

Let LLMs Take on the Latest Challenges !

A Chinese Dynamic Question Answering Benchmark

Anonymous ACL submission

Abstract

How to better evaluate the capabilities of Large Language Models (LLMs) is the focal point and hot topic in current LLMs research. Previous work has noted that due to the extremely high cost of iterative updates of LLMs, they are often unable to answer the latest dynamic questions well. To promote the improvement of Chinese LLMs’ ability to answer dynamic questions, in this paper, we introduce **CDQA**, a Chinese **D**ynamic **Q**A benchmark containing question-answer pairs related to the latest news on the Chinese Internet. We obtain high-quality data through a pipeline that combines humans and models, and carefully classify the samples according to the frequency of answer changes to facilitate a more fine-grained observation of LLMs’ capabilities. We have also evaluated and analyzed mainstream and advanced Chinese LLMs on *CDQA*. Extensive experiments¹ and valuable insights suggest that our proposed *CDQA* is challenging and worthy of more further study. We believe that the benchmark we provide will become the key data resource for improving LLMs’ Chinese question-answering ability in the future.

1 Introduction

Due to the excellent emergence capabilities and unified task paradigm, Large Language Models (LLMs) are undoubtedly the more popular stars in the field of Natural Language Processing (NLP) or Artificial Intelligence (Wei et al., 2022; Li et al., 2023; Shanahan, 2024). To promote the improvement of LLMs capabilities, more and more researchers have invested in building various LLMs evaluation benchmarks (Chang et al., 2023). In the era of LLMs, high-quality evaluation benchmarks allow researchers to better understand the true capabilities of LLMs, thereby stimulating further research on how to enhance LLMs.

¹Code and dataset will be available for reproducibility.

Static Question	ACL 主会每年举办几次? How many times does the ACL annual meeting take place each year?
GPT-4’s Answer	一年一次。 Once a year. ✓
Dynamic Question	下一次ACL 将在哪里举办? Where will the next ACL be held?
GPT-4’s Answer	我无法提供相关信息。 I can’t provide the information. ✗

Table 1: Examples of static and dynamic questions. The GPT-4 is on Feb 11, 2024.

Question answering is an important and long-standing topic in NLP (Rajpurkar et al., 2016; Joshi et al., 2017; He et al., 2018). Especially for LLMs, QA tasks have almost become the indispensable basic task in LLMs research (Pan et al., 2024). Various forms of QA benchmarks can be used to measure the capabilities of LLMs in different dimensions (Adlakha et al., 2022; Bosselut et al., 2022; Rein et al., 2023; Huang et al., 2023). Recently, the introduction of English FreshQA (Vu et al., 2023) has attracted widespread attention. It challenges LLMs through questions with dynamically changing answers, aiming to test LLMs’ mastery of the latest factual knowledge. Obviously, being able to answer the latest questions determines to some extent whether LLMs can truly move towards large-scale daily applications. **Urgently, we note that there is still no such benchmark in the Chinese community, although LLMs in the Chinese scenario still face the same challenges and dilemmas**, as shown in Table 1.

To let LLMs in Chinese scenarios take on the latest challenges and empower them to answer dynamic questions, in this work, we present **CDQA**, a Chinese **D**ynamic **Q**A benchmark. Specifically, we design a semi-automatic data production pipeline to construct our benchmark. In this pipeline, we first automatically generate a large number of raw

queries with the help of two LLMs with different roles, one is to extract key entities from the latest Chinese news, and the other is to automatically generate question queries based on the extracted entities that will be as the corresponding answers. Then we ask the well-trained annotators to filter, rewrite, and classify the automatically generated question samples to ensure the quality of *CDQA*. Through such a semi-automatic data construction method with human participation, we obtain 1,339 question-answer pairs for *CDQA*, classified by how frequently their answers change (i.e., fast-changing, slow-changing, and never-changing). The purpose of classifying *CDQA* samples by the frequency of answer changes is to provide finer-grained evaluation for LLMs, facilitating researchers to better perceive the true performance of LLMs.

Based on our constructed *CDQA*, we select a series of widely used and advanced LLMs in the Chinese community for evaluation. Results show that **GPT-4** still ranks at the top with searched results from search engines, surpassing at most **nearly 10 F1-recall scores** with less hallucination than the second-best model, i.e., **Deepseek-67B-Chat**, although Chinese-oriented LLMs tend to have more internal knowledge than OpenAI models. Besides, **in-context learning** and **prompting methods** like Chain-of-Thought generally increase performances with searched evidence but also elicit more hallucinations in LLMs. For **search engines** in the open-book scenario, Google consistently takes advantage over Bing for all baseline models, showing great strength as a retriever for LLMs.

In summary, our contributions could be summarized as follows:

1. We first introduce the idea of using dynamic questions to challenge Chinese LLMs, which provides a new direction for the development of LLMs in Chinese community.
2. We construct the high-quality *CDQA* benchmark composed of dynamic questions, which will become an important data resource for promoting the progress of Chinese LLMs.
3. Extensive experiments and detailed analyses based on *CDQA* provide valuable insights and discoveries, which are instructive for subsequent research about how to enhance LLMs to handle dynamic questions.

2 Chinese Dynamic Question Answering (CDQA)

2.1 Overview

Our dataset, **CDQA**, mainly originates from latest news in Chinese Internet from different areas such as finance, daily life, politics, technology and so on. Besides, there are also queries collected from Chinese labelers. They represent the information-seeking cases of Chinese people. The generation pipeline could be illustrated in Figure 1. The dataset currently consists of 1,339 questions covering a range of topics with evolving answers which are mostly extracted entities from the raw corpus scraped from Chinese Internet and it is being regularly updated.

2.2 Data Collection

We collect *CDQA* dataset in two stage. **The first stage is automatic generations with Entity Extraction and Doc2Query**, for which we use SeqGPT (Yu et al., 2023), and GPT-4 (OpenAI, 2023), which could give great amount of raw question-answering pairs as SeqGPT extracts entities from latest Chinese news and GPT-4 is prompted into generating corresponding questions. For GPT-4 prompts, we use few-shot prompting in generating diverse questions from entities. **The second stage is manual labeling from crowd-sourced workers**. The Chinese labelers not only filter questions which are answered with biases, ambiguities and obsolete² knowledge but also annotate with **tags**, check the correctness and **rewrite** the question answer pairs to be more time-related and dynamic. At the very beginning, the labelers are shown with pre-annotation examples and annotation guides.

