    "Acc": 0.8612385321100917,
    "Weighted Precision": 0.8821528346701518,
    "Weighted Recall": 0.8612385321100917,
    "Weighted F1": 0.8589593215900104
  }
}
```



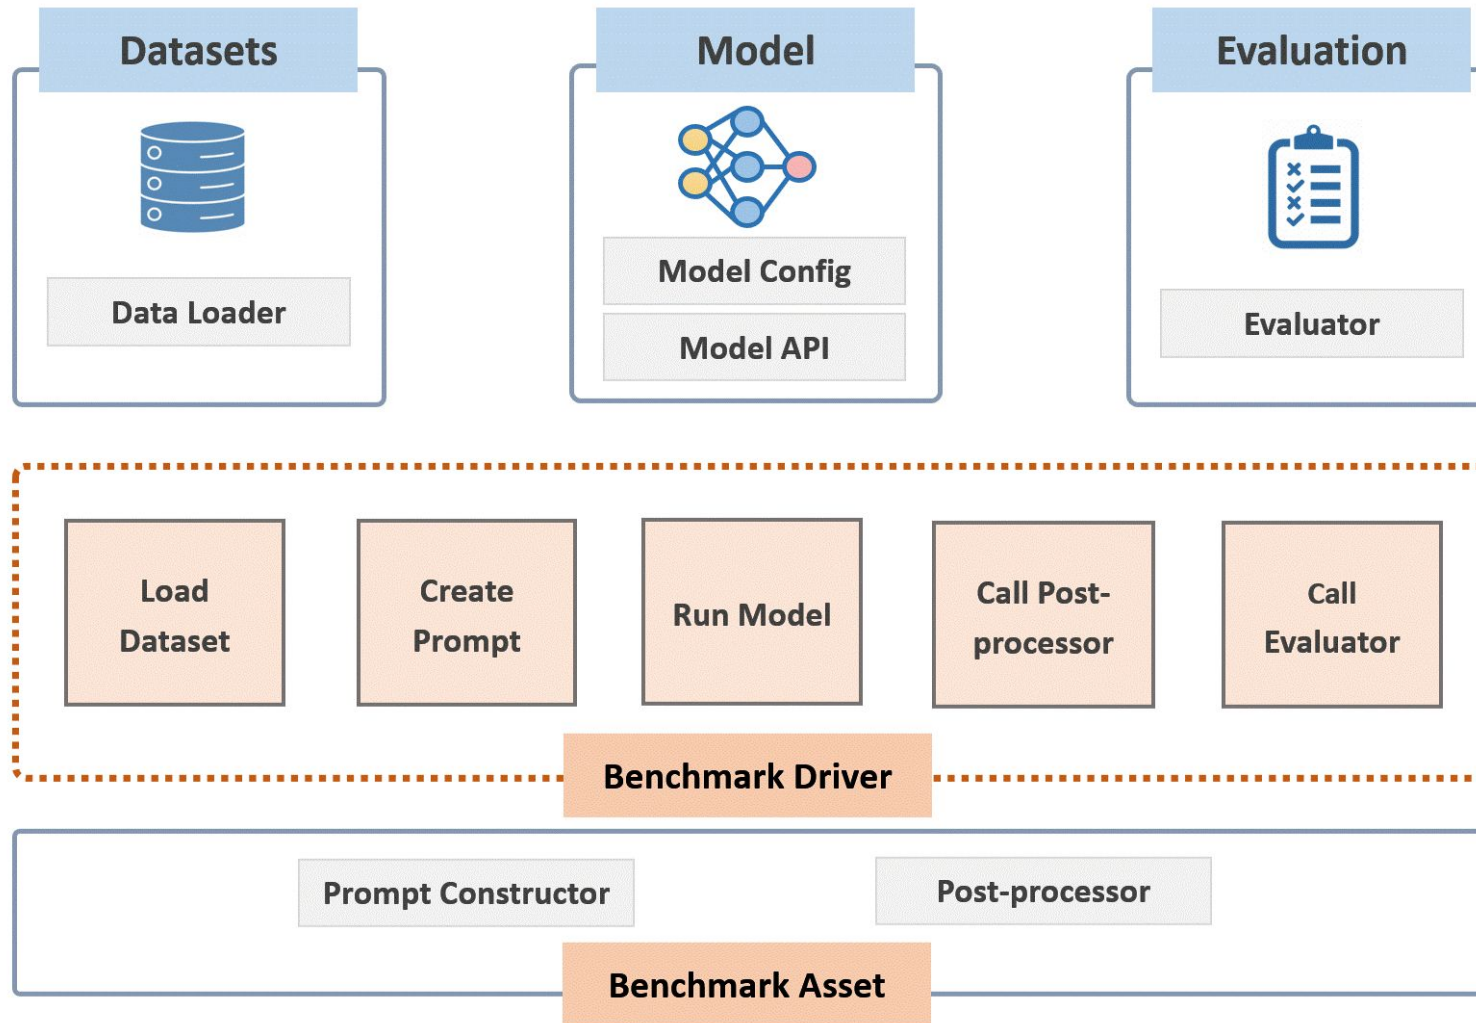


# LLMeBench Features

- ~**300 assets** across **12 languages**
- Extensive support for **reading datasets**
  - HuggingFace datasets + generic data loaders (csv, tsv, json)
  - Over 50 dataset-specific loaders
  - Automatic downloading of data (when allowed)
- Supports popular **task types** (Classification, regression etc.)
- Supports popular **model providers** (OpenAI, FastChat, Petals, HuggingFace Inference API)
- Extensive caching
- **Extensible and Plug-and-play!**
  - Easily add new datasets, tasks, evaluation metrics and model providers



# LLMeBench: Technical Overview



# Large Scale Experimentation Across:

| TASKS                                                                                                                                                       | DATASETS                                                                                                                                                                                                             | EVALUATION                                                                                                                                                                                                          | MODELS                                                                                   |          |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|----------|
| <ul style="list-style-type: none"> <li>Word Segmentation, Syntax &amp; Information Extraction (e.g., POS tagging)</li> </ul>                                | <ul style="list-style-type: none"> <li>XNLI</li> <li>XGLUE</li> <li>XQuAD</li> <li>ASAD</li> <li>Aqmar</li> <li>SANAD</li> <li>MADAR</li> <li>QASR</li> <li>WikiNews</li> <li>Conll2006</li> <li>ANERcorp</li> </ul> | <ul style="list-style-type: none"> <li>Accuracy</li> <li>F1</li> <li>Macro-F1</li> <li>Micro-F1</li> <li>Weighted-F1</li> <li>BLEU</li> <li>WER</li> <li>Pearson Correlation</li> <li>Jaccard Similarity</li> </ul> | <ul style="list-style-type: none"> <li>GPT-3.5</li> <li>GPT-4</li> <li>BLOOMZ</li> </ul> |          |
| <ul style="list-style-type: none"> <li>Factuality, Disinformation &amp; Harmful Content Detection (e.g., Hate Speech &amp; Propaganda Detection)</li> </ul> |                                                                                                                                                                                                                      |                                                                                                                                                                                                                     | <th>LEARNING</th>                                                                        | LEARNING |
| <ul style="list-style-type: none"> <li>Semantics (e.g., Semantic Textual Similarity and Natural Language Inference)</li> </ul>                              |                                                                                                                                                                                                                      |                                                                                                                                                                                                                     | <ul style="list-style-type: none"> <li>Zero-shot</li> <li>Few-shot</li> </ul>            |          |
