

Reinforcement Retrieval Leveraging Fine-grained Feedback for Fact Checking News Claims with Black-Box LLM

Xuan Zhang, Wei Gao

School of Computing and Information Systems, Singapore Management University
80 Stamford Road, Singapore 178902
xuanzhang.2020@phdcs.smu.edu.sg, weigao@smu.edu.sg

Abstract

Retrieval-augmented language models have exhibited promising performance across various areas of natural language processing (NLP), including fact-critical tasks. However, due to the black-box nature of advanced large language models (LLMs) and the non-retrieval-oriented supervision signal of specific tasks, the training of retrieval model faces significant challenges under the setting of black-box LLM. We propose an approach leveraging Fine-grained Feedback with Reinforcement Retrieval (FFRR) to enhance fact-checking on news claims by using black-box LLM. FFRR adopts a two-level strategy to gather fine-grained feedback from the LLM, which serves as a reward for optimizing the retrieval policy, by rating the retrieved documents based on the non-retrieval ground truth of the task. We evaluate our model on two public datasets for real-world news claim verification, and the results demonstrate that FFRR achieves significant improvements over strong LLM-enabled and non-LLM baselines.

Keywords: Claim Verification, Reinforcement Retrieval, Fine-Grained Feedbacks, Large Language Model

1. Introduction

Recent advancements in large language models (LLMs), such as GPT-4 (OpenAI, 2023) and Palm 2 (Anil et al., 2023), have demonstrated impressive capabilities in generating coherent, informative, and fluent verbal reasoning (Wei et al., 2022; Wang et al., 2022; Zhou et al., 2022). This opens up the possibility for applying LLMs in fact-checking related tasks, such as predicting claim veracity and generating explainable rationales (Zhao et al., 2023; Yao et al., 2023; Pan et al., 2023; Bubeck et al., 2023; Zhang and Gao, 2023). However, these tasks heavily rely on up-to-date information and reliable knowledge, and LLMs may fall short in these aspects, potentially affecting their downstream performance with bias and hallucination (OpenAI, 2022; Ye and Durrett, 2022; Bang et al., 2023).

Inspired by retrieval-augmented language modeling approach (Lewis et al., 2020; Guu et al., 2020), recent efforts have attempted to incorporate retrieved external knowledge with LLMs (Nakano et al., 2021; Khattab et al., 2022; Press et al., 2023; Yao et al., 2023; Peng et al., 2023). A common idea is to integrate a plug-and-play retrieval model with the LLM for providing relevant context. However, the quality of retrieval results can compromise the effectiveness of LLM-based reasoning and verification (Yao et al., 2023). A key unresolved issue is how to tune the retrieval model for returning the “smoking gun” evidence that can assist the LLM in determining the veracity of the claim reliably.

Meanwhile, advanced LLMs typically exhibit a black-box nature and often provide service via APIs (Sun et al., 2022), preventing end-users

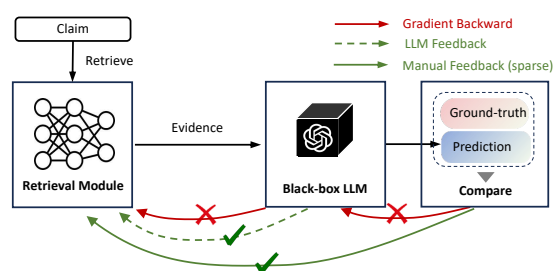


Figure 1: The black-box LLM handicaps end-to-end task training through gradient backward. However, manual (sparse) and LLM (fine-grained) feedback can be used to optimize the retrieval model.

from accessing their internal mechanisms, such as model parameters and gradients. Therefore, it is infeasible to employ the retrieval-augmented approach (Lewis et al., 2020; Guu et al., 2020) for end-to-end optimization because it is difficult to backpropagate prediction loss to the retrieval model through the LLM, as depicted in Figure 1. In addition, the ground truth is typically not retrieval-oriented, which may not be used directly for training the retrieval model.

To optimize the retrieval model for adaptation to LLM, a recent REtrieve-and-PLUG (REPLUG) approach for open-domain QA employs the LLM to rate retrieved top- K documents with reference to ground truth answers and updates the retriever via minimizing Kullback-Leibler Divergence (KLD) between retrieval score distribution and the rated score distribution (Shi et al., 2023). From the factuality standpoint, however, such an approach oversimplifies the interaction between the retrieval

model and LLM in document selection. Using top- K documents may incorporate unnecessary noise or miss useful evidence, and may fail to account for complex real-world claims. For example, for a claim like “Bruce Lee died from drinking too much water”, top search results predominantly stem from a report that Lee may have had too much water on the day of his death, and his kidneys lacked the ability to sufficiently shed it, which might not be key evidence¹. We contend that it is imperative to flexibly explore a larger pool of candidates to prioritize documents that present alternative evidence in diverse perspectives.

In this paper, we propose a new approach by employing Fine-grained Feedback with Reinforcement Retrieval, named FFRR, to enable more informed and accurate evidence retrieval for enhancing LLM-based fact checking on real-world news claims. We first prompt LLM to generate intermediate questions from various perspectives of a claim (Pan et al., 2023; Chen et al., 2023) to retrieve relevant documents. Then, FFRR collects fine-grained feedback from LLM on the retrieved documents as rewards at two levels: 1) Document-level rewards are used to refine the list of retrieved documents, such that individual documents that are more likely to contain evidence can have a better chance of getting selected through the exploration and exploitation process. 2) Question-level rewards are used to promote alternative evidence that cannot be surfaced by any single-angle inquiry, for which we optimize top-1 document ranking in each intermediate question. With the two-level rewards combined, we optimize the policy model of FFRR based on policy gradient to select evidence documents autonomously. Our contributions are three-fold:

- This is the first work using fine-grained LLM feedback to reward policy optimization of reinforcement retrieval for black-box LLM-enabled fact checking on real-world news claims.
- We turn the sparse, non-retrieval-oriented claim-level supervision signals to fine-grained rewards on candidate documents and intermediate questions, which facilitates retrieval policy optimization, without adding any overhead on inference.
- Results on two public news claim verification datasets demonstrate that FFRR outperforms strong LLM-enabled and non-LLM baselines by a large margin.²

¹ Fact-checking websites indicate the case is unproven because most researchers believe that Lee had several risk factors for the disease, such as swelling on the brain and heatstroke, and the claim is merely a recent hypothesis put forward by several medical researchers. <https://www.snopes.com/fact-check/bruce-lee-die-drinking-too-much-water/>

² Code and prompts data are available at [https://](https://github.com/jadeCurl/FFRR)

2. Related Work

Advanced LLMs have demonstrated impressive capabilities in generating human-like verbal reasoning (Wei et al., 2022; Wang et al., 2022; Zhou et al., 2022). Recent works (Press et al., 2023; Yao et al., 2023; Jiang et al., 2023; Pan et al., 2023; Zhao et al., 2023; Zhang and Gao, 2023; Zeng and Gao, 2024) find that combining LLM’s reasoning capability with accessibility to external knowledge is helpful in factuality-related tasks, such as fact-checking on FEVER dataset (Thorne et al., 2018). However, existing works overlook the optimization of retrieval models to better help LLMs. Compared to verifying the claims originated from Wikipedia facts in FEVER, verification of news claims is more realistic and complex, which needs more diverse and evidential information beyond LLM’s memory (Chen et al., 2023).

Claim decomposition and question generation have been employed for claim verification (Fan et al., 2020; Pan et al., 2021; Ousidhoum et al., 2022; Chen et al., 2022, 2023). Typically, these methods analyze the claim to identify entities (Ousidhoum et al., 2022) and focal points (Fan et al., 2020) to generate a series of corresponding questions, which are then answered via search engines to aid fact-checkers in making decisions. These models usually require a manual collection of claim-question datasets for training. Recently, Chen et al. (2023) prompt a black-box LLM (i.e., `text-davinci-003`) with in-context examples by carefully choosing four input-decomposition pairs from the human annotations of Chen et al. (2022) to form a few-shot prompt. We follow a similar approach to generate questions in our work.

Reinforcement learning (RL) has predominantly focused on treating LLMs as an RL agent and aligning them with human feedback (RLHF) (Ouyang et al., 2022; Xu et al., 2022; Bai et al., 2022a; Wu et al., 2023) or feedback generated by other LLMs (Bai et al., 2022b). Additional approaches diverge from this line of work by considering LLMs as black boxes and use RL to optimize models for choosing or generating suitable prompts for LLMs to enhance their outputs (Lu et al., 2023; Peng et al., 2023). This is motivated by the fact that LLMs’ massive parameter space makes it impractical for most users to fine-tune. Our FFRR method concentrates on optimizing the retrieval model by leveraging feedback from LLMs, targeting to improve retrieval results to enhance LLMs for news claim verification.

The REPLUG (Shi et al., 2023) approach is closely related to us, which optimizes the retrieval model for open-domain QA through KLD minimization with LLM feedback on top- K documents. We,

github.com/jadeCurl/FFRR.

however, optimize retrieval based on RL framework for news fact checking, which can autonomously explore and exploit any documents that are potentially evidential to avoid missing important evidence.

3. Problem Definition

A news claim verification dataset consists of a set of news claims $\mathcal{C} = \{c\}$ which have been verified by public fact-checking services, and a large corpus of documents $\mathcal{D} = \{d_1, d_2, \dots, d_{|\mathcal{D}|}\}$, which are pooled with the relevant documents gathered for each claim. It is worth to note that the documents from fact-checking sources and the raw documents published temporally after the corresponding fact-check reports should be strictly excluded from \mathcal{D} to avoid potential leak of ground truth. For any $c \in \mathcal{C}$, let $\mathcal{Q} = \{q_i\}_{i=1}^{|\mathcal{Q}|}$ denote a set of intermediate questions corresponding to c obtained by claim decomposition (or question generation), and we use $\mathcal{D}' \subseteq \mathcal{D}$ to denote a set of documents retrieved by a query q that is related to the claim c (e.g., $q = c$ or $q \in \mathcal{Q}$).

Our task is to train a retrieval policy model π_θ with trainable parameters θ to output top- K ($K \ll |\mathcal{D}'|$) documents for LLM to determine the veracity of a claim as one of the classes defined in benchmark datasets (Wang, 2017; Yang et al., 2022b), such as true, mostly-true, half-true, barely-true, false, and pants-fire. We use $\pi_\theta(d|q)$ to denote the retrieval probability of a document d with respect to q .

4. Our Methodology

In this section, we will describe our FFRR model, starting from the basic document-level policy, followed by the question-level policy, and finally, a hybrid model that combines them.

4.1. Document-level Retrieval Policy

We begin with a preliminary version of the FFRR model in Figure 2(a), where the claim is used directly as a query for retrieval, resulting in a ranking list of documents in the corpus. The final goal of our policy model is to obtain top- K documents that are most useful to LLM’s final class prediction.

To optimize the retrieval model, our policy aims to select a set of retrieved documents for gathering intermediate feedback on them from LLM. Specifically, we use a dense retrieval model π_θ to obtain \mathcal{D}' containing retrieved documents with respect to the given c . We employ an ϵ -greedy strategy to strike a balance between exploitation, i.e., choosing from the highly ranked documents such as top- K in \mathcal{D}' , and exploration, i.e., selecting any documents in \mathcal{D}' . Thus, a document \tilde{d}_i is sampled following one

of the two document distributions according to

$$\tilde{d}_i \sim \begin{cases} \text{Uniform}(\mathcal{D}'), & \text{explore with prob } \epsilon \\ \pi_\theta(c, K), & \text{exploit with prob } 1 - \epsilon \end{cases}$$

where $\text{Uniform}(\mathcal{D}')$ is a uniform distribution over all retrieved documents, and $\pi_\theta(c, K)$ is the probability distribution of top- K documents in \mathcal{D}' retrieved by π_θ , and K and ϵ are hyper-parameters. At each time step, we select a document followed by reward acquisition from LLM as described below.