Tags The tags are annotated for questions and answers. For questions, we have the same taxonomy as *FreshQA* (Vu et al., 2023). The questions are categorized as **fast-changing**, **slow-changing**, and **never-changing**. For answers, we categorize these entities or short texts as **person**, **location**, **time**, **event**, **artificial work**, **group**, **nature**, **quantity** and **other**. Therefore, we could evaluate the models' latest world knowledge from various perspectives. The taxonomy is illustrated in Appendix A.

Quality Control After getting the synthetic queries, the human annotators could rewrite and

²The answer should be only supported with the knowledge after January 1st 2019 except for static knowledge, i.e., never-changing.

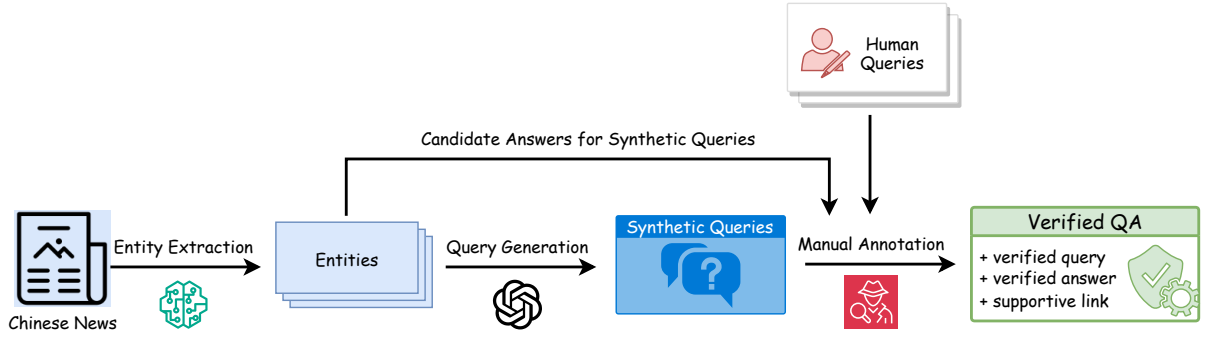


Figure 1: Data Generation Pipeline for CDQA dataset. We first collect Chinese News from Internet and then extract entities from these news passages. Based on GPT-4, we generate synthetic queries from passages and corresponding entities. Manual annotation is conducted to verify the synthetic data and extra human-crafted queries, providing the verified queries, answers and supportive evidence links.

	Acceptance	Question Tags	Answer Types
Ann1 v.s. Ann2	62.3	87.2	96.6
GT v.s. Ann1	79.6	59.1	100
GT v.s. Ann2	47.3	68.3	100
Avg	63.1	71.5	98.9

Table 2: Inter-annotator agreement for different annotation sections are calculated by **Cohen Kappa scores**. Ann1/2 represents Annotator1/2 respectively and GT represents Ground Truth. Our annotations could be considered as “substantial agreement” as the average scores are above 60.

calibrate the questions and answers to make the QA pairs correct, consistent and dynamic. For example, annotators are required to provide the supporting evidence URLs along with correct answers using search engines. This calibration process could solidify our answers with supplementary valid information and help us better iterate the dataset as the generation process in the previous stage is not well-evaluated with supportive documents, let alone the correctness. Moreover, in order to facilitate the periodic updates, we filter out the questions with more than one valid answer.

For inter-annotator agreement, we randomly sample 100 examples from synthetic question-answer pairs and annotations from two annotators in the same annotation vendor are measured by **acceptance** (*whether the pair is accepted or discarded*), **question tags** and **answer types**. The ground-truth labels are provided by authors. For each category, we calculate their Cohen Kappa scores (McHugh, 2012). From Table 2, the averaged score across all types of annotations are above 63.1, representing “substantial agreement” for our dataset annotations.

2.3 Regular Updates

Our dataset is highly sensitive to time since the ground truth is evolving along the world development. Therefore, we commit to updating the dataset regularly and researchers are strongly encouraged to stay tuned with our latest version for evaluation. And the datasets are mainly calibrated with information from Chinese Internet.

2.4 Data Statistics

Due to limitations in automatic query generation by GPT-4 and SeqGPT from the first stage, our dataset has low **retention rate** in which only 44.6% synthetic data are accepted by human annotators. Among the accepted data, 53.1% of them still need further modifications because of improper questions or wrong answers. For **question tags**, we have relatively balanced distributions between *fast-changing* and *slow-changing* questions with fewer *never-changing* questions. For **answer types**, we have biased distributions as nearly 70% of entities extracted from passages lie in “person” and “group” categories. This is because most of entities in first stage by automatic generation are “person” and “group”. However, question tags and answer types could be changed or calibrated over time by re-annotation of the dataset. These distribution graphs and more details about our dataset are in Appendix B.

2.5 Evaluation

As our dataset is constructed from Internet, our evaluation setup is based on **retrieval-augmented question answering** with search engines while the evaluation methods are *answer rate* and *F1-recall*. Our experiments consist of two settings: **close-book** and **open-book**. Overall, our evaluation pro-

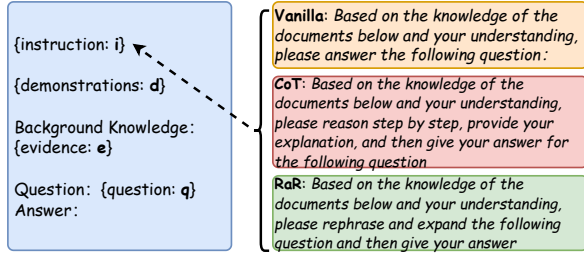


Figure 2: Our prompts are formulated under this framework. Different prompting methods are used with different instructions *i*. The Chinese version is in Appendix C.

vides a comprehensive understanding of current LLMs in factuality, especially for evolving knowledge. Besides, due to the safety implementation for different LLMs from helpful and harmless responses in training data (Bai et al., 2022), **F1-recall only counts on questions with effective responses** while **answer rate is used in representing the ratio of answered questions to the total questions**, which is a practical metric for the real world application of LLMs and could directly indicate the degree of hallucination in generated responses.

