

Untangle the *KNOT*: Interweaving Conflicting Knowledge and Reasoning Skills in Large Language Models

Yantao Liu^{1†*}, Zijun Yao^{2,3‡}, Xin Lv^{2,3}, Yuchen Fan^{5*}
Shulin Cao^{2,3}, Jifan Yu⁴, Lei Hou^{2,3‡}, Juanzi Li^{2,3}

¹University of Chinese Academy of Sciences, Beijing, China

²BNRist; ³KIRC, Institute for Artificial Intelligence; ⁴Institute of Education

⁵Beijing University of Posts and Telecommunications, Beijing, China

Tsinghua University, Beijing 100084, China

yaozj20@mails.tsinghua.edu.cn, {houlei,lijuanzi}@tsinghua.edu.cn

Abstract

Providing knowledge documents for large language models (LLMs) has emerged as a promising solution to update the static knowledge inherent in their parameters. However, knowledge in the document may conflict with the memory of LLMs due to outdated or incorrect knowledge in the LLMs' parameters. This leads to the necessity of examining the capability of LLMs to assimilate supplemental external knowledge that conflicts with their memory. While previous studies have explained to what extent LLMs extract conflicting knowledge from the provided text, they neglect the necessity to **reason** with conflicting knowledge. Furthermore, there lack a detailed analysis on strategies to enable LLMs to resolve conflicting knowledge via prompting, decoding strategy, and supervised fine-tuning. To address these limitations, we construct a new dataset, dubbed *KNOT*, for **knowledge** conflict resolution examination in the form of question answering. *KNOT* facilitates in-depth analysis by dividing reasoning with conflicting knowledge into three levels: (1) Direct Extraction, which directly extracts conflicting knowledge to answer questions. (2) Explicit Reasoning, which reasons with conflicting knowledge when the reasoning path is explicitly provided in the question. (3) Implicit Reasoning, where reasoning with conflicting knowledge requires LLMs to infer the reasoning path independently to answer questions. We also conduct extensive experiments on *KNOT* to establish empirical guidelines for LLMs to utilize conflicting knowledge in complex circumstances. Dataset and associated codes can be accessed at our [GitHub repository](#).

Keywords: Knowledge Conflicts, Large Language Model, Reasoning Dataset

1. Introduction

Providing large language models (LLMs) with supplemented knowledge documents has become a *de facto* solution for tackling tasks that require up-to-date knowledge. This strategy is widely adopted in various fields including knowledge-grounded dialogue (Shuster et al., 2022; Zhang et al., 2023) and knowledge question answering (Guu et al., 2020; Trivedi et al., 2022; Izacard and Grave, 2020; Shao et al., 2023).

However, supplemented knowledge can not only be unfamiliar to LLMs, but can also potentially conflict with their existing knowledge stored in parameters, *i.e.*, conflicting knowledge. As shown in Figure 1, the supplemented knowledge document states that `Messi plays for Inter Miami CF`, which conflicts with the fact that `Messi plays for FC Barcelona`, stored as parametric knowledge of LLMs that are not timely re-trained. Moreover, there are supplemented knowledge documents that conflict with the parametric knowledge indirectly with deduced

conflicting conclusions. For example, the provided document in Figure 1 entails that `Messi's club plays in Major League Soccer` and that `Messi lives in Miami`, which cannot be directly extracted from the text. It is thus vital to explore to what extent LLMs reconcile with conflicting knowledge.

To this end, various datasets composed of conflicting knowledge have been constructed and tested on LLMs, which require the latter to answer questions with the provided conflicting knowledge. For example, NQ-Swaq requires LLMs to extract conflicting entities from the provided document without solely referring to their parametric knowledge directly (Longpre et al., 2021). The study by Zhou et al. (2023b) further investigates relation extraction from documents with the same entities but conflicting relations. MemoTrap (McKenzie et al., 2023) instructs LLMs to rewrite well-known proverbs that deviate from their inherent parameters. Although these benchmarks evaluate the ability of LLMs to extract different types of conflicting knowledge from the input document, they overlook LLMs' capacity to deduce subsequent contradictory knowledge by **reasoning**.

In light of their limitations, this work aims to construct a new dataset, dubbed *KNOT*, to examine

[†] Liu and Yao contributes equally to *KNOT*.

[‡] Corresponding author.

* Work is done when Liu and Fan are interns at Tsinghua University.

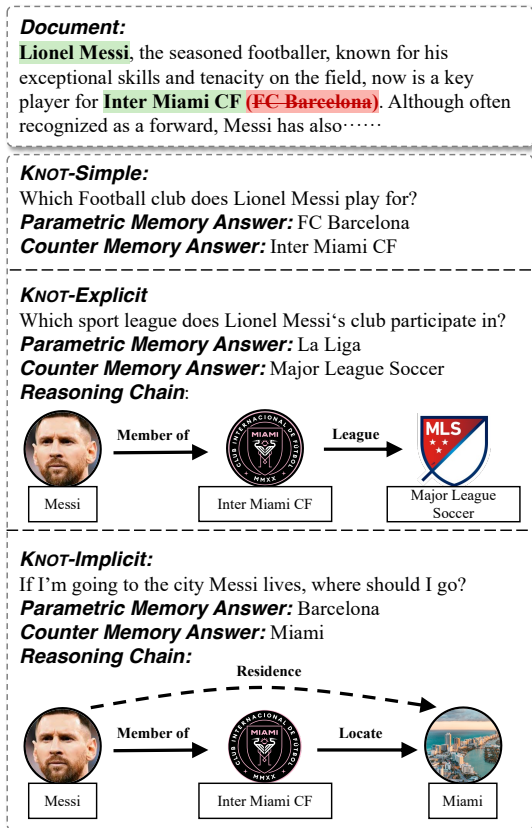


Figure 1: Example questions from KNOT where knowledge conflicts are resolved via extraction, explicit reasoning, and implicit reasoning.

how LLMs resolve **Knowledge conflicts** from the provided document when different levels of reasoning skills are required. As illustrated in Figure 1, each question in KNOT comes with a self-contained but contradictory document relative to the parametric knowledge of LLMs. To answer these questions, LLMs must coordinate the conflicting knowledge and may even need to deduce answers which are contradictory to the knowledge in their memory by reasoning.

To disentangle the subtle interaction between reasoning skills and knowledge conflicts, KNOT divides questions into three categories—KNOT-S, KNOT-E, KNOT-I—based on the level of reasoning skills required to reconcile the conflicting knowledge. **Simple** questions in KNOT-S do not require reasoning skills but merely the identification of the conflicting answer from the provided document. KNOT-E requires **explicit** reasoning skills by providing a clear reasoning path that weaves the conflicting knowledge directly in the questions. KNOT-I requires **implicit** reasoning skills, where LLMs need to induce the reasoning path implied by the question, and thereby deduce the answer from the conflicting knowledge according to the reasoning path. KNOT provides both a training set and a test

set for KNOT-S, KNOT-E, and KNOT-I.

With KNOT, this work also seek for the empirically best solution for LLMs to reconcile conflicting knowledge when different reasoning skills are required. We conclude conflicting knowledge resolving strategies into three main categories: decoding-based (Shi et al., 2023), prompting-based (Zhou et al., 2023b), and fine tuning-based (Zhang et al., 2023) methods. To thoroughly compare their effectiveness, we conduct experiments on commonly used LLMs of different model sizes and training strategies (*i.e.*, language models pre-trained only or assistant language models after alignment tuning).