| <ul style="list-style-type: none"> <li>Demographic &amp; Protected Attributes (e.g., Gender and User Country Detection)</li> </ul>                          |                                                                                                                                                                                                                      |                                                                                                                                                                                                                     |                                                                                          |          |
| <ul style="list-style-type: none"> <li>Sentiment, Stylistic &amp; Emotion Analysis (e.g., Stance Detection, Sarcasm Detection)</li> </ul>                   |                                                                                                                                                                                                                      |                                                                                                                                                                                                                     |                                                                                          |          |
| <ul style="list-style-type: none"> <li>Machine Translation (e.g., English-Arabic and Arabic dialects)</li> </ul>                                            |                                                                                                                                                                                                                      |                                                                                                                                                                                                                     |                                                                                          |          |
| <ul style="list-style-type: none"> <li>News Categorization</li> </ul>                                                                                       |                                                                                                                                                                                                                      |                                                                                                                                                                                                                     |                                                                                          |          |
| <ul style="list-style-type: none"> <li>Question Answering</li> </ul>                                                                                        |                                                                                                                                                                                                                      |                                                                                                                                                                                                                     |                                                                                          |          |



# LLMeBench

A Complete Video Tutorial



<https://rb.gy/6m6h2b>



# Language Model Evaluation Harness

*A framework to evaluate LLMs on a large number of tasks and datasets*

- Over 60 standard academic benchmarks for LLMs, with hundreds of subtasks and variants implemented.