Evaluation Metrics For **F1-recall**, we calculate *the ratio of common tokens between model-generated responses and ground truth to the ground truth*. Specifically, we first segment the generated text and golden text into token lists using word segmentation tools³, then calculate the ratio of tokens generated by models belonging to the golden token list to golden tokens. For **answer rate**, we directly calculate *the ratio of effectively answered questions to total questions*, i.e., responses of refusal are filtered out in our evaluation.

3 Experiments

3.1 Experiment Setup

Baselines We experiment with a series of baseline models pretrained with Chinese data, including representative OpenAI’s **ChatGPT** (*gpt-3.5-turbo-1106*) (OpenAI, 2022) and **GPT-4** (*gpt-4-1106-preview*) (OpenAI, 2023), open-sourced Chinese-oriented models such as **Internlm-20B-Chat** (InternLMTeam, 2023), **Aquila2-34B-Chat** (BAAI, 2023), **Yi-34B-Chat** (01-ai, 2023), **Deepseek-67B-Chat** (DeepSeek-AI et al., 2024), **Qwen-72B-Chat** (Bai et al., 2023). In the close-book scenario, we only use the standalone LLM in answering questions with zero-shot and few-shot settings. For the

³<https://github.com/fxsjy/jieba>

open-book scenario, search engines are used for retrieving question-related results on the Internet and then fed into LLMs for reading.

Search Engines Except for language models for information synthesis, we select three representative search engines to recall relevant passages from the Chinese Internet namely **Google** and **Bing**. These search engines are mainly used by Chinese people for information seeking. **Baidu** is omitted due to the difficulty in scraping its contents. The **Top-5** and **Top-10** searched results are provided to models in the **open-book setting**.

Prompt Design Our prompt framework, which is in Chinese, could be framed as concatenation of (*i*, *d*, *e*, *q*), in Figure 2 where *i* represents the instruction, *d* for question-answer pairs from crowdsourced labelers, *e* for search results and *q* for current question. Different instructions *i* are used with three widely adopted prompting styles, **Vanilla**, **Chain-of-Thought (CoT)** (Wei et al., 2023) and **Rephrase-and-Respond (RaR)** (Deng et al., 2023). **Vanilla** instruction is directly asking models to answer questions with the context. **CoT** instruction is asking models to first explain and analyze the question *q* step by step and then give their answers. **RaR** instruction, however, is asking models to first rephrase and expand the question *q* and then give their answers, which could be viewed as a complement of CoT (Deng et al., 2023) as CoT is for diving deeper while RaR is for exploring broader. Besides, for demonstrations *d*, we have used zero-shot and different few-shot settings, i.e., 5-shot and 16-shot. More specifically, our few-shot demonstrations are made up of human written questions and answers similar to *CDQA* dataset without contexts or other explanations as it costs longer time without any improvement.

3.2 Results and Analyses

Table 3, 4, 5 summarize best performances over few-shot prompting across different baselines for Vanilla, CoT and RaR prompts respectively. Our default search engine for analysis is **Google** as it is most widely used around the world.

Baseline Comparison In Table 4, we see that **GPT-4** achieves the best across all questions. Only GPT-4 reaches nearly or over 65, 70 and 80 in F1-recall for *fast-changing*, *slow-changing* and *never-changing* questions respectively with searched results which indicates that there is still room of

Model	fast-changing			slow-changing			never-changing		
	+0	+5	+10	+0	+5	+10	+0	+5	+10
Interlm-20B-Chat	16.1 [16]	49.7 [0]	50.5 [0]	19.8 [16]	61.2 [16]	61.8 [0]	34.0 [16]	68.1 [16]	71.7 [16]
Aquila2-34B-Chat	14.9 [16]	49.4 (99.6%) [0]	51.5 [0]	17.7 [16]	60.4 [0]	62.5 [0]	35.6 [5]	69.2 [0]	69.4 [0]
Yi-34B-Chat	22.9 [16]	54.9 [16]	56.5 [16]	30.8 [16]	67.0 [16]	68.8 [16]	46.9 [16]	76.0 [16]	76.9 [16]
Deepseek-67B-Chat	24.3 [0]	57.5 [16]	58.4 [16]	37.2 [16]	68.3 [16]	70.0 [16]	53.1 [0]	78.8 [16]	79.2 [5]
Qwen-72B-Chat	18.6 [16]	51.7 [16]	54.7 [16]	16.9 [16]	63.1 [0]	64.6 [16]	41.5 [16]	74.6 [16]	75.4 [16]
ChatGPT	18.1 (96.6%) [0]	56.2 (98.3%) [0]	59.2 (98.3%) [0]	14.1 (93.3%) [0]	64.7 (97.9%) [0]	66.3 (98.3%) [0]	34.7 (99.0%) [0]	73.2 [0]	73.7 (99.7%) [0]
GPT-4	35.1 (13.5%) [0]	59.6 (96.0%) [0]	61.2 (96.4%) [0]	33.8 (25.4%) [0]	65.8 (97.5%) [0]	68.4 (96.5%) [0]	54.4 (56.1%) [0]	76.4 (99.3%) [5]	78.8 (98.6%) [5]

Table 3: Best performance over few-shot settings of baselines for **Vanilla** prompt with searched results from Google. +5 and +10 represent different numbers of searched results appended in the inputs. We report in the form of *F1-recall (answer rate) [best number of few shot examples]* and omit the answer rate if it is 100%. Data with the highest F1-recall scores are marked in bold.