The primary findings of the experiments include: (1) Mainstream LLMs are adept at resolving conflicting knowledge when no reasoning is required. However, they struggle to answer questions when multi-hop reasoning is required, suggesting that their ability to infer knowledge across multiple interconnected facts is not as robust; (2) Training-free method for knowledge conflicting resolution, *i.e.*, prompting-based and decoding-based methods, are not universally effective. LLMs show sensitivity to the prompting strategy, thus the prompting-based method delivers a far-from-ideal outcome on our dataset. Meanwhile, decoding methods result in mild amnesia of other background knowledge. In this case, the fallback is fine-tuning LLMs with our training data. (3) For the complex questions in KNOT-E and KNOT-I, increasing the model size of LLMs can improve performance by enhancing instruction-following capabilities. However, for implicit reasoning questions in KNOT-I, larger LLMs tend to answer the question in a shortcut with their more abundant parametric knowledge, rather than conduct implicit reasoning from the conflicting knowledge in the context.

In summary, our contributions are three-fold: (1) We extend the task of resolving conflicting knowledge by incorporating different levels of reasoning skills. (2) We construct KNOT, a dataset for testing the reasoning skills in knowledge conflict resolution. (3) Comprehensive experiments on KNOT provide empirical guidelines for LLMs to resolve knowledge conflicts.

2. Preliminaries

In this section, we provide definitions for the probing task designed to resolve knowledge conflicts, and formulate 3 levels of reasoning skills involved in dealing with conflicting knowledge.

2.1. Task Definition

KNOT examines the capacity of LLMs to resolve knowledge conflicts in the form of question answer-

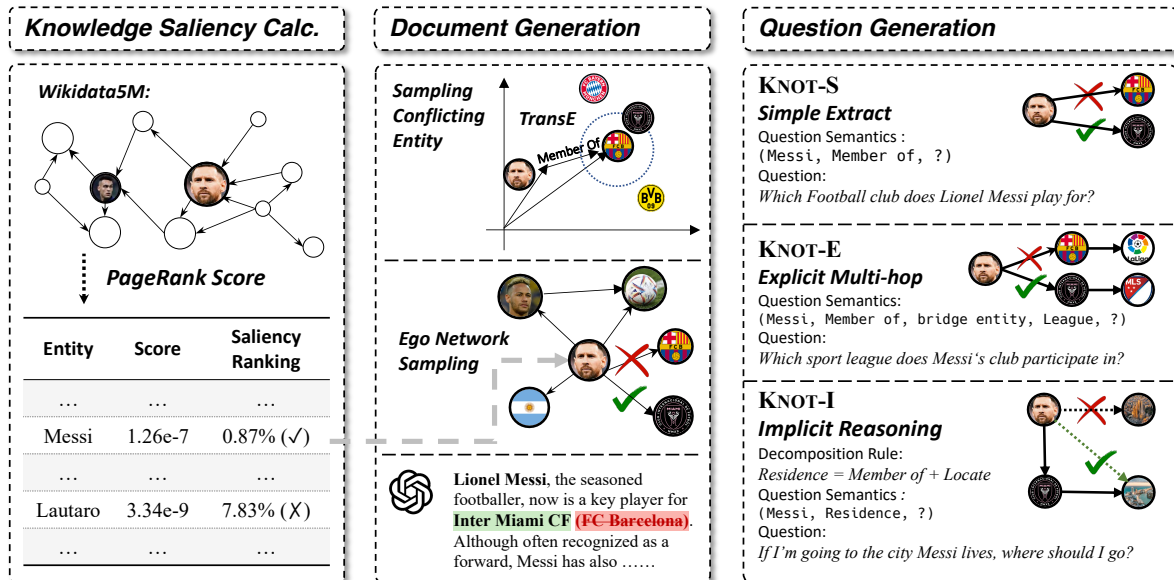


Figure 2: The overall framework for constructing KNOT.

ing with supplemented documents. Specifically, we define $\text{KNOT} = \{(q, \{a_i\}, d)\}$, where q represents the natural language question. $\{a_i\}$ constitutes the set of correct answer, with each a_i being an alias of the correct answer. d denotes the supplemented document containing the necessary knowledge $k \in d$ for the question.

The necessary knowledge $k \in d$ can be represented as a triple $k = (h, r, t)$, which contradicts the memory of LLMs. For example, in Figure 1, $k = (\text{Messi}, \text{Member of}, \text{Inter Miami CF})$. This contradicts with the memory of LLMs, which stores $k' = (\text{Messi}, \text{Member of}, \text{FC Barcelona})$.

We provide the document d and question q as input, and assess whether LLMs output correct answer: $\text{LLMs}(d, q) \in \{a_i\}$. LLMs make mistakes when they persist in utilizing the knowledge k' in their memory instead of switching to use $k \in d$.

2.2. Reasoning Skills

KNOT assigns reasoning skills into 3 levels. For questions in KNOT-S, the answer entity is an element of the knowledge k . For example, the KNOT-S question in Figure 1 is the tail entity of k . For questions in KNOT-E and KNOT-I, the answer does not directly appear in the document, but requires a multi-hop reasoning process involving k as one of the steps. The difference between KNOT-E and KNOT-I is that explicit questions depict the reasoning path, while implicit questions require LLMs to decompose questions into the reasoning path.

3. KNOT Construction

KNOT is constructed in three steps, as shown in Figure 2. First, we calculate the saliency score of each entity in the knowledge base (KB). Next, we create documents containing conflicting knowledge around salient entities. Finally, we generate questions based on the previously constructed documents and the knowledge triples in the KB.

3.1. Knowledge Saliency Calculation

As LLMs are trained on large-scale corpora, they tend to have a better memory of more important parts of knowledge (Petroni et al., 2019). Thus, it is more challenging for LLMs to reconcile facts that conflict with salient knowledge. To this end, we construct the questions in KNOT around the most salient knowledge. We apply the PageRank algorithm (Brin and Page, 1998) on Wikidata5M (Wang et al., 2021), a subset of Wikidata consisting of all entities in Wikipedia before 2021, and use PageRank values for entities as their saliency measure.

3.2. Conflicting Document Construction

We construct documents to present knowledge that conflicts with that stored in the parameters of LLMs. To this end, we first sample ego networks from the KB and then convert these networks into documents using a data-to-text approach.

Ego Network Sampling. The ego network is sampled with the following criteria to ensure the saliency of its provided knowledge: (1) The focal node entities are sampled from the top 1% salient entities, i.e., entities with PageRank score higher

than 99% of other entities; (2) The first-hop neighbors of the focal node entities are sampled from the top 20% salient entities.

Next, we randomly select an entity e from the ego network and replace it with a different entity as the conflicting entity. To obtain highly confusing conflicting entities, we employ TransE (Bordes et al., 2013) to sample an entity e' that is most similar to the selected entity. We then swap e with the conflicting entity e' from the ego network. To keep the ego network self-contained, e is chosen from the dangling entities that are only connected to the focal node entity. This process ensures that only one knowledge triple that conflicts with the memory of LLMs is introduced into the ego network.

Document Generation. Given that LLMs present a remarkable ability to generate coherent and fluent continuations of a prompt, we leverage LLMs to perform data-to-text conversion based on the previously constructed ego networks. Specifically, we follow Xie et al. (2022) and transform the knowledge-conflicting ego network into sets of triples. We prompt `gpt-3.5-turbo`¹ to generate documents with more than 100 words according to these triples.

3.3. Question Generation

Finally, we generate questions for KNOT, whose answers require the coordination of conflicting knowledge within the provided documents. Each question has a topic entity e_t , which is the topic of the natural language question, and an answer entity e_a , which serves as the correct answer. The question is generated by prompting `gpt-3.5-turbo` with prototypical questions as demonstrations (Yao et al., 2023). In this section, we introduce how we select e_t and e_a , and how we design prototypical questions for KNOT-S, KNOT-E, and KNOT-I.

KNOT-S. The semantics of the simple questions in KNOT-S are equivalent to a link prediction query $(e_t, r, ?)$ for KB (Rossi et al., 2021). In particular, the triple (e_t, r, e_a) is contained in the ego network and is thus inherited by the document. We choose e_t to be either the focal entity or the conflicting entity in the ego network, and e_a as its counterpart. The questions are generated by mimicking questions like “Which entity has relation r with e_t ?”