Training Rewards. During training, we consider the contribution of each selected document to making a *correct* veracity judgment based on it as the *document-level* intermediate reward. Specifically, for any sampled \tilde{d}_i , we compute a reward $r_{\tilde{d}_i} = \text{score}_{\text{LLM}}(y|c, \tilde{d}_i)$, where $\text{score}(\cdot)$ is the score function based on LLM’s output score of generating the *ground-truth label token* y given c and \tilde{d}_i as input, such as the probability or perplexity of the output token.

One might argue that the rewards generated by LLMs are likely biased by their preconceptions (Törnberg, 2023; Shaikh et al., 2023), potentially hampering retrieval performance. Why does it work provided that LLM might provide biased feedback due to such inherent limitation? Intuitively, the score function $r_{\tilde{d}_i}$ provides a natural facilitation mechanism to mitigate LLM bias on its veracity judgment, by telling the retriever what kind of evidence it returns will be helpful for LLM to make a right decision, which is governed by the *ground-truth* veracity label. Take the claim “Bruce Lee’s death is due to drinking too much water” with a ground-truth label “Unproven” as an example. If the LLM is strongly biased towards a different class, the reward score LLM gives to a retrieved document tends to be low since the posterior is “Unproven”. In order to get a higher reward on the correct class, the retriever will be encouraged to select stronger evidence documents to counterbalance LLM’s biased preperceptions, so that LLM can face more and more concrete evidence returned that indicates the claim is unproven.

Note that the sampling process can terminate automatically. For the termination condition, we utilize the score distribution $\text{score}_{\text{LLM}}(\mathbf{y}|c, \tilde{d}_i)$ over *all classes* \mathbf{y} that LLM outputs given c and a sampled \tilde{d}_i . The selection stops if $\text{argmax}_{\mathbf{y}' \in \mathbf{Y}} \sum_{i=1}^{\kappa} \text{score}(\mathbf{y}'|c, \tilde{d}_i) = y$, where y is the ground-truth label and κ is a free number of *actually sampled documents* upon termination.

Let us denote $k = \min(K, \kappa)$. Then we can propose the *final* prediction reward as $r_g = \text{score}_{\text{LLM}}(y|c, \tilde{d}_1, \tilde{d}_2, \dots, \tilde{d}_k)$, where $\tilde{d}_1, \tilde{d}_2, \dots, \tilde{d}_k$ are the k sampled documents being used by LLM to make the final prediction.

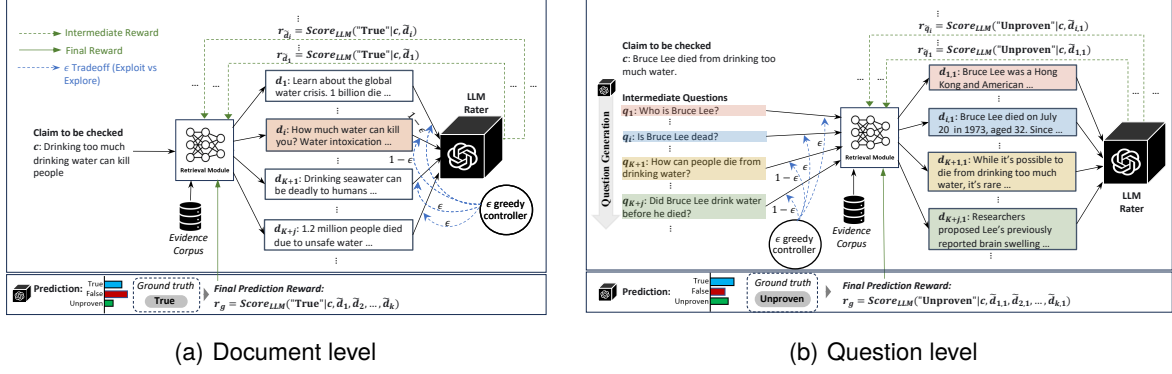


Figure 2: FFRR policy models with different fine-grained rewards from LLM. In Figure 2(a), d_i is shaded to indicate that it is the relevant document that potentially aids in verifying the truthfulness of the claim.

Inference. It is worth noting that the intermediate and final rewards are not needed once the retrieval model is trained. During inference, we just output the class token with the highest or the lowest score the LLM predicts depending on the specific metric used, e.g., probability or perplexity, by feeding top- K retrieved documents into it. Hence, our approach does not introduce additional overhead to inference. It indeed costs more training time, but not significantly high, which is primarily caused by network traffic. We believe that a minor training overhead with much better effectiveness is still cost-effective. Such latency tends to decrease with the advancement of hardware infrastructure.

Prompting LLM. We prompt LLM to rate the selected documents for providing intermediate and final feedback for policy rewarding. We use the following prompt template for getting the rewards: “The following evidence is given: [DOC(, DOC,...)]. Among [LABEL SET], the claim [CLAIM] is classified as:”. With the specific prompt text, LLM’s API will return a score distribution of generating each class label as the next token. We then obtain the score corresponding to ground-truth label for calculating the reward. The illustrative example can be found in Table 4 in Appendix A.

4.2. Question-level Retrieval Policy

It has been recognized that claim verification often requires the consideration of multiple focal points to facilitate multi-hop reasoning (Yao et al., 2023; Pan et al., 2023; Chen et al., 2023). To this end, we elicit LLM to decompose the claim c into intermediate questions $\{q_i\}$ by employing in-context learning (ICL) following (Chen et al., 2023). We instruct LLM to determine the number of questions automatically using demonstration examples. The detailed demonstrations are omitted here for space which can be found in Table 5 in Appendix A.