Model	fast-changing			slow-changing			never-changing		
	+0	+5	+10	+0	+5	+10	+0	+5	+10
Interlm-20B-Chat	15.0 [16]	46.8 [5]	49.5 [5]	18.1 [16]	59.3 [16]	61.7 [0]	34.7 [16]	67.2 [16]	68.9 [16]
Aquila2-34B-Chat	14.5 [16]	50.0 (99.6%) [0]	51.9 [0]	17.1 [16]	60.0 [0]	61.4 [0]	35.6 [5]	69.5 [0]	69.8 [0]
Yi-34B-Chat	23.2 [16]	54.7 [5]	57.4 [16]	30.4 [16]	67.8 [16]	68.5 [16]	47.0 [16]	75.3 [16]	77.3 [16]
Deepseek-67B-Chat	22.9 [16]	58.0 [16]	59.2 [16]	37.0 [5]	67.8 [16]	70.6 [16]	53.0 [16]	79.8 [16]	80.2 [16]
Qwen-72B-Chat	17.8 [16]	63.6 [5]	64.2 [16]	17.8 [16]	63.1 [0]	63.6 [5]	40.8 [16]	74.2 [16]	76.1 [16]
ChatGPT	17.9 (97.3%) [0]	57.6 (97.0%) [0]	61.4 (96.6%) [0]	13.9 (98.3%) [0]	65.5 (97.5%) [0]	65.7 (98.7%) [0]	36.0 (99.7%) [0]	75.2 (98.6%) [0]	74.9 (98.6%) [0]
GPT-4	22.1 (82.9%) [16]	67.3 (89.0%) [0]	68.0 (89.0%) [5]	19.8 (86.7%) [0]	73.9 (93.3%) [0]	74.7 (90.4%) [5]	48.2 (56.1%) [0]	81.6 (98.3%) [5]	83.5 (98.3%) [0]

Table 4: Best performance over few-shot settings of baselines for **CoT** prompt with searched results from Google. +5 and +10 represent different numbers of searched results appended in the inputs. We report in the form of *F1-recall (answer rate) [best number of few shot examples]* and omit the answer rate if it is 100%. Data with the highest F1-recall scores are marked in bold.

improvement for other open-sourced models with Chinese instructions under the retrieval-augmented-generation scenario. Besides, the second-best model, **Deepseek-67B-Chat**, has shown great performance as it even surpasses GPT-4 on *slow-changing* and *never-changing* questions by Vanilla prompt with searched results in Table 3 and performs better (both in answer rate and F1-recall) than GPT-4 without searched results in *slow-changing* and *never-changing* questions especially in CoT prompt. This indicates that Deepseek-67B-Chat has stored more Chinese Internet knowledge as its alignment corpora mainly focus on Chinese and English while GPT-4 aligns with multilingual data. Moreover, it is worth noting that answer rates of ChatGPT and GPT-4 are often lower than 100% across all different prompts and question types especially for GPT-4 in close-book question answering which indicates that **there are more hallucination reduction measures such as refusal of questions in GPT-4 than other models.**

How do different styles of prompts work in LLMs? As our evaluation scenarios are made up of close-book and open-book, different prompts just elicit different behaviors in LLMs. To rule out the influence of few-shot demonstrations, we use zero-shot setting in this analysis. We use GPT-4 and Deepseek models as closed and open-sourced representative models in the following analysis. In


Figure 3, under **close-book** scenario, we see that Deepseek-67B-Chat has same answer rates over different prompts on different questions while there are different answering behaviors in GPT-4. Specifically, GPT-4 answers with great care in vanilla prompts with lowest answer rates but high F1-recall scores while GPT-4 suffers from hallucination in CoT and RaR prompts with at most +522% and +176% in answer rates but -43% and -17% in F1-recall scores compared to Vanilla prompt. For both models, Vanilla prompts outperform the other two kinds of prompts. **This indicates that verbose explanation or expansion could increase hallucination especially when without evidence.**

In Figure 4, under **open-book** scenario⁴, we see that Deepseek-67B-Chat still shares same answer rates across different prompts and question types and vanilla prompt is the best for it, which indicates that CoT and RaR take risks in inducing more hallucinated answers. For GPT-4, answer rates have less gaps between different prompts (all over 90%) and F1-recall scores are all increasing dramatically compared to close-book counterparts, representing adding contextual information elicits LLMs in answering questions more efficiently. Besides, with search results, CoT and RaR both outperform Vanilla prompt and CoT performs the best in GPT-


⁴We use Top-10 searched results as more searched results generally improve the results of LLMs from Table 3, 4, 5.

Model	fast-changing			slow-changing			never-changing		
	+0	+5	+10	+0	+5	+10	+0	+5	+10
Internlm-20B-Chat	16.9 [16]	49.5 [0]	52.4 [16]	20.0 [16]	61.2 [5]	63.3 [0]	34.6 [5]	68.2 [16]	71.3 [16]
Aquila2-34B-Chat	15.5 [16]	48.4 [16]	51.4 [0]	17.5 [16]	59.6 [16]	61.9 [0]	36.1 [16]	68.0 [16]	69.5 [0]
Yi-34B-Chat	22.8 [16]	54.6 [16]	57.0 [16]	30.6 [16]	68.2 [16]	68.5 [0]	47.7 [16]	77.1 [5]	76.8 [16]
Deepseek-67B-Chat	23.3 [5]	57.6 [16]	58.9 [5]	37.7 [5]	68.5 [16]	70.7 [16]	54.2 [16]	79.6 [5]	79.8 [5]
Qwen-72B-Chat	18.3 [16]	52.5 [16]	55.4 [16]	17.1 [0]	62.9 [16]	65.4 [16]	39.5 [16]	74.1 [16]	75.7 [16]
ChatGPT	19.2 [0]	58.9 [0]	61.7 [0]	15.9 [0]	65.8 (99.4%) [0]	67.6 (99.6%) [0]	35.6 [0]	75.2 (99.7%) [0]	76.5 (99.7%) [0]
GPT-4	29.1 (37.3%) [0]	61.8 (94.5%) [0]	64.9 (95.2%) [0]	28.3 (54.62%) [0]	71.6 (94.2%) [0]	71.6 (96.2%) [0]	54.9 (72.5%) [0]	80.7 (99.0%) [0]	82.6 (99.3%) [0]


Table 5: Best performance over few-shot settings of baselines for **RaR** prompt with searched results from Google. +5 and +10 represent different numbers of searched results appended in the inputs. We report in the form of *F1-recall(answer rate)[best number of few shot examples]* and omit the answer rate if it is 100%. Data with the highest F1-recall scores are marked in bold.



(a) Deepseek-67B-Chat




(a) Deepseek-67B-Chat



(b) GPT-4

Figure 3: F1-recall scores and Answer Rates of **different prompts** for LLMs in **close-book** scenario under zero-shot setting. We represent F1-recall scores with bar plots and answer rates with dotted lines.