KNOT-E. The semantics of the questions requiring explicit reasoning in KNOT-E are equivalent to multi-hop reasoning queries (Yang et al., 2018), in the form $(e_t, r_1, e_b, r_2, ?)$, where e_b is the bridge entity. Here, we require e_t and e_b to be either the focal entity and the conflicting entity, or vice versa. e_a is an entity in Wikidata5M but does not belong to

the ego network. The demonstration questions resemble “Which entity has relation r_2 with the entity that has relation r_1 with e_t ?”. To answer a question from KNOT-E, LLMs need to extract the conflicting knowledge (e_t, r_1, e_b) from the document, and recall (e_b, r_2, e_a) from its parameters to reach e_a .

KNOT-I. The semantics of questions that require implicit reasoning in KNOT-I are equivalent to a triple queries $(e_t, r, ?)$, while r can be further decomposed into $r = r_1 + \dots + r_n$ (Geva et al., 2021; Yao et al., 2023). Thus, the triple query entails a multi-hop reasoning path in the form $(e_t, r_1, e_b, \dots, r_n, ?)$. We assign one of e_t and e_b to be the conflicting entity and the other to be the focal entity of the ego network. For relation decomposition, we use BMR (Lv et al., 2021), which provides high-quality decomposition rules from human annotation. The questions are generated by mimicking the same question template similar to KNOT-S.

As (e_t, r, e_a) is not contained in the provided document, LLMs need to first decompose the question into multi-hop reasoning path and then reach the answer entity accordingly to answer these questions. It should be noted that directly recalling (e_t, r, e_a) from memory will lead to an incorrect answer, as this shortcut process does not utilise conflicting knowledge (e_t, r_1, e_b) .

3.4. Human Annotation

To prevent potentially ill-generated data points and reduce the bias towards question-generating model (a.k.a., `gpt-3.5-turbo`), we employ human annotators to (1) filter out low-quality data points and (2) annotate questions answers with chain-of-thoughts (rationales) (Wei et al., 2022).

Particularly, for **human filtering**, we ask the annotators to check from two aspects: (1) whether the document contains the conflicting knowledge k' by demand; (2) whether the question is faithful to the given prototypical question. After filtering, 8.0% questions are discarded due to their low-quality.

For **annotating question answers**, we ask human annotators to answer the question and provide their rationale on how conflicting knowledge from the text and their memories are combined and resolved to infer the final answer. Meanwhile, to obtain rationales with various syntax structures, we manually construct 20 templates (e.g., *We can infer from the text that ...*; *The author articulates the idea that ...*; etc.) as reference for our annotators. We also allow human annotators to reject a question when they believe the given question cannot be answered based on the given document. This further filters out a small portion of low-quality data points. The annotated data are used in the training set.

¹<https://platform.openai.com/docs/models/gpt-3-5>.

Split	KNOT-S	KNOT-E	KNOT-I
#Questions for evaluation	3,887	1,136	510
A.S. of Topic Entity	8.66	7.51	5.86
A.S. of Answer Entity	0.83	0.19	0.18
A.S. of Answer in Memory	0.56	0.18	0.19
Average Reasoning Hops	1	2	2.7
#Questions for training	603	190	26

Table 1: The overall statistics of KNOT. A.S stands for average saliency ranking in percentile.




Figure 3: Accuracy of LLaMA-2-70B-Chat with regard to question topic entity saliency on KNOT without providing the documents. The accuracy is positively correlated with the saliency.

3.5. Dataset Analysis

General Statistics: Table 1 shows the general statistics of KNOT. From a saliency perspective, all the entities directly related to the questions in KNOT are among the top 10% salient entities, ensuring quality of the corresponding knowledge remembered by LLMs. From the perspective of reasoning paths, the hop number of the reasoning chain of the question in KNOT-S and KNOT-E to answers is fixed. However, for the questions in KNOT-I, the hop number of the corresponding reasoning chains is dynamic, with an average of 2.70. Therefore, this indicates that a higher level of knowledge utilization skills are required.

Consistency between saliency and parametric knowledge. To investigate whether salient knowledge act as a good proxy for the parametric knowledge of LLMs, we evaluate the performance of LLaMA-2-70B-Chat (Touvron et al., 2023b) on KNOT without providing the documents. In this way, the model can only rely on its memory to answer the questions. As figure 3 shows, the accuracy of LLaMA-2-70B-Chat increases as the topic entity becoming more salient, which testifies that the saliency of entities in Wikidata5M is highly correlated with the memory of mainstream LLMs.

4. Experiment Setup

In this section, we introduce the evaluated model and the evaluated knowledge conflict resolving methods. We also setup the evaluation metric.

4.1. Implemented Models

We split evaluated LLMs into two categories: (1) **Pre-trained Only LLMs**, including encoder-decoder models—UL2 (20B) (Tay et al., 2023)—and decoder-only models—GPT-J (6B) (Wang and Komatsuzaki, 2021), GPT-NeoX (20B) (Black et al., 2022), LLaMA (7B, 13B, 30B, 65B) (Touvron et al., 2023a) and LLaMA-2 (7B, 13B, 70B) (Touvron et al., 2023b). (2) **Assistant LLMs**, which are further fine-tuned to align to human instructions, including Flan-T5 (11B), Flan-UL2 (20B) (Chung et al., 2022), GPT-JT (6B) (Computer, 2022), GPT-NeoXT-Chat-Base (20B) (Computer, 2023), Alpaca (7B,13B) (Taori et al., 2023), Vicuna (7B,13B) (Chiang et al., 2023), Tulu (7B,13B,30B) (Wang et al., 2023), LLaMA-2-Chat (7B, 13B, 70B) (Touvron et al., 2023b), and variants of InstructGPT (Ouyang et al., 2022). They are denoted as `text-curie-001` (6.7B), `text-davinci-002` (175B), `gpt-3.5-turbo`², and `text-davinci-003` (175B).

We evaluate the overall performance of these LLMs in the setting of one-shot prompting. Specially, we prepare 10 human-annotated examples for each level of KNOT. To mitigate the sensitivity of example selection in in-context learning of LLMs, we randomly choose one example for each data in KNOT. As one-shot prompting prevents LLMs from learning knowledge conflict resolution skills from more than 1 training data, the achieved performance approximates the intrinsic capacity to resolve knowledge conflicts.

4.2. Evaluation Metric

Considering the flexibility of the output of LLMs, The prediction is marked correct when it matches any of the aliases of the answer entity from Wikidata. We use accuracy as the evaluation metric.

In particular, to validate the feasibility of this metric, we randomly sample 100 questions from KNOT and corresponding answers from 22 different LLMs. Then we ask human annotators to label the correctness of the answers. We calculate the correlation between human-evaluation score and the proposed alias-based evaluation accuracy for each LLMs with regard to the sampled questions. The correlation coefficient is 0.87, which indicates that the automatic score is a good estimation for model performance in our task.

²The model size of `gpt-3.5-turbo` is unclear.

5. Results and Analysis

5.1. Overall Results

We report one-shot prompting accuracy in Table 2 as the overall results, which keeps the minimum volume of exposed training data to the LLMs. Overall, we have 3 observations: (1) Mainstream LLMs, especially LLaMA-2-Chat and InstructGPT variants, achieve satisfactory results on KNOT-S, which indicates that LLMs have the capacity to extract conflicting knowledge directly. (2) However, performance drops by a large margin when reasoning with conflicting knowledge in KNOT-E and KNOT-I is required, which suggests that advanced utilization of conflicting knowledge is still challenging for LLMs. (3) Compared to pre-trained only LLMs, *i.e.*, LLaMA, LLaMA-2, GPT-J, GPT-NeoX, UL2 and T5, their counterpart assistant models *i.e.*, Alpaca, LLaMA-2-Chat, GPT-JT, GPT-NeoXT, Flan-UL2 and Flan-T5, promise an improvement in performance, due to their better understanding of the instructions, which requires utilizing the supplemented document.