- Support for models loaded via [transformers](#) (including quantization via [AutoGPTQ](#)), [GPT-NeoX](#), and [Megatron-DeepSpeed](#), with a flexible tokenization-agnostic interface.
- Support for fast and memory-efficient inference with [vLLM](#).
- Support for commercial APIs including [OpenAI](#), and [TextSynth](#).
- Support for evaluation on adapters (e.g. LoRA) supported in [HuggingFace's PEFT library](#).
- Support for local models and benchmarks.
- Evaluation with publicly available prompts ensures reproducibility and comparability between papers.
- Easy support for custom prompts and evaluation metrics.

<https://github.com/EleutherAI/lm-evaluation-harness>



# Language Model Evaluation Harness

## *Pros*

- Does not require explicit prompting
- Evaluation is based on log-likelihood
- Good for fast evaluation of LLMs

## *Cons*

- Evaluation is not based on token(s) to represent candidate answer
- Lack of chat-templates

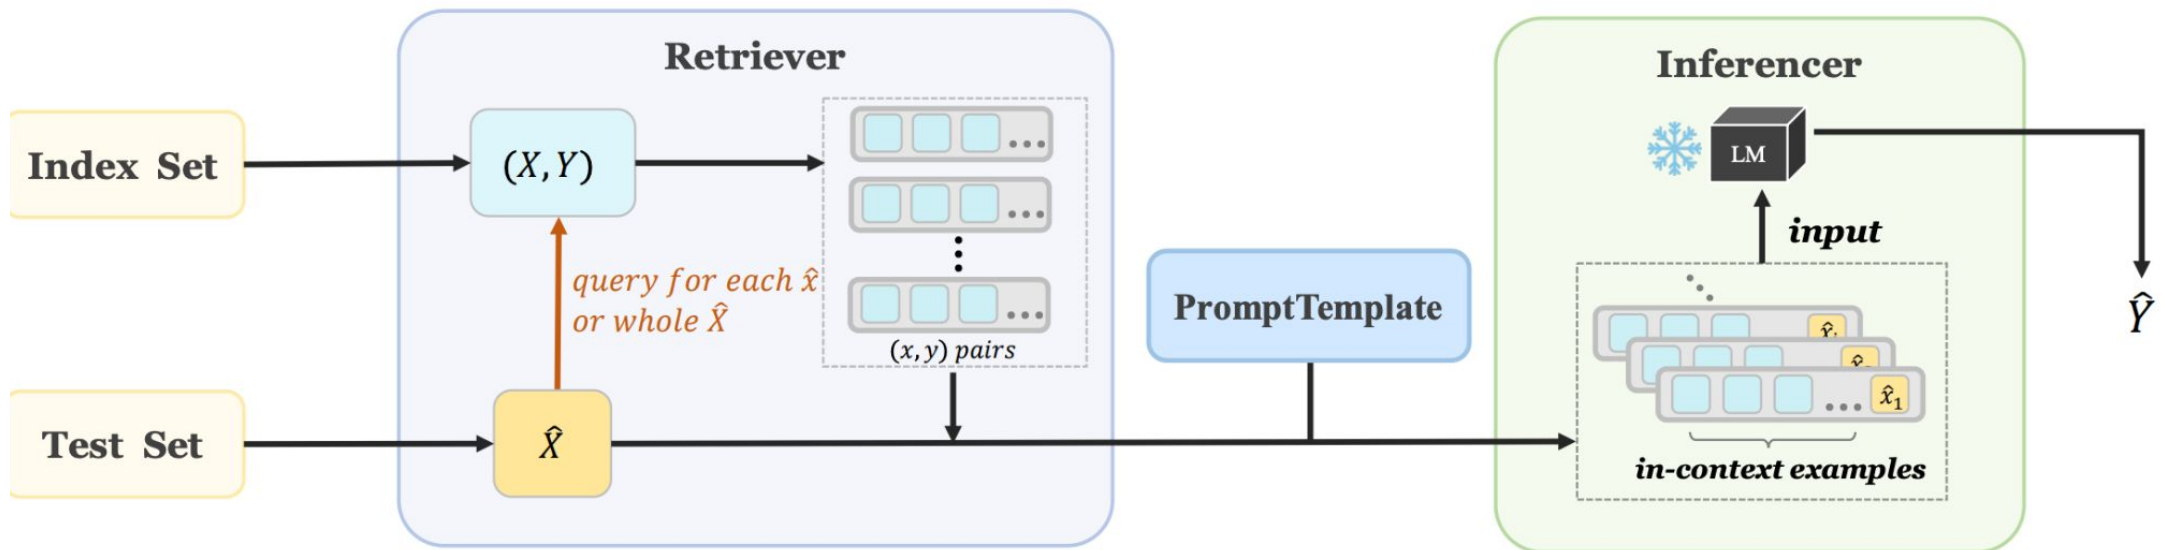
<https://github.com/EleutherAI/lm-evaluation-harness>