(b) GPT-4

Figure 4: F1-recall scores and Answer Rates of **different prompts** for LLMs in **open-book** scenario under zero-shot setting. We represent F1-recall scores with bar plots and answer rates with dotted lines.

4 with less hallucination, i.e., lower answer rate, and higher F1-recall score. **This indicates that CoT and RaR could improve LLMs on complex tasks but CoT elicits more reasoning abilities to directly improve the answering.**

Nevertheless, model sizes and training data are both fundamental for these prompts to work. In Figure 5, **not every model improves with CoT or RaR compared to Vanilla prompt.** For example, Deepseek-34B-Chat’s and Qwen-72B-Chat’s performance decrease in CoT and RaR; ChatGPT and Internlm-20B-Chat prefer RaR while GPT-4 and

Yi-34B-Chat prefer CoT for larger gains; Aquila2-34B-Chat is robust to all prompt types.

Does few-shot prompting always work for all LLMs? For better analyzing the influence of few-shot prompting, we collect close-book and open-book results of all LLMs in **zero-shot** setting and **vanilla** prompt. In Figure 6, based on nearly 100% answer rate, five in close-book scenario and four in open-book scenario out of all five open-sourced models have better performance with more few-shot demonstrations, which are sampled in the same data distribution during generating *CDQA*.





Figure 5: F1-recall scores averaged over all three different questions for all models with **different prompts** in open-book scenario under zero-shot setting. We present F1-recall score only since all answer rates $\geq 90\%$.

Models	close-book			open-book		
	0-shot	5-shot	16-shot	0-shot	5-shot	16-shot
ChatGPT	96.3	95.0	96.7	98.8	99.7	99.9
GPT-4	31.7	52.2	64.7	97.4	97.7	98.0


Table 6: Answer rates for ChatGPT and GPT-4 averaged on all types of questions in close-book and open-book scenarios with **different few-shot settings**.

However, we also notice that ChatGPT and GPT-4 have shown different trends compared to open-sourced models, i.e., more few-shot examples lead to decreases in F1-recall scores. Therefore, we check the answer rates of ChatGPT and GPT-4 in Table 6 where ChatGPT stays in fairly high answer rates ($\geq 95\%$) and GPT-4 increases its answer rates with more few-shot examples. Combined with their monotonic decrease in F1-recall scores, we show that ChatGPT and GPT-4 hallucinate more with more few-shot examples in *CDQA*. **This indicates the challenge of *CDQA* and different LLMs with different capabilities struggle in different ways.**

How do different search engines help? For a more detailed comparison between search engines across all baselines, we use vanilla prompt since CoT performs better in LLMs and zero-shot setting with Top-10 searched results from different search engines. In Figure 7, searched results from Google consistently outperforms Bing among all baseline models, which indicates that the **Google currently provides the most useful retrieved evidence for question answering about Chinese news.**



(a) Close-book



(b) Open-book (Top-10 searched results)

Figure 6: F1-recall scores averaged over all types of questions for different models with **different few-shot settings** in close-book and open-book scenarios.

How do LLMs perform across different answer types? As answers in *CDQA* are mainly entities from news, we conduct analysis across different answer types for two representative open-sourced and closed LLMs, i.e., Deepseek-67B-Chat and GPT-4. In Figure 8, we observe that **GPT-4’s internal knowledge is poorer than Chinese-oriented models such as Deepseek-67B-Chat for Chinese users.** However, with enough retrieved evidence, **GPT-4 has stronger abilities in learning from contexts than Deepseek-67B-Chat** where this “learning efficiency”, i.e., the ratio of gaps between open-book scores and close-book scores to the close-book could reach at most 1370% compared to 219% in Deepseek-67B-Chat. Moreover, from Figure 8, we also could notice that the “quantity” group is hardest for models to figure out the correct answers, which is due to the granularity of answers and the need of reasoning abilities.

4 Related Work

Temporal and Dynamic QA Benchmark *StreamingQA* (Liska et al., 2022) is a QA dataset where questions are human written or LM-generated on given dates, showing how




Figure 7: F1-recall scores averaged over all questions for different models with **different search engines**.




Figure 8: F1-recall scores on **different answer types** for Deepseek-67B-Chat and GPT-4 in close-book and open-book scenarios with Vanilla prompt. We use Top 10 searched results from Google. Deepseek-67B-Chat holds 100% answer rate in all time while GPT-4 drastically increase its answer rate to 100% with searched results from Google.

open-book and close-book QA models adapt to new knowledge over time and importance of retrieval augmentation in up-to-date search space. *TimeQA* (Chen et al., 2021) are formed from extracted evolving facts in *WikiData* by manual extraction and verification while we extract entities to directly formulate them as answer candidates based on the documents. *RealTimeQA* (Kasai et al., 2022), a dynamic QA benchmark with automatic weekly updates from the weekly News Quiz section in social media such as CNN, is most related to our semi-automatic question generation with latest Chinese news corpus.

English QA benchmark Question answering is a long-standing task in NLP area, ranging from classic single-turn benchmarks such as *SQuAD* (Rajpurkar et al., 2016, 2018), *TriviaQA* (Joshi et al., 2017) and *Natural Questions* (Kwiatkowski et al., 2019) to conversational QA like *TopiOCQA* (Ad-

lakha et al., 2022). Also, it could also be categorized by normal human intelligence and human experts intelligence. For example, *CommonsenseQA* (Talmor et al., 2019) comes up with questions from *ConceptNet* in discriminating the target concepts with commonsense while *GPQA* (Rein et al., 2023) consists of questions generated from graduate-level students. Although these benchmarks have provided efficient and effective evaluation metrics and covered a variety of topics, they are mostly static with a focus on less-evolving knowledge, which might already be or going to be included in the pre-training data of large language models.