Intriguingly, we find that the 13B variant of LLaMA and LLaMA-2 performs worse than the 7B variant. By further examining the output from 13B variants, we find that they prefer continuing asking similar questions to answering the question. This observation indicates that scaling up pre-trained only LLMs can enhance the ability to continue writing, but the level of knowledge mastery is not guaranteed to improve accordingly, nor are the task-solving and instruction-following abilities. By contrast, this is not observed in Alpaca-13B, Vicuna-13B, Tulu-13B and LLaMA-2-13B-Chat which suggests that instruction alignment tuning helps to resolve knowledge conflicts by improving the ability to follow instructions.

Moreover, although Flan-T5 and Flan-UL2 present competitive performance on KNOT-S, they fall short of other assistant language models on KNOT-I. This is because Flan-T5 and Flan-UL2 are fine-tuned with existing datasets such as CommonsenseQA (Talmor et al., 2019), while the instruction datasets of Alpaca and Vicuna are generated from analogous or real dialog between human and LLMs, which is more diverse. Therefore, it is necessary to diversify instruction datasets.

5.2. Conflicting Knowledge Resolving Methods Analysis

We thoroughly explore 3 main existing conflict-resolving aspects, **Prompting**, **Decoding** and **Fine-tuning**.

Prompting is the most common way to fit LLMs into down-stream tasks as it is straightforward and easy to use. In terms of the conflicting knowledge

Type	LLM	KNOT-S	KNOT-E	KNOT-I	
Pre-trained Only	UL2	52.66	13.03	3.33	
	GPT-J	48.57	29.84	18.63	
	GPT-NeoX	55.65	30.72	20.39	
	LLaMA-7B	64.19	38.20	21.18	
	LLaMA-13B	49.65	27.20	17.45	
	LLaMA-30B	70.36	40.49	23.33	
	LLaMA-65B	67.89	38.82	24.51	
	LLaMA-2-7B	63.88	35.65	20.59	
	LLaMA-2-13B	77.18	42.52	24.51	
	LLaMA-2-70B	59.63	32.75	20.20	
Assistant LLMs	Flan-T5-XXL	85.95	40.67	15.69	
	Flan-UL2	84.23	46.48	15.29	
	GPT-JT	60.59	32.57	17.25	
	GPT-NeoXT	70.03	38.73	27.65	
	Alpaca-7B	82.74	46.83	29.22	
	Alpaca-13B	85.82	43.75	28.82	
	Vicuna-7B	75.92	42.43	26.67	
	Vicuna-13B	84.20	46.30	25.29	
	Tulu-7B	86.72	51.50	27.06	
	Tulu-13B	87.34	53.26	26.67	
	Tulu-30B	89.61	53.61	24.31	
	LLaMA-2-7B-Chat	90.66	50.18	28.82	
	LLaMA-2-13B-Chat	91.66	53.43	35.88	
	LLaMA-2-70B-Chat	94.93	58.54	38.24	
		text-curie-001	74.22	42.96	29.41
		text-davinci-002	87.45	55.02	30.59
	gpt-3.5-turbo	83.66	34.95	10.20	
	text-davinci-003	95.11	63.56	35.49	

Table 2: One-shot prompting results on KNOT. All the results are presented in accuracy (%).

	LLMs	KNOT-S	KNOT-E	KNOT-I
Instr+Opin	LLaMA-7B	57.0 (− 7.2)	40.9 (+ 2.7)	17.7 (− 3.5)
	Alpaca-7B	78.2 (− 4.6)	46.8 (0.0)	26.1 (− 3.1)
	Vicuna-7B	32.6 (−43.4)	25.0 (−17.4)	12.9 (−13.7)
	Tulu-7B	83.5 (− 3.3)	47.6 (− 3.9)	26.3 (− 0.8)
	Tulu-13B	88.8 (+ 1.5)	46.8 (− 6.5)	26.7 (0.0)
Instr+Attr	LLaMA-7B	63.0 (− 1.2)	41.7 (+ 3.5)	20.8 (− 0.4)
	Alpaca-7B	81.4 (− 1.4)	48.7 (+ 1.9)	25.7 (− 3.5)
	Vicuna-7B	43.2 (−32.8)	34.5 (− 7.9)	12.4 (−14.3)
	Tulu-7B	82.6 (− 4.1)	49.6 (− 1.9)	24.1 (− 2.9)
	Tulu-13B	88.2 (+ 0.9)	46.9 (− 6.3)	24.7 (− 2.0)

Table 3: Accuracy (%) with prompting strategies. We show absolute accuracy changes compared with results from Table 2 in parenthesis.

resolving task, Zhou et al. design two elaborated prompting strategies to improve the ability of LLMs to handle knowledge conflicting problem.

Decoding is a more subtle way to control the output of LLMs compared to prompting. Some previous work (Welleck et al., 2019b, 2020) has proved text generated with proper decoding strategies like beam-search to be more fluent. For resolving conflicting knowledge, Shi et al. (2023) recently propose **Context Aware Decoding (CAD)** by calibrating the output logits, which forces LLMs

LLMs	KNOT-S	KNOT-E	KNOT-I
LLaMA-7B	66.5 (+2.3)	37.1 (-1.1)	9.2 (-12.0)
Alpaca-7B	82.7 (+0.0)	50.3 (+3.4)	20.4 (-8.8)
Vicuna-7B	77.3 (+1.4)	41.6 (-0.8)	18.8 (-7.9)
Tulu-7B	87.7 (+0.9)	42.9 (-8.6)	11.4 (-15.7)
Tulu-13B	91.5 (+4.2)	44.3 (-9.0)	15.9 (-10.8)

Table 4: Accuracy (%) with context-aware decoding. We show absolute accuracy changes compared with results from Table 2 in parenthesis.

to pay more attention to the provided document.

Fine-Tuning is the most direct way to change the behaviour of LLMs. Recent studies (Zhou et al., 2023a) point out that a small-scale well-curated supervised training dataset is enough to teach models to produce high quality output. Following this idea, we fine-tune LLMs on the training set.

5.2.1. Prompting Analysis

We use regular question answering prompting (*read the passage and answer the following question*) as the baseline methods. We report the performance of two prompting strategies proposed by Zhou et al. (2023b) (*i.e.*, Instr+Opin and Instr+Attr) as well as their absolute performance improvement *w.r.t.* baseline prompt in Table 3.

The results do not present a promising improvement in performance. We argue that the reason is that LLMs are sensitive to the prompting strategy. One well-designed strategy on one dataset may not be suitable for another dataset. The huge performance drop of Vicuna-7B indicates that for the dialogue assistant LLMs, which are fine-tuned with specific prompt templates, have difficulties in adopting our prompt strategy directly.

5.2.2. Decoding Strategy Analysis

We use greedy search as the baseline decoding strategy to examine the effectiveness of context-aware decoding (Shi et al., 2023). The performance of CAD and its absolute improvement is shown in Table 4.

The results show that CAD helps LLMs to trust the supplement document and achieve better performance in KNOT-S. However, we hypothesize that *CAD diminishes LLMs' ability to access intrinsic knowledge*. When the question requires both the conflicting knowledge, and related information in the memory, as required by KNOT-E and KNOT-I, performance drops inevitably. The experiment results testifies our hypothesis.