Next, we optimize document retrieval for each intermediate question with a question-level policy, as shown in Figure 2(b). We focus on improving top-1 retrieval for each question to gather question-level rewards, thereby promoting the discovery of alternative evidence that is not surfaced through a single query alone. Hence, we sample top-1 documents across different questions. The question-level intermediate reward is thus calculated as $r_{q_i} = \text{score}_{\text{LLM}}(y|c, \tilde{d}_{i,1})$, where $\tilde{d}_{i,1}$ is the top-1 document retrieved by \tilde{q}_i that is sampled out.

Moreover, we want to promote the questions that should be better addressed by top-1 documents than those that are straightforward. For example, compared to “Who is Bruce Lee”, questions like “Did Bruce Lee drink water before he died” and “How can people die from drinking water” are more complex and important, and it is more difficult to get the right top-1 documents for them. Thus, top-1 document sampling is intentionally biased towards the questions that receive lower r_{q_i} score to give them a higher chance of getting selected. That said, we sort the questions reversely by the score and use the ϵ -greedy method to select k questions based on top-1 documents of all the questions.

Finally, LLM utilizes the top-1 documents of k selected intermediate questions to make the final prediction. The final prediction reward is given as $r_g = \text{score}_{\text{LLM}}(y|c, \tilde{d}_{1,1}, \tilde{d}_{2,1}, \dots, \tilde{d}_{k,1})$, where the top-1 documents are from the k sampled questions with intermediate rewards. The model is trained to better rank top-1 documents for those questions that need to be better addressed. During inference, we return the top-1 documents of K questions that have the highest retrieval probability scores.

4.3. Retrieval Policy with Hybrid Rewards

We combine question- and document-level rewards to ensure that the model takes into account both breadth and depth of evidence retrieval. The hybrid policy is shown in Figure 3, where each optimiza-

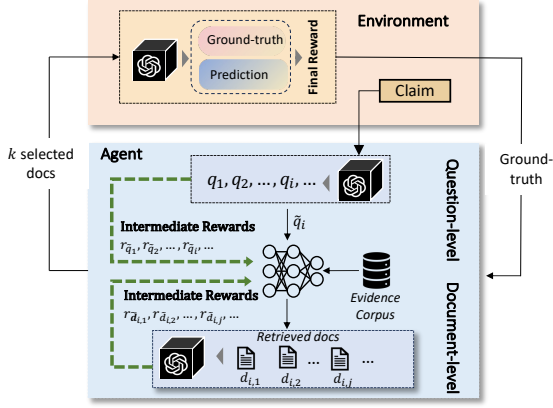


Figure 3: FFRR with two-level policies combined.

tion step at question level is followed by a sequence of steps at document level and the two-level optimization proceeds alternately.

We use two separate ϵ -greedy controllers for the policy models of different levels. The document-level intermediate rewards is similar as described in § 4.1, i.e., $r_{\tilde{d}_{i,j}} = \text{score}_{\text{LLM}}(y|c, \tilde{d}_{i,j})$ for any sampled intermediate question \tilde{q}_i . The question-level intermediate reward, and the final prediction reward are similar as described in § 4.2. In inference, using top- K documents from all questions for final prediction is limited by LLM input constraints. So, we only take top-1 document for each question, mirroring the question-level strategy.

5. Optimizing Retrieval Policy

Following prior works (Gao et al., 2022; Karpukhin et al., 2020), we use a dense retrieval model based on the dual-encoder architecture, where the encoder is a pre-trained RoBERTa (Liu et al., 2019) with parameter θ , to map a query q and a candidate document d into vectors $\mathbf{h}_\theta(q)$ and $\mathbf{h}_\theta(d)$, respectively. The goal is to optimize the encoder \mathbf{h}_θ so that the query-document similarity serves as an effective retrieval function.

Specifically, we index all the documents in \mathcal{D} offline using FAISS (Johnson et al., 2019) which is a highly efficient open-source library for similarity search and clustering of dense vectors. The retrieval probability is defined as:

$$\pi_\theta(d | q) = \frac{\exp(\mathbf{h}_\theta(q)^\top \cdot \mathbf{h}_\theta(d)/\tau)}{\sum_{d' \in \mathcal{D}} \exp(\mathbf{h}_\theta(q)^\top \cdot \mathbf{h}_\theta(d')/\tau)},$$

where τ is the temperature hyper-parameter and \mathcal{D}' is the set of retrieved documents. In practice, we use a widely-used open-source toolkit Tevatron (Gao et al., 2022) to build and optimize the dense retrieval model.

For optimizing retrieval policy π_θ , with the goal of assisting the LLM, we adopt a policy gradient RL

framework. Each RL step corresponds to sampling a document provided to LLM for rewarding. We formulate the retriever-LLM interaction as a Markov Decision Process (MDP) described by:

- S is an infinite set of states, which involves the claim, the documents selected in prior steps, and the current query to the retrieval model;
- A is a set of actions consisting of all the documents in \mathcal{D} ; An action corresponds to selecting a candidate document d returned by π_θ .
- $R(s, a)$ is the reward received after taking action $a \in A$ in state $s \in S$. This reward incorporates both the fine-grained intermediate and final prediction rewards, as detailed in prior sections.

Our goal is to maximize $\mathbb{E}_{d \sim \pi_\theta(d|q)}[R(q, d)]$, the expected reward under the policy, with respect to the parameters using REINFORCE algorithm (Williams, 1992). The gradient of the expectation is approximated as:

$$\begin{aligned} & \nabla \mathbb{E}_{d \sim \pi_\theta(d|q)} [R(q, d)] \\ &= \mathbb{E}_{d \sim \pi_\theta(d|q)} [\nabla_\theta \log(\pi_\theta(d | q)) R(q, d)] \\ &\approx \frac{1}{|\{(q, d)\}|} \sum_{(q, d)} \nabla_\theta \log(\pi_\theta(d | q)) R(q, d), \end{aligned}$$

which is the mean of reward-weighted gradients over the selected documents. We update the retrieval policy so that model parameters can be learned via the rewards received by the documents that have been sampled.