Chinese QA benchmark In contrast to prosperous English QA benchmarks, Chinese counterparts are still under-explored. *DuReader* (He et al., 2018) is a classic free-form QA benchmark collected by Baidu from its own products and *CLUE* (Xu et al., 2020) is the first large scale NLU benchmark in Chinese. After the recent debut of powerful large language models, a series of Chinese QA benchmarks are proposed for better evaluating them. *C-Eval* (Huang et al., 2023) is a multiple-choice questions answering dataset from Chinese Standard Exams. *WebCPM* (Qin et al., 2023) collects questions from web forums through web searching and browsing and *SuperCLUE* (Xu et al., 2023) is a comprehensive Chinese benchmark for question answering in aligning users needs. But they all suffer from either data leakage or the risk of saturated performance which hinders the accurate evaluation on questions requiring fresh knowledge to answer as static questions are readily overfitted.

5 Conclusion

The creation of *CDQA* addresses the urgent need for the evaluation of Chinese LLMs, thereby improving LLM-driven applications for Chinese users. Given the cultural influences in LLMs’ training data, it is our aspiration that *CDQA* will foster development in various capabilities of LLMs, particularly within Chinese contexts. While *CDQA* progresses further with a semi-automatic generation pipeline with more data than *FreshQA*, we acknowledge that it is far from a perfect LLM evaluation. Other critical dimensions, including tool learning, LLMs safety, and robustness, remain to be explored. However, we believe that our constructed *CDQA* and the series of insights obtained based on it will provide valuable resources and guidance for subsequent research on Chinese LLMs.

Limitations

One of the limitations of our work is that the language we study is Chinese only. As the two most widely used languages in the world, English and Chinese have always been equally valued and widely concerned in the NLP community. In fact, our work is inspired by previous *FreshQA* in the English scenario and aims to provide similar data resources to Chinese LLMs researchers. We also encourage and welcome more researchers from other languages to engage in similar research.

In addition, another limitation that cannot be ignored is how to keep our *CDQA* updated. Because *CDQA* focuses on questions whose answers change dynamically, it is critical to ensure that the answers to questions in *CDQA* are always correct and up-to-date. Therefore, we also commit to updating our *CDQA* regularly and providing researchers with the latest version of *CDQA* for LLMs evaluation.

Ethics Statement

The task we focus on is the evaluation of LLMs, and the LLMs we evaluate are all public and widely used LLMs, so they do not bring potential ethical risks. The data samples of *CDQA* that we collect have been manually cleaned and pre-processed to ensure that they do not contain any data that will cause moral risks, such as politically sensitive, violent, and private data. In addition, we also have signed legal labor contracts with the human annotators we employ, and pay them higher than market prices based on their workload.

References

- 01-ai. 2023. Yi series large language models. <https://github.com/01-ai/Yi>.
- Vaibhav Adlakha, Shehzaad Dhuliawala, Kaheer Suleman, Harm de Vries, and Siva Reddy. 2022. *Topicqa: Open-domain conversational question answering with topic switching*.
- BAAI. 2023. Aquila2. <https://github.com/FlagAI-Open/Aquila2>.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu,

Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. *Qwen technical report*.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. *Training a helpful and harmless assistant with reinforcement learning from human feedback*.

Antoine Bosselut, Xiang Li, Bill Yuchen Lin, Vered Shwartz, Bodhisattwa Prasad Majumder, Yash Kumar Lal, Rachel Rudinger, Xiang Ren, Niket Tandon, and Vilém Zouhar, editors. 2022. *Proceedings of the First Workshop on Commonsense Representation and Reasoning (CSRR 2022)*. Association for Computational Linguistics, Dublin, Ireland.

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2023. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*.

Wenhu Chen, Xinyi Wang, and William Yang Wang. 2021. *A dataset for answering time-sensitive questions*.

DeepSeek-AI, :, Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qishui Du, Zhe Fu, Huazuo Gao, Kaige Gao, Wenjun Gao, Ruiqi Ge, Kang Guan, Daya Guo, Jianzhong Guo, Guangbo Hao, Zhewen Hao, Ying He, Wenjie Hu, Panpan Huang, Erhang Li, Guowei Li, Jiashi Li, Yao Li, Y. K. Li, Wenfeng Liang, Fangyun Lin, A. X. Liu, Bo Liu, Wen Liu, Xiaodong Liu, Xin Liu, Yiyuan Liu, Haoyu Lu, Shanghao Lu, Fuli Luo, Shirong Ma, Xiaotao Nie, Tian Pei, Yishi Piao, Junjie Qiu, Hui Qu, Tongzheng Ren, Zehui Ren, Chong Ruan, Zhangli Sha, Zhihong Shao, Junxiao Song, Xuecheng Su, Jingxiang Sun, Yaofeng Sun, Minghui Tang, Bingxuan Wang, Peiyi Wang, Shiyu Wang, Yaohui Wang, Yongji Wang, Tong Wu, Y. Wu, Xin Xie, Zhenda Xie, Ziwei Xie, Yiliang Xiong, Hanwei Xu, R. X. Xu, Yanhong Xu, Dejian Yang, Yuxiang You, Shuiping Yu, Xingkai Yu, B. Zhang, Haowei Zhang, Lecong Zhang, Liyue Zhang, Mingchuan Zhang, Minghua Zhang, Wentao Zhang, Yichao Zhang, Chenggang Zhao, Yao Zhao, Shangyan Zhou, Shunfeng Zhou, Qihao Zhu, and Yuheng Zou. 2024. *Deepseek llm: Scaling open-source language models with longtermism*.

Yihe Deng, Weitong Zhang, Zixiang Chen, and Quanquan Gu. 2023. *Rephrase and respond: Let large language models ask better questions for themselves*.