LLMs	KNOT-S	KNOT-E	KNOT-I
LLaMA-7B	90.1 (+25.9)	60.2 (+22.0)	32.0 (+10.8)
Alpaca-7B	92.5 (+9.8)	59.2 (+12.4)	32.2 (+2.9)
Vicuna-7B	94.6 (+18.6)	60.0 (+17.6)	32.9 (+6.3)
Tulu-7B	90.2 (+3.4)	54.1 (+2.6)	27.5 (+0.4)
Tulu-13B	93.5 (+6.2)	56.7 (+3.4)	31.2 (+4.5)

Table 5: Accuracy (%) and their absolute changes (compared with Table 2) after fine-tuning.

5.2.3. Fine-Tuning Analysis

Previous analysis show that there is no silver bullet that resolves knowledge conflicts with LLMs without fine-tuning. In order to further address the issue of knowledge conflicts that LLMs may encounter, we employ human annotated data to fine-tune LLMs to resolve knowledge conflicts.

Fine-Tuning Setup. The general purpose of fine-tuning LLMs with the annotated data is two-fold. (1) To enable LLMs to understand instructions to resolve knowledge conflicts; (2) To enable LLMs to generate rationale as the chain-of-thought in more complicated reasoning.

To prevent catastrophic forgetting, we mix the fine-tuning dataset with Wikitext (Merity, 2016) and other supervised fine-tuning dataset, *e.g.*, Alpaca dataset (Taori et al., 2023). Finally, we fine-tune the 7B version of LLaMA, Alpaca, Vicuna, and Tulu on the mixed data. We train these models for two epochs at a learning rate of 2×10^{-5} . We use 3% of the total training steps for linear learning rate warmup and adopt a linear learning rate decay.

Fine-Tuning Results. The experiment results after fine-tuning LLMs are shown in Table 5. Generally, a promising performance improvement from the extra training process can be observed. LLaMA-7B has more performance gain compared with the assistant LLMs. It even reaches a comparable capability to resolve knowledge conflicts with complicated reasonings. Moreover, after fine-tuning, Vicuna-7B reaches a comparable result with LLaMA-2-70B-Chat in Table 2, indicating the role of fine-tuning relatively small scale models in resolving knowledge conflicts as the fallback.

5.3. Effect of Reasoning Types

We aim to understand how LLMs react differently to conflicting knowledge across 3 tasks: direct extraction in KNOT-S, explicit reasoning in KNOT-E, and implicit reasoning in KNOT-I. To this end, we evaluate the tendency of LLMs to rely on their inherent parametric knowledge k' versus adapting to the conflicting knowledge k provided in the context. Particularly, we denote the set of aliases of answer to the question when consider parametric knowl-

LLMs	KNOT-S	KNOT-E	KNOT-I
LLaMA-2-7B	0.76	14.79	15.69
LLaMA-2-13B	1.22	14.00	15.88
LLaMA-2-70B	0.65	13.83	14.90
LLaMA-2-7B-chat	0.42	14.88	18.24
LLaMA-2-13B-chat	0.53	13.47	23.33
LLaMA-2-70B-chat	0.57	13.38	28.24

Table 6: The Accuracy (%) of LLMs w.r.t $\{a'_i\}$ on KNOT. $\{a'_i\}$ is the set of aliases of answer to the question when the parametric knowledge k' is considered, rather than the conflicting knowledge k .

edge k' , as $\{a'_i\}$. Following the same metric we introduce in Section 4.2, we evaluate the accuracy of LLM with regard to $\{a'_i\}$. This accuracy score thus reflects LLMs' tendency for shortcut reasoning based on their parametric knowledge, rather than carefully reading the conflicting knowledge to deduce the answer accordingly.

Our analysis includes both pre-trained only LLMs (LLaMA-2-7B, LLaMA-2-13B and LLaMA-2-70B) and assistant LLMs that are further developed with Reinforcement Learning from Human Feedback (RLHF) (LLaMA-2-7B-chat, LLaMA-2-13B-chat and LLaMA-2-70B-chat) under one-shot setting, with results presented in Tables 6. There are three main observations:

Knowledge conflicting with direct extraction is almost resolved by mainstream LLMs. As in Table 6, LLaMA-2-13B displays the highest score, 1.22, in answering the question directly based on parametric knowledge on KNOT-S. Meanwhile, the accuracy of all the RLHF variants of the LLaMA-2 surpass 90%, as shown in Table 2. These results imply that LLMs, especially assistant LLMs, are able to resolve knowledge conflicts with direct extraction. This is because direct extraction only requires LLMs to extract the answer from the context, which is a common align task in developing assistant LLMs. Therefore, LLMs are able to extract the conflicting knowledge k from the context and ignore the parametric knowledge k' .

Resolving conflicting knowledge with complex reasoning requirement stimulates LLMs to stick to their parametric knowledge. We hypothesize that LLMs, instead of performing complex reasoning, tend to find a shortcut to answer the question, which is to sticking to their parametric knowledge k' . It is obvious from Table 2 that the accuracy of LLMs on KNOT-E and KNOT-I is significantly lower than that on KNOT-S. Since the answer of KNOT-E and KNOT-I does not appear in the context, LLMs are required to perform complex reasoning to find conflicting knowledge k as clues to answer the question. Meanwhile, both KNOT-E and KNOT-I have a higher chance than KNOT-S in

generating e'_a as the answer, as shown in Table 6. These phenomena validate our hypothesis.

The impact of scaling up LLMs is double-edged: While the stickiness to parametric knowledge mildly alleviates in KNOT-E, it exacerbates in KNOT-I. As shown in Table 6, the accuracy of LLMs w.r.t the answer $\{a'_i\}$ from parametric knowledge k' decreases as the size of LLMs increases for KNOT-E, but increases for KNOT-I. This phenomenon mainly roots in the difference between explicit and implicit reasoning skills. Although KNOT-E and KNOT-I share the feature that answers cannot be directly extracted from their context, questions in KNOT-E provide explicit reasoning paths, while KNOT-I calls for LLMs to seek out the reasoning path by themselves. As size escalates in LLMs, they are more capable of following reasoning path provided in questions. Therefore, the accuracy w.r.t $\{a'_i\}$ decreases as the size of LLMs increases for KNOT-E. Meanwhile, as LLMs increase in size, they accumulate parametric knowledge, which fosters their inclination to find shortcuts to answer in KNOT-I, instead of seeking the reasoning path by themselves, which shows more complexity.

6. Related Works

There are mainly two threads of works focusing on conflicting knowledge. One thread identifies various tasks associated with conflicting knowledge, while the other thread concentrates on context-faithful generation with the presence of conflicting knowledge.

Conflicting Knowledge Tasks. Research into knowledge conflicts can be traced back to studies on backdoor attacks, where LLMs are susceptible to errors when provided conflicting knowledge (Niu and Bansal, 2018; Kassner et al., 2021; Du et al., 2022). More recently, knowledge conflicts, which has been demonstrated by entity conflicting (Longpre et al., 2021) or relation conflicting (Zhou et al., 2023b) in the form of question answering (Pan et al., 2021), stands in the spotlight. This also occurred in downstream applications, such as knowledge-grounded dialogue systems (Zhang et al., 2023). However, these works often neglect the necessity of complex reasoning processes with conflicting knowledge.

Context-faithful Generation is a number of strategies adopted to resolve conflicting knowledge in generative LLMs. One strategy is constraint decoding, where LLMs generate answers from a constrained dictionary (Post and Vilar, 2018; Welleck et al., 2019a; Nye et al., 2021). Prompting-based methods design sophisticated prompt templates to persuade LLMs to adhere faithfully to the provided document (Si et al., 2022;

Zhou et al., 2023b). Decoding-based methods modify the output logits of each generated token conditioned on the provided documents (Shi et al., 2023; Chen et al., 2022). However, context-faithful generation methods only take effect when advanced reasoning is not required.