6. Experiments and Results

6.1. Experimental Setup

6.1.1. Datasets.

We conduct experiments on two public English news datasets referring to different fact-checking websites for ground-truth veracity labels: 1) **RAWFC** (Yang et al., 2022b) is based on claims and labels from Snopes in three classes (True/False/Half); 2) **LIAR-RAW** (Yang et al., 2022b) contains claims and labels from PolitiFact in six classes (True/Mostly-true/Half-true/Barely-true/False/Pants-fire), based on the LIAR-PLUS dataset (Alhindi et al., 2018).

We chose these two datasets because of their relatively lower ground-truth leakage than their counterparts like MultiFC (Augenstein et al., 2019) which was found 17.2% leaked evidence by our analysis using the leak detection method (Glockner et al., 2022). We used the raw documents in the two datasets collected by Yang et al. (2022b) for our corpus \mathcal{D} . These documents were retrieved from

the web by excluding a range of fact-checking websites, and all of them were dated before publication of the fact-check reports that correspond to the claims. Using the detection method proposed by Glockner et al. (2022), we further removed 4.7% documents from RAWFC and 2.4% documents from LIAR-RAW that are identified with leaks of ground truth. Table 1 displays the statistics of the two datasets.³

Following Yang et al. (2022b), we use macro-average precision (P), recall (R), and F_1 ($F_1 = \frac{2RP}{R+P}$) scores for evaluation. We utilize the supplied division of train-validate-test for both data sets. Specifically, they employed an 8/1/1 split ratio for the train/validation/test for the two datasets. Therefore, the corresponding number of samples in the RAWFC data for train/valid/test is 1,612/200/200, and the number of samples in the LIAR-RAW dataset is 10,065/1,274/1,251.

6.1.2. Implementation Details.

We use GPT-3.5 API `text-davinci-003`⁴ as a backbone LLM and generate outputs using greedy decoding. Following Wei et al. (2022), we tune the shot number hyper-parameter within the set $\{1, 2, 4, 6, 8\}$ on the validation dataset. We observe that, although the model’s performance experiences marginal gains with more than 2 shots, the increase in cost is substantial. Therefore, to balance cost and performance, we set the shot number of demonstration examples to 2.

We use the open-source fine-tuning toolkit Tevatron (Gao et al., 2022) to build the dense retrieval model. We pre-calculate the embeddings of documents \mathcal{D} and create a FAISS index (Johnson et al., 2019) to facilitate fast similarity search. We retrieve a set of 20 documents as \mathcal{D}' for each q from the FAISS index, and set the temperature $\tau = 1$. The selection of ϵ and K is based on the validation set. We examine a range of ϵ from 0 to 1 and find that an ϵ value of 0.1 gives a good balance between exploration and exploitation. Specifically, when $\epsilon > 0.1$, the model explores overly frequently, making it fo-

³We didn’t choose FEVER (Thorne et al., 2018) and its variants (Jiang et al., 2020; Schuster et al., 2021) because 1) their claims are crafted based on Wikipedia articles, which don’t encapsulate the complexities inherent in real-world news-related assertions that we aim to verify; 2) Wikipedia data are commonly known as an important part of the training data used to pre-train many LLMs, which is likely to leak ground truth of claims. We also excluded the datasets consisting of only human-compiled evidence documents referenced in fact-check reports, such as WatClaimCheck (Khan et al., 2022) and UKP-Snopes (Hanselowski et al., 2019), since the evidence documents need to be retrieved out of noise.

⁴<https://platform.openai.com/docs/models/gpt-3-5>.

	RAWFC	LIAR-RAW
Claim	2,012	12,590
# true	695	2,021
# mostly-true	-	2,439
# half-true*	671	2,592
# barely-true	-	2,057
# false	646	2,465
# pants-fire	-	1,012
Avg. doc per claim	20.0	12.0
Avg. sentence per doc	7.1	5.5

Table 1: Datasets statistics after removing the documents with detected ground truth leaks. # half-true* is also denoted as # half in RAWFC.

cus on random documents and hard to converge; when $\epsilon < 0.1$, the model seems to always exploit, failing to select any new evidence. We adjust K within the range of $[1, 8]$ and fix it as 3 based on the observation that model performance generally declines when $K > 3$. We train the retriever using Adam optimizer (Kingma and Ba, 2014) with a learning rate of $1e-5$, a batch size of 4, and a warm-up ratio of 0.1.

6.1.3. Compared Methods

Baselines. 1) *End-to-End optimized SOTA* (Yang et al., 2022b), which is not LLM-based. Both its retrieval module and classifier are tunable. 2) *Direct Prompting* (Brown et al., 2020), which directly prompts the LLM to determine the class label of the claim. 3) *Question-based Prompting*, which asks the LLM to decompose the claim into questions and then predict the label by answering the questions. 4) *ReAct* (Yao et al., 2023), which originally combines LLM’s reasoning ability and external Wikipedia API in an interleaved manner. We opt for Google Search API⁵ instead of Wikipedia API since news claims typically require extensive and diverse information (e.g., blog posts, forum threads, and authoritative news) for verification. 5) *Verify-and-Edit* (Zhao et al., 2023), which post edits CoT-style reasoning chains with external knowledge. Since it showed that using Google Search outperformed other retrieval methods on a fact-checking dataset, we adopt this setting as our baseline. 6) *REPLUG* (Shi et al., 2023), which adapts the dense retrieval model to a black-box LLM by minimizing the KLD between the distributions of the likelihoods of retrieval and language model. We tailor REPLUG to our document-level (d), question-level (q), and

⁵<https://serpapi.com>. The choice aligns with previous work in news claim verification that predominantly employs search engines for retrieving relevant information (Popat et al., 2018; Ma et al., 2019; Yang et al., 2022a).