610	Wei He, Kai Liu, Jing Liu, Yajuan Lyu, Shiqi Zhao,	Yujia Qin, Zihan Cai, Dian Jin, Lan Yan, Shihao	666
611	Xinyan Xiao, Yuan Liu, Yizhong Wang, Hua Wu,	Liang, Kunlun Zhu, Yankai Lin, Xu Han, Ning Ding,	667
612	Qiaoqiao She, Xuan Liu, Tian Wu, and Haifeng	Huadong Wang, Ruobing Xie, Fanchao Qi, Zhiyuan	668
613	Wang. 2018. Dureader: a chinese machine reading	Liu, Maosong Sun, and Jie Zhou. 2023. Webcpm: In-	669
614	comprehension dataset from real-world applications.	teractive web search for chinese long-form question	670
		answering.	671
615	Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei	Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018.	672
616	Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu,	Know what you don't know: Unanswerable questions	673
617	Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu,	for squad.	674
618	Maosong Sun, and Junxian He. 2023. C-eval: A		
619	multi-level multi-discipline chinese evaluation suite	Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and	675
620	for foundation models.	Percy Liang. 2016. Squad: 100,000+ questions for	676
		machine comprehension of text.	677
621	InternLMTeam. 2023. Internlm: A multilingual lan-	David Rein, Betty Li Hou, Asa Cooper Stickland, Jack-	678
622	guage model with progressively enhanced capabili-	son Petty, Richard Yuanzhe Pang, Julien Dirani, Ju-	679
623	ties. https://github.com/InternLM/InternLM .	lian Michael, and Samuel R. Bowman. 2023. Gpqa:	680
624	Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke	A graduate-level google-proof q&a benchmark.	681
625	Zettlemoyer. 2017. Triviaqa: A large scale distantly		
626	supervised challenge dataset for reading comprehen-	Murray Shanahan. 2024. Talking about large language	682
627	sion.	models. <i>Communications of the ACM</i> , 67(2):68–79.	683
628	Jungo Kasai, Keisuke Sakaguchi, Yoichi Takahashi,	Alon Talmor, Jonathan Herzig, Nicholas Lourie, and	684
629	Ronan Le Bras, Akari Asai, Xinyan Yu, Dragomir	Jonathan Berant. 2019. Commonsenseqa: A question	685
630	Radev, Noah A. Smith, Yejin Choi, and Kentaro Inui.	answering challenge targeting commonsense knowl-	686
631	2022. Realtime qa: What's the answer right now?	edge.	687
632	Tom Kwiatkowski, Jennimaria Palomaki, Olivia Red-	Tu Vu, Mohit Iyyer, Xuezhi Wang, Noah Constant, Jerry	688
633	field, Michael Collins, Ankur Parikh, Chris Alberti,	Wei, Jason Wei, Chris Tar, Yun-Hsuan Sung, Denny	689
634	Danielle Epstein, Illia Polosukhin, Matthew Kelcey,	Zhou, Quoc Le, and Thang Luong. 2023. Freshllms:	690
635	Jacob Devlin, Kenton Lee, Kristina N. Toutanova,	Refreshing large language models with search engine	691
636	Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob	augmentation.	692
637	Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natu-	Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel,	693
638	ral questions: a benchmark for question answering	Barret Zoph, Sebastian Borgeaud, Dani Yogatama,	694
639	research. <i>Transactions of the Association of Compu-</i>	Maarten Bosma, Denny Zhou, Donald Metzler, et al.	695
640	<i>tational Linguistics.</i>	2022. Emergent abilities of large language models.	696
641	Yinghui Li, Haojing Huang, Shirong Ma, Yong Jiang,	<i>arXiv preprint arXiv:2206.07682.</i>	697
642	Yangning Li, Feng Zhou, Hai-Tao Zheng, and Qingyu	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	698
643	Zhou. 2023. On the (in)effectiveness of large lan-	Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and	699
644	guage models for chinese text correction. <i>CoRR</i> ,	Denny Zhou. 2023. Chain-of-thought prompting elic-	700
645	abs/2307.09007.	its reasoning in large language models.	701
646	Adam Liska, Tomas Kocisky, Elena Gribovskaya, Tay-	Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao,	702
647	fun Terzi, Eren Sezener, Devang Agrawal, D'Autume	Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong	703
648	Cyprien De Masson, Tim Scholtes, Manzil Zaheer,	Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi,	704
649	Susannah Young, et al. 2022. Streamingqa: A bench-	Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang,	705
650	mark for adaptation to new knowledge over time in	Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian,	706
651	question answering models. In <i>International Confe-</i>	Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao,	707
652	<i>rence on Machine Learning</i> , pages 13604–13622.	Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang	708
653	PMLR.	Yang, Kyle Richardson, and Zhenzhong Lan. 2020.	709
654	Mary McHugh. 2012. Interrater reliability: The kappa	Clue: A chinese language understanding evaluation	710
655	statistic. <i>Biochemia medica : časopis Hrvatskoga</i>	benchmark.	711
656	<i>društva medicinskih biokemičara / HDMB</i> , 22:276–		
657	82.	Liang Xu, Anqi Li, Lei Zhu, Hang Xue, Changtai Zhu,	712
658	OpenAI. 2022. Chatgpt. https://chat.openai.com/	Kangkang Zhao, Haonan He, Xuanwei Zhang, Qiyue	713
659	chat.	Kang, and Zhenzhong Lan. 2023. Superclue: A com-	714
660	OpenAI. 2023. Gpt-4 technical report.	prehensive chinese large language model benchmark.	715
661	Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Ji-	Tianyu Yu, Chengyue Jiang, Chao Lou, Shen Huang,	716
662	apu Wang, and Xindong Wu. 2024. Unifying large	Xiaobin Wang, Wei Liu, Jiong Cai, Yangning Li,	717
663	language models and knowledge graphs: A roadmap.	Yinghui Li, Kewei Tu, Hai-Tao Zheng, Ningyu	718
664	<i>IEEE Transactions on Knowledge and Data Engi-</i>	Zhang, Pengjun Xie, Fei Huang, and Yong Jiang.	719
665	<i>neering.</i>	2023. Seqgpt: An out-of-the-box large language	720
		model for open domain sequence understanding.	721

A Tag Taxonomy of CDQA

The tag taxonomy of *CDQA* is presented in Table 7.

B Dataset Distributions

Knowledge types for queries and answer types are visualized in the following Figure 9, 10. More specifically, we have further visualized the answer type distributions in each question tag. From Figure 11, we see that nearly 80% of slow changing questions are about person and group. Except for person and group, artificial work should be the third largest category for answers, which includes jobs, titles, knowledge and so on. These observations are all consistent with our data sources as information for the protagonists, places and events are compulsory and most frequent in news reports. Besides, percentages of time reach the maximum in never-changing tag as currently most of questions answered with time are about the frequencies.

C Translated Chinese Prompts

The translated prompt framework is illustrated in Figure 12.