7. Conclusion and Future Directions

This paper introduces ΚΝΟΤ, a dataset designed to examine conflicting knowledge resolution capability when advanced reasoning skills are required. The dataset does not just assess the ability to extract counter-parameter knowledge, but also evaluates the capacity to reason with conflicting knowledge. We conduct evaluation on mainstream LLMs alongside existing methods for resolving knowledge conflicts, culminating in guidelines for selecting an appropriate solution for knowledge conflict resolution.

In the future, we will explore a broader range of knowledge conflicting scenarios, such as counter-commonsense knowledge. Apart from fine-tuning, we are also interested in developing more lightweight solutions to resolve knowledge conflicts in a more general condition.

Code and Data Availability Statement

The artifacts associated with this paper include both datasets and experiment codes.

For dataset construction, it is desirable to enable humans to generate the documents and to generate the questions. However, this approach is extremely time-consuming and labor-intensive. To strike a balance between cost and quality, we employ human annotators to label the training data and filter out low-quality data points for the test set. As we demonstrate in Section 4.2, evaluation results on the automatically generated data are consistent with human evaluation, satisfying our requirements. Thus, the whole dataset include (1) automatically generated dataset; (2) human annotated dataset; and (3) human filtered dataset.

Our codes include scripts that are used to generate data, and codebases that are used to establish the initial baselines.

Per the request of anonymous protocol, we will release both the automatically generated dataset, human annotated dataset, and human filter dataset as soon as this paper is accepted after the reviewing process.

Ethical Consideration

We primarily focus on two ethical issues: the treatment of annotators and the usage of data resources.

For the treatment of annotators, they are all students undertaking post-secondary educations. The annotator work is conducted in their part-time. We provide hourly wages that meet the local average for data work, ensuring equal pay for equal work across all employees.

Data resources include both the checkpoint of models and the training data. We use these data in accordance with the requirements of the data providers, and this research has no commercial intentions. The usage of this portion of data is purely for academic purposes.

Limitations

We would like to highlight that our data annotation process relies on the help of `gpt-3.5-turbo`. Although it has been shown to be effective in the synthesis of datasets by appropriately prompting large language models (Taori et al., 2023; Chiang et al., 2023), there are also observations pointing out that this could potentially introduce bias. To resolve this bias, we hire and train a group of annotators with secondary education experience to filter out low-quality data and construct 850 dataset from scratch (which means they are generated by human rather than LLMs).

Acknowledgement

This work is supported by a grant from the Institute for Guo Qiang, Tsinghua University (2019GQB0003), Tsinghua University Initiative Scientific Research Program, and Zhipu AI.

Bibliographical References

- Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonnell, Jason Phang, Michael Pieler, Usven Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. [GPT-NeoX-20B: An open-source autoregressive language model](#). In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*.
- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting*

- held December 5-8, 2013, Lake Tahoe, Nevada, United States.
- Sergey Brin and Lawrence Page. 1998. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*.
- Hung-Ting Chen, Michael Zhang, and Eunsol Choi. 2022. [Rich knowledge sources bring complex knowledge conflicts: Recalibrating models to reflect conflicting evidence](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).
- Together Computer. 2022. [Releasing GPT-JT powered by open-source AI](#).
- Together Computer. 2023. [OpenChatKit: An Open Toolkit and Base Model for Dialogue-style Applications](#).
- Yibing Du, Antoine Bosselut, and Christopher D. Manning. 2022. [Synthetic disinformation attacks on automated fact verification systems](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. [Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies](#). *Transactions of the Association for Computational Linguistics*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Papat, and Ming-Wei Chang. 2020. [REALM: Retrieval augmented language model pre-training](#). In *ICML*.
- Gautier Izacard and Edouard Grave. 2020. [Distilling knowledge from reader to retriever for question answering](#). *arXiv preprint arXiv:2012.04584*.
- Nora Kassner, Oyvind Tafjord, Hinrich Schütze, and Peter Clark. 2021. [BeliefBank: Adding memory to a pre-trained language model for a systematic notion of belief](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. [Entity-based knowledge conflicts in question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Xin Lv, Yixin Cao, Lei Hou, Juanzi Li, Zhiyuan Liu, Yichi Zhang, and Zelin Dai. 2021. [Is multi-hop reasoning really explainable? towards benchmarking reasoning interpretability](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Ian McKenzie, Alexander Lyzhov, Alicia Parrish, Ameya Prabhu, Aaron Mueller, Najoung Kim, Sam Bowman, and Ethan Perez. 2023. [Inverse scaling prize: Second round winners](#).
- Stephen Merity. 2016. The wikitext long term dependency language modeling dataset. *Salesforce Metamind*.
- Tong Niu and Mohit Bansal. 2018. [Adversarial over-sensitivity and over-stability strategies for dialogue models](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*.
- Maxwell I. Nye, Michael Henry Tessler, Joshua B. Tenenbaum, and Brenden M. Lake. 2021. [Improving coherence and consistency in neural sequence models with dual-system, neuro-symbolic reasoning](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. [Training language models to follow instructions with human feedback](#). *Advances in Neural Information Processing Systems*.
- Liangming Pan, Wenhui Chen, Min-Yen Kan, and William Yang Wang. 2021. [Contraqa: Question](#)

- answering under contradicting contexts. *ArXiv preprint*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Matt Post and David Vilar. 2018. [Fast lexically constrained decoding with dynamic beam allocation for neural machine translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*.
- Andrea Rossi, Denilson Barbosa, Donatella Firmani, Antonio Matinata, and Paolo Merialdo. 2021. Knowledge graph embedding for link prediction: A comparative analysis. *ACM Transactions on Knowledge Discovery from Data (TKDD)*.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. [Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy](#). *ArXiv preprint*.
- Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Scott Wen-tau Yih. 2023. [Trusting your evidence: Hallucinate less with context-aware decoding](#). *ArXiv preprint*.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, et al. 2022. [Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage](#). *ArXiv preprint*.
- Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuo-hang Wang, Jianfeng Wang, Jordan Boyd-Graber, and Lijuan Wang. 2022. [Prompting gpt-3 to be reliable](#). *ArXiv preprint*.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. [Alpaca: A strong, replicable instruction-following model](#). *Stanford Center for Research on Foundation Models*.
- Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Dara Bahri, Tal Schuster, Steven Zheng, et al. 2023. [Ul2: Unifying language learning paradigms](#). In *The Eleventh International Conference on Learning Representations*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. [Llama: Open and efficient foundation language models](#). *ArXiv preprint*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). *ArXiv preprint*.
- Ben Wang and Aran Komatsuzaki. 2021. [GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model](#). <https://github.com/kingoflolz/mesh-transformer-jax>.
- Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021. [KEPLER: A unified model for knowledge embedding and pre-trained language representation](#). *Transactions of the Association for Computational Linguistics*.
- Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [How far can camels go? exploring the state of instruction tuning on open resources](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Sean Welleck, Kianté Brantley, Hal Daumé III, and Kyunghyun Cho. 2019a. [Non-monotonic sequential text generation](#). In *Proceedings of*

the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, Proceedings of Machine Learning Research.

Sean Welleck, Ilya Kulikov, Jaedeok Kim, Richard Yuanzhe Pang, and Kyunghyun Cho. 2020. Consistency of a recurrent language model with respect to incomplete decoding. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5553–5568.

Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019b. Neural text generation with unlikelihood training. *arXiv preprint arXiv:1908.04319*.

Tianbao Xie, Chen Henry Wu, Peng Shi, Ruiqi Zhong, Torsten Scholak, Michihiro Yasunaga, Chien-Sheng Wu, Ming Zhong, Pengcheng Yin, Sida I. Wang, Victor Zhong, Bailin Wang, Chengzu Li, Connor Boyle, Ansong Ni, Ziyu Yao, Dragomir Radev, Caiming Xiong, Lingpeng Kong, Rui Zhang, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. 2022. [UnifiedSKG: Unifying and multi-tasking structured knowledge grounding with text-to-text language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.