Model	RAWFC			LIAR-RAW		
	$P(\%)$	$R(\%)$	$F_1(\%)$	$P(\%)$	$R(\%)$	$F_1(\%)$
End-to-End Optim. SOTA (Yang et al., 2022b)	52.8	51.0	51.9	29.0	28.9	28.9
Direct Prompting (Brown et al., 2020)	48.5	48.5	48.5	29.1	25.1	27.0
Question-based Prompting	48.3	47.9	48.1	27.1	24.5	25.7
ReAct (Yao et al., 2023)	51.2	48.5	49.8	33.2	29.0	31.0
Verify-and-Edit (Zhao et al., 2023)	53.1	52.8	52.9	32.7	28.1	30.2
REPLUG (d) (Shi et al., 2023)	51.8	56.9	54.2	30.6	27.9	29.2
REPLUG (q) (Shi et al., 2023)	52.4	56.0	54.1	30.0	28.7	29.3
REPLUG (d+q) (Shi et al., 2023)	52.5	56.9	54.6	30.3	28.9	29.6
FFRR-frozen (d)	53.6	57.9	55.7	30.7	29.5	30.1
FFRR-frozen (q)	51.9	54.2	53.0	29.9	27.3	28.5
FFRR-frozen (d+q)	51.9	54.2	53.0	29.9	27.3	28.5
FFRR (d)	55.5	56.2	55.8	33.0	29.8	31.3
FFRR (q)	54.2	56.0	55.1	32.8	28.4	30.4
FFRR (d+q)	56.5*	57.4	57.0*	34.5*	32.6*	33.5*

Table 2: Results of news claim verification on the two datasets where potential leaks are removed. Bold denotes the best performance. * means significantly better than the best baseline with $p < 0.01$. LLMs involved in different models are based on GPT-3.5 (text-davinci-003).

hybrid (d+q) retrieval settings. For REPLUG (q), we optimize the KLD composed of retrieved top-1 documents from K questions of each claim. For REPLUG (d+q), we optimize the KLD of top- K documents of each question for all the K questions of each claim, which are then combined with the KLD of top-1 documents of the K associated questions.

FFRR Variants. 1) *FFRR-frozen*: Our FFRR with the retrieval model frozen. 2) *FFRR*: Our FFRR with tunable retrieval model as detailed in § 4. Each of the two settings corresponds to the document-level (d), question-level (q), and hybrid (q+d) variants. Note that during inference, both FFRR (q) and FFRR (d+q) use top-1 document for each question, as described in § 4.3. Therefore, FFRR-frozen (d+q) and FFRR-frozen (q) are the same because retriever’s parameters and retrieved documents are fixed, leading to the same set of top-1 documents of the questions between the two settings.

6.2. Results of Verification

Table 2 shows our results with the following findings:

- **Simple zero-shot prompting methods already show competitive performance.** Direct Prompting lags behind by only 7.0% in F_1 score compared to the non-LM-based end-to-end optimized SOTA model. This suggests that LLM is promising in zero-shot news claim verification, thanks to LLM’s ability acquired from its pre-training data, which embeds a substantial amount of factual information. However, we find that Question-based

Prompting performs slightly worse. This might be attributed to the fact that Question-based Prompting demands more knowledge to answer intermediate questions, which LLM might not have, implying that more accurate information is needed from external sources.

- **Retrieval benefits LLM.** The performance improves when combining the LLM with methods for acquiring external knowledge from the web. ReAct and Verify-and-Edit integrate the multi-step reasoning capabilities of the LLM with an open-domain search engine to obtain relevant information, thus outperforming Direct Prompting and Question-based Prompting. Interestingly, FFRR-frozen utilizing the frozen dense retrieval model clearly boosts the performance even higher on RAWFC and is just marginally lower than ReAct and Verify-and-Edit on LIAR-RAW. This suggests the strong potential of combining small retrieval models with black-box LLMs for the task.
- **Tuning the dense retrieval model with FFRR is superior.** Our FFRR consistently provides improvements. Compared to FFRR-frozen, the FFRR (d), (q), and (d+q) demonstrate an increase of 1.7%, 5.3%, and 12.5% on the F_1 score, respectively. REPLUG cannot compete with it because REPLUG is unable to flexibly explore and exploit all the retrieved documents in \mathcal{D}' but sticks to its top- K results. Furthermore, the claim decomposition of FFRR allows for searching evidence documents from various question perspectives. Despite this expanding the coverage, the questions as queries lead to more question-




Figure 4: Effect of number of documents K .

oriented evidence, which may not necessarily be helpful for the final prediction. This also account for the reason why FFRR (q) is slightly worse than FFRR (d), since the latter can flexibly explore the retrieved documents to the greater depth of the ranking list. When combining both document-level and question-level rewards, FFRR (d+q) enjoys the advantages of both, thereby outperforming all these baselines.

6.3. Result Analysis

Effect of the number of documents K . Firstly, we examine the effect of the number of documents K by adjusting its value in the range from 1 to 5 for FFRR (d), FFRR (q), and FFRR (d+q) based on the test sets. As shown in Figure 4, the increase of K improves the performance consistently when $K \leq 3$. Our analysis indicates that while a too small K (e.g., $K = 1$) might still be able to find some documents needed to cover necessary evidence, a large K (e.g., $K = 5$) tends to include more irrelevant or noisy information generally leading to no further improvement. This suggests that the best option for guiding the LLM-based verification would be just using the top 3-4 documents.

Effect of final reward r_g . To analyze the impact of the final prediction reward of FFRR, we conduct an ablation analysis as shown in Figure 5. For all three settings, i.e., d, q, and d+q, the performance of the model decreases on both datasets after removing the final prediction reward. This indicates that the final prediction reward is necessary and helpful, because it can measure the usefulness of the set of selected documents put altogether to influence LLM’s judgement. However, if we only use the final reward for tuning, the performance is unsatisfactory. This confirms our hypothesis that optimizing the retrieval model with the sparse final feedback alone is insufficient, and thus our more fine-grained feedback scheme is necessary.

Error analysis of the retrieval process. To delve deeper into the FFRR model’s performance in retrieval, we conduct a more detailed analysis.