Category	Description	Example
fast-changing	The answer to the question is prone to changing within one year	(How many sessions has the Maritime Silk Road Cultural Heritage Forum been held?, Four)
slow-changing	The answer to the question is prone to changing in several years	(Which ancient city site in China has recently been recognized as a UNESCO World Cultural Heritage?; the Liangzhu Ancient City Site)
never-changing	The answer to the question is from static knowledge such as scientific theories, historical facts and so on	(In rural areas during winter heating, it is necessary to guard against the risk of poisoning from which gas?; Carbon monoxide)
person	Specific individual, usually referring to a human being.	(Who among the current representatives of the Fuxin County People's Congress was one of the first batch of anti-epidemic heroes to rush to support Wuhan?; Xin Li)
location	Geographical position.	(Which province has recently strengthened the regulation of the intellectual property agency industry?; Hainan)
time	Points or intervals of a continuous sequence of events or conditions.	(In which year was the recent "Haikou Cup" sailing competition held?; 2023)
event	Something that happens, which can be planned or spontaneous.	(What themed event was recently launched in Suzhou High-speed Railway New Town to promote the development of private enterprises?; "Suzhou Sentiments, Private Enterprises Connected at Heart")
artificial work	Items or intellectual achievements created by humans, which have artistic, academic, or practical value.	(What is the latest TV series aired starring Xin Jiang?; As Long As We Are Together)
group	Entities formed by multiple individuals for a specific purpose.	(Which undergraduate university is recently established in the Ningxia Hui Autonomous Region recently?; Ningxia Minjiang Institute of Applied Technology)
nature	Phenomena or entities in the natural world.	(Please explain to me what Nucleases is?; Small RNA molecules with catalytic function, belonging to the category of biological catalysts?; capable of degrading specific mRNA sequences.)
quantity	Numeric value for times or stuff.	(How many base pairs in human Y chromosome have been observed from the latest sequencing results?; More than 30 million)
other	Other answer not classified to the above categories.	(Is there any fee for withdrawing WeChat balance to bank card?; Yes)

Table 7: Descriptions and examples of **question tags** (first three rows) and **answer types** (last nine rows). We represent (*<question>*; *<answer>*) as examples. Original language for these examples is Chinese. We translate them here for better preview.

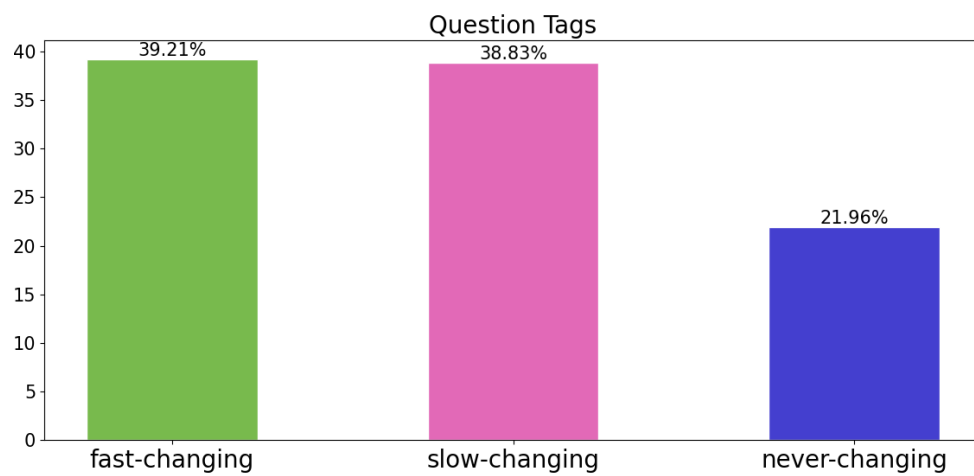


Figure 9: Distributions of question tags for full data







Figure 10: Distribution of answer types for full data



(a) Fast Changing



(b) Slow Changing



(c) Never Changing

Figure 11: Distributions of answer types for different question tags.

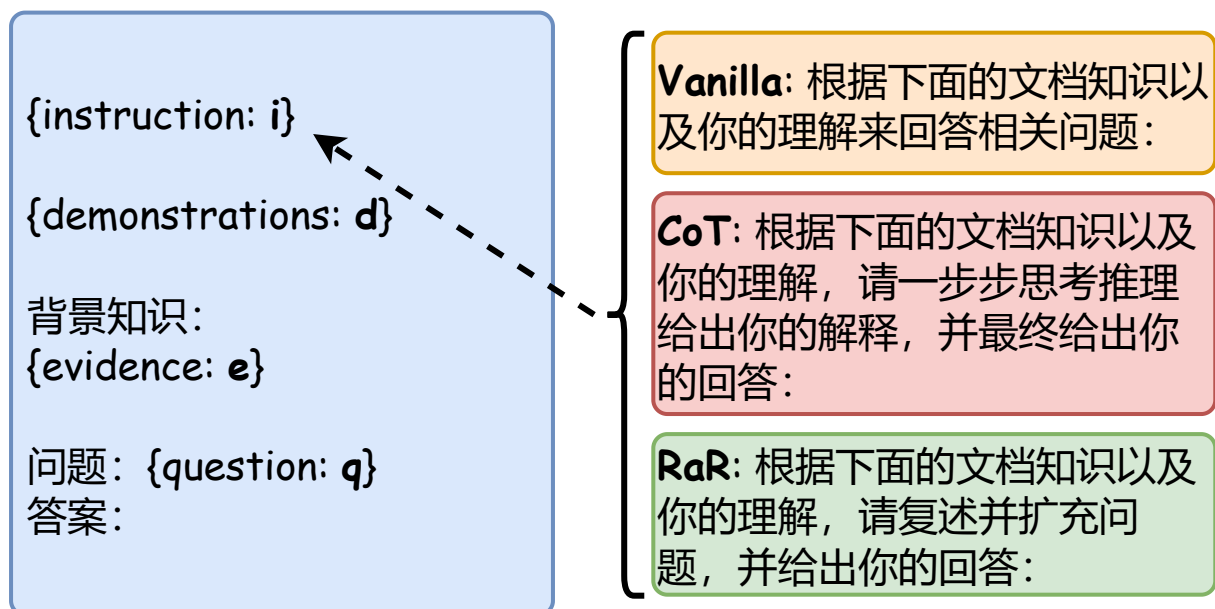


Figure 12: The Chinese prompt framework for Figure 2.