Zijun Yao, Yantao Liu, Xin Lv, Shulin Cao, Jifan Yu, Juanzi Li, and Lei Hou. 2023. [Korc: Knowledge oriented reading comprehension benchmark for deep text understanding](#). In *Findings of ACL*.

Jing Zhang, Xiaokang Zhang, Daniel Zhangli, Jifan Yu, Zijun Yao, Zeyao Ma, Yiqi Xu, Haohua Wang, Xiaohan Zhang, Nianyi Lin, Sunrui Lu, Juanzi Li, and Jie Tang. 2023. [Glm-dialog: Noise-tolerant pre-training for knowledge-grounded dialogue generation](#). *ArXiv preprint*.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivas Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2023a. Lima: Less is more for alignment. *arXiv preprint arXiv:2305.11206*.

Wenxuan Zhou, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2023b. [Context-faithful prompting for large language models](#). *ArXiv preprint*.

A. Prompts for Data Construction

A.1. Document Generation

We use the following prompt to generate document from the ego network.

INSTRUCTION: Imagine you are a journalist. And you are given some semantic triples from Wikidata about one specific topic. And you are planning to write one short passage provided to your reader based on the information that these semantic triples. As its name indicates, a semantic triple is a set of three entities that codifies a statement about semantic data in the form of subject–predicate–object expressions (e.g., "Bob is 35", or "Bob knows John"). Your task is to convert these semantic triples into a concise and coherent passage with approximately 200 words to help your readers understand the specific topic.

Attention 1: In the resulting passage, DO NOT USE any text directly contains entity, semantic triple, or Wikidata, etc.

Attention 2: In the resulting passage, all the statements presented in the given semantic triples should be included.

Attention 3: If there is any counterfactual info contained in the given semantic triples, you should just accept it and faithfully describe this counterfactual info. DO NOT USE any words like "interesting", or "surprising" to indicate this info is counterfactual to passage readers.

INPUT:

Semantic Triples From Wikidata:
(Lionel Messi, place of birth, Rosario)
(Lionel Messi, sex or gender, male)
(Lionel Messi, date of birth, 24 June 1987)
(Lionel Messi, member of sports team, Inter Miami CF)
(Lionel Messi, occupation, Footballer)
(Lionel Messi, position played on team / speciality, Forward)
(Lionel Messi, position played on team / speciality, Midfielder)
(Lionel Messi, award received, Ballon d'Or)
(Lionel Messi, award received, FIFA World Player of the Year)
(Lionel Messi, sponsor, Adidas)

OUTPUT: Lionel Messi, hailing from the city of Rosario, was born on the 24th of June, 1987. This seasoned footballer, known for his exceptional skills and tenacity on the field, is now a key player for Inter Miami CF. Although predominantly recognized as a forward, Messi has also showcased his versatility by frequently occupying the midfielder position. The Argentine star's performance over the years has not gone unnoticed. His trophy cabinet includes accolades such as the Ballon d'Or and the FIFA World Player of the Year, testifying to his profound impact on the sport. Besides his sports-related commitments, Messi has also embarked on commercial ventures, endorsing global brands like Adidas. His enduring legacy continues to influence and reshape the world of football.

Table 7: The instruction and an example of Document Generation

A.2. Question Generation

A.2.1. Question Generation for KNOT-S and KNOT-I

As questions in both KNOT-S and KNOT-I are equivalent to a triple query, we use the following prompt to generate these questions.

INSTRUCTION: A semantic triple describe the relation between one head entity and one tail entity. For example, Job Biden -> native language -> English is one semantic triple which means Job Biden (head entity)'s native language (relation) is English (tail entity), now you are given one incomplete semantic triple where the tail entity is missing and one hint which would tell what all the possible missing entity is. your task is to design 5 questions based on the given semantic triple and the hint to find out the missing tail entity. Notice: the given hint could be utilized to design more accurate questions with respect to the given possible missing entity, but any part of the hint should not be contained in the generated question!

Example:

Input: incomplete semantic triple: Faiyum -> located in the administrative territorial entity (the item is located on the territory of the following administrative entity.) -> missing entity Hint: the missing entity is "Egypt"

Output:

1. Which country is Faiyum a part of?
2. What is the sovereign state that contains Faiyum?
3. Faiyum is in which country?
4. In which country would I be if I was visiting Faiyum?
5. If I were to send a letter to Faiyum, which country's name should I write on the envelope?

Example 2:

INPUT: incomplete semantic triple:Paradise ->Country (sovereign state that this item is in (not to be used for human beings)) -> missing entity

OUTPUT:

1. Which sovereign state does Paradise belong to?
 2. Is Paradise located in a province or a territory of Canada?
 3. What is the name of the country where Paradise is located?
 4. Which country has Paradise as a part of its sovereign state?
 5. If I were to travel to Paradise, which country would I need to visit
-

Table 8: The instruction and an example of One hop question generation.

A.2.2. Multi-hop question generation for KNOT-E

We use the following prompt to generate questions explicitly stating multiple hops in KNOT-E.

INSTRUCTION: You are given one multi-hop reasoning chain over Wikidata. This chain is related to a multi-hop KGQA question. Your task is to generate several multi-hop KGQA questions in natural language based on the given multi-hop reasoning chain.
Input 1: Michel Verne -> place of birth (P19) -> Middle Entity -> country (P17) -> Answer Entity

Output 1:

1. Which country does Michel Verne's birth city belong to?
2. In which country can we find the city where Michel Verne was born
3. To which country is the city of Michel Verne's birthplace associated with?
4. The city where Michel Verne was born is in which country?
5. What is the nationality of the country where Michel Verne was born?

Input 2: Buick LaCrosse -> brand (P1716) -> Middle Entity -> location of formation (P740) -> Answer Entity

Output 2:

1. In which city was the company that created the Buick LaCrosse brand established?
2. What is the city of origin of the brand associated with the Buick LaCrosse?
3. The brand that owns Buick LaCrosse was formed in which city?
4. Where was the company that produced the Buick LaCrosse initially established?
5. The Buick LaCrosse is from a brand that was formed in what city?

INPUT 3: Lionel Messi -> Member of sports team (P54) -> Middle Entity -> League (P118) -> Answer Entity

OUTPUT 3:

1. What league is Lionel Messi's sports team a part of?
 2. In which league does Lionel Messi's sports team compete?
 3. To which league does the sports team Lionel Messi belongs to belong?
 4. Which league is associated with Lionel Messi's sports team?
 5. What is the name of the league that Lionel Messi's sports team plays in?
-

Table 9: The instruction and an example of Two hop question generation.

B. Prompts for Model Evaluation

B.1. Prompt Template for InstructGPT variants, LLaMA, GPT-J, GPT-JT, and GPT-NeoX

For pre-trained only models, they are not aligned to a specific format of instructions. We use a unified prompting template for InstructGPT variants, LLaMA, GPT-J, GPT-JT, and GPT-NeoX.

INSTRUCTION:

read the passage and answer the question. You may do some analysis on the passage and question before answering. But you should directly give out in the answer in the end and in the form of answer.

EXAMPLES:

Passage: [[example_passage_1]]

Question: [[example_question_1]]

Answer: [[example_answer_1]]

...

EXAMPLE K+1: Passage: [[passage]]

Question: [[question]]

Answer:

Table 10: Prompt template for InstructGPT variants, LLaMA, GPT-J, GPT-JT, and GPT-NeoX.

B.2. Prompt Template for Alpaca

We use the format in the supervised fine-tuning data of Alpaca (Taori et al., 2023).

INSTRUCTION:

Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.

Instruction: Instruction: Read the passage and answer the question. You may do some analysis on the passage and question before answering. But you should directly give out in the answer in the end and in the form of answer.