# Open ICL

*An easy-to-use and extensible in-context-learning (ICL) framework for zero-/few-shot evaluation of LLMs*



- Random
- Heuristic method (BM25, TopK, VoteK)
- Model based approach
- Tokens in candidate answer
- Perplexity

<https://github.com/Shark-NLP/OpenICL>



# Open ICL

## Features

- Supports many state-of-the-art retrieval methods
- A unified and flexible interface for the development and evaluation of new ICL methods
- Implements data parallelism to improve the performance of both the retrieval and inference steps
- Model parallelism that users can easily parallelize their models with minimal modification to the code.

<https://github.com/Shark-NLP/OpenICL>



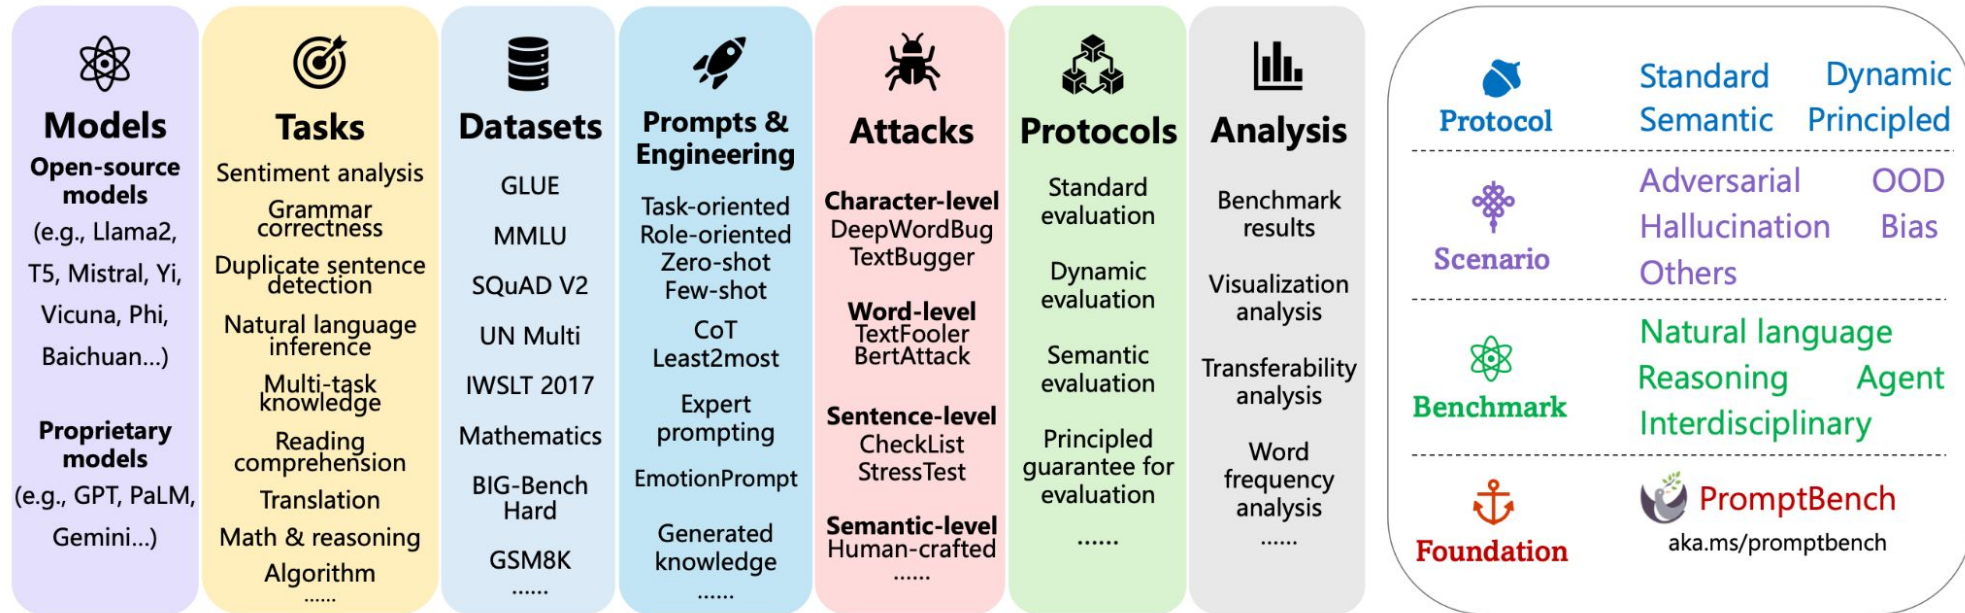
# Prompt Bench

*A Unified Library for Evaluating and Understanding LLMs.  
A comprehensive benchmark designed for assessing the robustness of LLMs to adversarial prompts*



(a) Components

(b) Supported evaluation research topics



<https://github.com/microsoft/promptbench>



# Prompt Bench

## Features

- Quick model performance assessment
- Prompt Engineering
- Evaluating adversarial prompts
- Dynamic evaluation to mitigate potential test data contamination

<https://github.com/microsoft/promptbench>



# LLM-as-a-Judge

*MT-bench is a challenging multi-turn question set designed to evaluate the conversational and instruction-following ability of models*

- 80 high-quality, multi-turn questions
- automated evaluation pipeline based on GPT-4

```
[System]
Please act as an impartial judge and evaluate the quality of the response provided by an
AI assistant to the user question displayed below. Your evaluation should consider factors
such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of
the response. Begin your evaluation by providing a short explanation. Be as objective as
possible. After providing your explanation, please rate the response on a scale of 1 to 10
by strictly following this format: "[[rating]]", for example: "Rating: [[5]]".
```

```
[Question]
{question}
```

```
[The Start of Assistant's Answer]
{answer}
[The End of Assistant's Answer]
```


prompt for single answer grading





**QA**

# Thank You



<https://llm-low-resource-lang.github.io/>