EXAMPLES:

Input 1: Passage: [[example_passage_1]]
Question: [[example_question_1]]
Response 1: Answer: [[example_answer_1]]

...

Input K:
Passage: [[example_passage_K]]
Question: [[example_question_K]]
Response K:
Answer: [[example_answer_K]]

EXAMPLE K+1:

Input K+1 Passage: [[passage]]
Question: [[question]]
Response K+1:

Table 11: Prompt template for Alpaca.

B.3. Prompt Template for Vicuna

We use the format in the supervised fine-tuning data of Vicuna (Chiang et al., 2023).

INSTRUCTION:

A chat between a curious user and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the user's questions. USER: Read the passage and answer the question. You may do some analysis on the passage and question before answering. But you should directly give out in the answer in the end and in the form of answer. ASSISTANT: OK, I got it.</s>

EXAMPLES:

USER: Passage: [[example_passage_1]]
Question: [[example_question_1]] ASSISTANT: Answer: [[example_answer_1]]</s>

...

USER: Passage: [[example_passage_K]]
Question: [[example_question_K]] ASSISTANT: Answer: [[example_answer_K]]</s>

EXAMPLE K+1:

USER: Passage: [[passage]] Question: [[question]] ASSISTANT: Answer:

Table 12: Prompt template for Vicuna.

B.4. Prompt Template for Tulu

We use the following format for prompting Tulu.

INSTRUCTION: < user > Read the passage and answer the question. You may do some analysis on the passage and question before answering. But you should directly give out in the answer in the end and in the form of answer < assistant > OK, I got it. < user >
EXAMPLES: Passage: [[example_passage_1]] Question: [[example_question_1]] < assistant > Answer: [[example_answer_1]] ... < user > Passage: [[example_passage_K]] Question: [[example_question_K]] < assistant > Answer: [[example_answer_K]]
EXAMPLE K+1: < user > Passage: [[passage]] Question: [[question]] < assistant >

Table 13: Prompt template for Tulu.

B.5. Prompt Template for GPT-NeoXT

GPT-NeoXT is also allowed with specific format, as shown below.

INSTRUCTION: ***Read the passage and answer the question.*** ***You may do some analysis on the passage and question before answering.*** ***But you should directly give out in the answer in the end and in the form of answer.***
EXAMPLES: <human> Passage: [[example_passage_1]] Question: [[example_question_1]] <bot>: [[example_answer_1]] ... <human> Passage: [[example_passage_K]] Question: [[example_question_K]] <bot>: [[example_answer_K]]
EXAMPLE K+1: <human> Passage: [[passage]] Question: [[question]] <bot>:

Table 14: Prompt template for Chat-NeoXT.

C. Dataset Diversity Analysis

One natural question arise from the construction process of KNOT is that, does `text-davinci-003` generate high quality questions? We examine the quality by their diversity and plot the trigrams in Figure 4. As we can see, although the questions are generated automatically, they are potentially different with each other.

C.1. Question Diversity

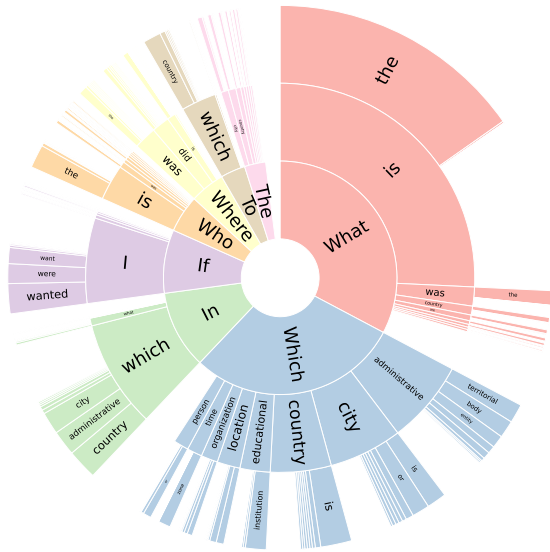


Figure 4: Trigram distribution of questions in KNOT.

C.2. Rationale Diversity

The rationales are generated from human annotations for LLMs fine-tuning. We plot the trigram of the rationales in Figure 5, which verifies that our human annotators provides diverse rationales.

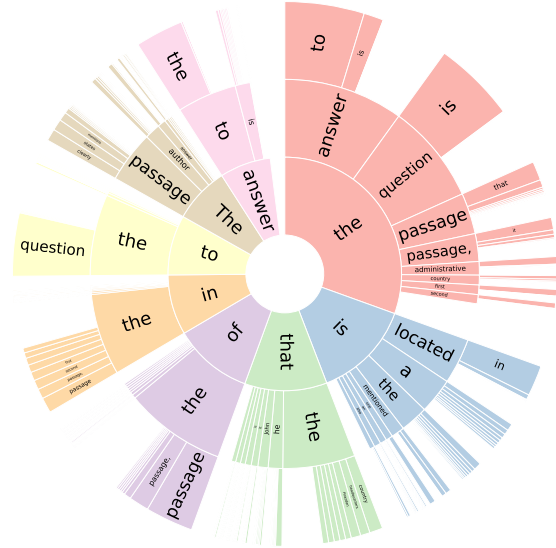


Figure 5: Trigram distribution of the annotated rationale.

D. Rationale Case Study

We show a rationale labeled by our annotators in Table 15.

PASSAGE:

The London Borough of Waltham Forest is an administrative territorial entity located in Greater London. It is situated in the UTC+01:00 time zone and shares borders with the London Borough of Enfield. Waltham Forest is home to the capital city of Walthamstow, a vibrant area with a rich culture and history. It is the birthplace of two renowned writers, Eliza Parsons and Tom Hood, and the place of passing of writer Jane Ray. Additionally, the borough also contains the remote hamlet of Niaqornat.

Waltham Forest has a lot to offer in terms of its culture and history. It has a variety of museums, parks, schools and other cultural venues. It also has a number of historic buildings, such as the Grade I listed St Mary's Church. The borough is also home to a number of landmarks, such as the William Morris Gallery and Epping Forest. It is also home to a number of festivals and events, such as the annual Walthamstow Garden Party. The London Borough of Waltham Forest is an area of great diversity and culture, and it is easy to see why so many people choose to make it their home. With its rich history and vibrant culture, it is a great place to live, work, and explore.

QUESTION:

In which country is Niaqornat located?

RATIONALE:

We can get the information that Niaqornat is located in the London Borough of Waltham Forest from the passage from the passage. And we all know that the London Borough of Waltham Forest is located in the United Kingdom. So we can infer that Niaqornat is located in the United Kingdom.

Table 15: One sampled rationale for KNOT-I.

E. Evaluation Metric Analysis




Figure 6: A scatter plot comparing human-annotated scores and automatic scores for 100 randomly sampled answers from each model in KNOT

In order to investigate the feasibility of automatic evaluate strategy we proposed in section 4.2, we randomly sample out 100 answers of each model in KNOT, asking human annotator to mark the correctness of model answers. Figure 6 shows the scatter plots of human annotated score verse the automatic score, which presents the highly positive linear correlation between two scores, suggesting that the automatic evaluation strategy is plausible when compare the performance between LLMs.

F. Effect of Example Number

In table 2 and table 3, we report the performance of different LLMs with different prompting strategies under one-shot setting. To thoroughly investigate the effect of example number, we test the performance of different LLMs from zero-shot to four-shot settings. Figure 7 shows the results.




Figure 7: Accuracy (%) with different example numbers.

These results are averaged over KNOT-E, KNOT-S, and KNOT-I, with dataset sizes of 3887, 1136, and 510 respectively serving as weights. As observed, performance does improve with an increase in the number of shots, but this improvement tends to converge when the number of shots exceeds two. Additionally, due to the text window size limit (2048 max), the number of shots is restricted to four.