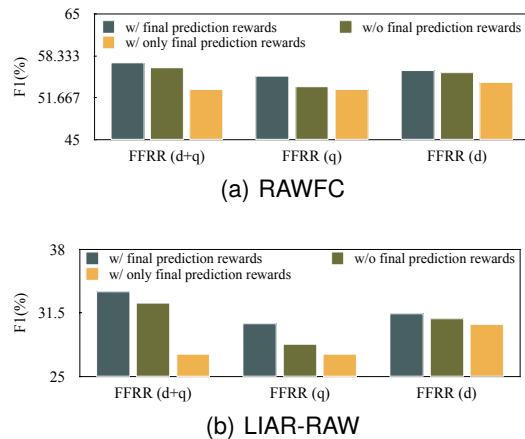


Figure 5: Effect of final reward r_g .

Error Types	FFRR(d+q)
Irrelevant Questions	16%
Insufficient Coverage	34%
Redundant Questions	14%
Document Mismatch	48%

Table 3: Distribution of errors based on 50 examples from RAWFC, where FFRR(d+q) gives incorrect verification results.

Specifically, we randomly sample 50 instances that are predicted incorrectly by FFRR (i.e., failure cases) and manually identify four types of errors: (1) *Irrelevant questions*: this occurs when the generated intermediate questions do not pertain to or can not validate the claim; (2) *Insufficient Coverage*: the generated intermediate questions can not adequately cover the required information; (3) *Redundant questions*: the generated questions are repetitive or unnecessary for the verification; and (4) *Document Mismatch*: the documents retrieved can not answer the generated intermediate questions. In other words, there are clear mismatches or information gaps between the retrieved documents and the questions.

As shown in Table 3, we manually label the error types for these 50 instances. We clearly find that the most prevalent challenge lies in the document mismatch issue, which comprises a notable 48% of the failure cases. This indicates that accurately retrieving evidential documents that can answer the raised questions is critical but desperately needs further improvement for providing necessary relevant context to the LLM. Following closely, the insufficient coverage issue, contributing to 34% of the failures, suggests that sometimes specific nuances or details are not sufficiently obtained from the retrieval process. This can lead to situations where not all facets of a claim are adequately corroborated, resulting in incorrect predictions. On the

brighter side, irrelevant and redundant questions appear to be less of a main concern. This gives us confidence that the LLM is generally adept at formulating intermediate questions that pertain directly to the given claim without being repetitive.

7. Conclusion and Future Work

In this paper, we study the reinforcement optimization of the retrieval model to retrieve evidence for promoting the capability of black-box LLM on fact verification of real-world news claims. We propose a method called FFRR utilizing Fine-grained Feedback for Reinforcement Retrieval, which adopts a two-level strategy to gather different rewards from the LLM based on retrieved documents to optimize the retrieval policy. Evaluated on two public datasets, our LLM-based verification model leveraging FFRR significantly outperforms state-of-the-art LLM-enabled and non-LLM baselines. In the future, we will develop a conversational news claim verification model based on LLMs that can receive fine-grained human feedback used to better align retrieval models with professional fact-checkers. Furthermore, we can also employ FFRR to optimize tunable retrieval models working in conjunction with LLMs on other NLP tasks.

8. Limitations and Ethics Statement

While our FFRR has demonstrated promising results, several limitations should be acknowledged.

Firstly, although the fine-grained reward is based on the final ground-truth label, our method greatly relies on the LLM's capabilities to assess the usefulness of selected documents to the prediction of the ground truth, as well as to make the final prediction. The performance of the LLM has a direct impact on our method's outcomes. The LLM's inherent limitations, such as gaps in its understanding of specific contexts or language nuances, can influence the accuracy of our retrieval model and subsequent verification results. Similarly, any intrinsic bias in the LLM stemming from its training data can potentially skew the reward towards certain perspectives, thereby affecting the verification outcomes. We need further study on the influence of these issues on our task.

Secondly, the efficiency of training the reinforcement retrieval process could be handicapped by the frequent interactions for receiving feedback from the LLM. On the one hand, we may need a systematic approach to optimize the feedback loop using more advanced exploration strategies to reduce the number of interactions needed. On the other hand, accelerating LLM inference can be achieved by integrating techniques such as non-autoregressive transformer (Gu et al., 2017; Du et al., 2021, 2022)

and speculative decoding (Leviathan et al., 2023; Du et al., 2024).

Thirdly, the domain of this study is primarily focused on the verification of news claims only, which is just a portion of broader problem of misinformation. Our concentration on news claims can hopefully path for future studies, which could be extended to check other forms of misinformation, such as fact verification (Zhang et al., 2023; Zeng and Gao, 2023) and infodemic surveillance (Zhang and Gao, 2024).

The potential ethical concern associated with this research might be due to the generally low accuracy of news claim verification. The reliability of the system would be a primary concern when applied in the real world. Mislabeling false information or vice versa could have potential societal harm. Therefore, we advocate for the careful application of our model in real world. While our research contributes to the fight against misinformation, we believe the battle should also be supported by responsible journalism and informed public discourse.

Acknowledgement

We thank the anonymous reviewers for their helpful comments during the review of this paper.

Tariq Alhindi, Savvas Petridis, and Smaranda Muresan. 2018. Where is your evidence: Improving fact-checking by justification modeling. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, Brussels, Belgium. Association for Computational Linguistics.

Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*.

Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims. In *EMNLP*. Association for Computational Linguistics.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022a. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.

- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022b. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *NeurIPS*, volume 33, pages 1877–1901.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Jifan Chen, Grace Kim, Aniruddh Sriram, Greg Durrett, and Eunsol Choi. 2023. Complex claim verification with evidence retrieved in the wild. *arXiv preprint arXiv:2305.11859*.
- Jifan Chen, Aniruddh Sriram, Eunsol Choi, and Greg Durrett. 2022. Generating literal and implied subquestions to fact-check complex claims. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3495–3516, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Cunxiao Du, Jing Jiang, Xu Yuanchen, Jiawei Wu, Sicheng Yu, Yongqi Li, Shenggui Li, Kai Xu, Liqiang Nie, Zhaopeng Tu, et al. 2024. Glide with a cape: A low-hassle method to accelerate speculative decoding. *arXiv preprint arXiv:2402.02082*.
- Cunxiao Du, Zhaopeng Tu, and Jing Jiang. 2021. Order-agnostic cross entropy for non-autoregressive machine translation. In *International Conference on Machine Learning*, pages 2849–2859. PMLR.
- Cunxiao Du, Zhaopeng Tu, Longyue Wang, and Jing Jiang. 2022. ngram-OAXE: Phrase-based order-agnostic cross entropy for non-autoregressive machine translation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5035–5045, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Angela Fan, Aleksandra Piktus, Fabio Petroni, Guillaume Wenzek, Marzieh Saeidi, Andreas Vlachos, Antoine Bordes, and Sebastian Riedel. 2020. Generating fact checking briefs. In *EMNLP*, pages 7147–7161.
- Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2022. Tevatron: An efficient and flexible toolkit for dense retrieval. *arXiv preprint arXiv:2203.05765*.
- Max Glockner, Yufang Hou, and Iryna Gurevych. 2022. Missing counter-evidence renders nlp fact-checking unrealistic for misinformation. In *EMNLP*, pages 5916–5936.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2017. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pappas, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *ICLR*, pages 3929–3938. PMLR.
- Andreas Hanselowski, Christian Stab, Claudia Schulz, Zile Li, and Iryna Gurevych. 2019. A richly annotated corpus for different tasks in automated fact-checking. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 493–503, Hong Kong, China. Association for Computational Linguistics.
- Yichen Jiang, Shikha Bordia, Zheng Zhong, Charles Dognin, Maneesh Singh, and Mohit Bansal. 2020. HoVer: A dataset for many-hop fact extraction and claim verification. In *Findings of EMNLP*, pages 3441–3460, Online. Association for Computational Linguistics.
- Zhengbao Jiang, Frank F Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Active retrieval augmented generation. *arXiv preprint arXiv:2305.06983*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547.

- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *EMNLP*, pages 6769–6781.
- Kashif Khan, Ruizhe Wang, and Pascal Poupart. 2022. WatClaimCheck: A new dataset for claim entailment and inference. In *ACL*, pages 1293–1304, Dublin, Ireland. ACL.
- Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts, and Matei Zaharia. 2022. Demonstrate-search-predict: Composing retrieval and language models for knowledge-intensive nlp. *arXiv preprint arXiv:2212.14024*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, pages 19274–19286. PMLR.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *NeurIPS*, 33:9459–9474.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. 2023. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. In *ICLR*.
- Jing Ma, Wei Gao, Shafiq Joty, and Kam-Fai Wong. 2019. Sentence-level evidence embedding for claim verification with hierarchical attention networks. In *ACL*.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- OpenAI. 2022. [Introducing chatgpt](#).
- OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Nedjma Ousidhoum, Zhangdie Yuan, and Andreas Vlachos. 2022. Varifocal question generation for fact-checking. In *EMNLP*, pages 2532–2544.
- 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. *NeurIPS*, 35:27730–27744.
- Liangming Pan, Wenhui Chen, Wenhui Xiong, Min-Yen Kan, and William Yang Wang. 2021. Zero-shot fact verification by claim generation. In *ACL (Short)*, pages 476–483.
- Liangming Pan, Xiaobao Wu, Xinyuan Lu, Anh Tuan Luu, William Yang Wang, Min-Yen Kan, and Preslav Nakov. 2023. [Fact-checking complex claims with program-guided reasoning](#).
- Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, et al. 2023. Check your facts and try again: Improving large language models with external knowledge and automated feedback. *arXiv preprint arXiv:2302.12813*.
- Kashyap Popat, Subhabrata Mukherjee, Andrew Yates, and Gerhard Weikum. 2018. Declare: Debunking fake news and false claims using evidence-aware deep learning. In *EMNLP*, pages 22–32.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. 2023. Measuring and narrowing the compositionality gap in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, Singapore. Association for Computational Linguistics.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. Get your vitamin C! robust fact verification with contrastive evidence. In *NAACL*, Online. ACL.
- Omar Shaikh, Hongxin Zhang, William Held, Michael Bernstein, and Diyi Yang. 2023. On second thought, let’s not think step by step! bias and toxicity in zero-shot reasoning. In *ACL*, pages 4454–4470, Toronto, Canada. ACL.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2023. Replug: Retrieval-augmented black-box language models. *arXiv preprint arXiv:2301.12652*.

- Tianxiang Sun, Yunfan Shao, Hong Qian, Xuanjing Huang, and Xipeng Qiu. 2022. Black-box tuning for language-model-as-a-service. In *ICML*, pages 20841–20855.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. In *NAACL*, pages 809–819.
- Petter Törnberg. 2023. Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning. *arXiv preprint arXiv:2304.06588*.
- William Yang Wang. 2017. “liar, liar pants on fire”: A new benchmark dataset for fake news detection. In *ACL (short)*, pages 422–426.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. In *ICLR*.
- 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. *NeurIPS*, 35:24824–24837.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256.
- Zequi Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A Smith, Mari Ostendorf, and Hannaneh Hajishirzi. 2023. Fine-grained human feedback gives better rewards for language model training. *arXiv preprint arXiv:2306.01693*.
- Jing Xu, Megan Ung, Mojtaba Komeili, Kushal Arora, Y-Lan Boureau, and Jason Weston. 2022. Learning new skills after deployment: Improving open-domain internet-driven dialogue with human feedback. *arXiv preprint arXiv:2208.03270*.
- Jing Yang, Didier Vega-Oliveros, Taís Seibt, and Anderson Rocha. 2022a. Explainable fact-checking through question answering. In *ICASSP*, pages 8952–8956.
- Zhiwei Yang, Jing Ma, Hechang Chen, Hongzhan Lin, Ziyang Luo, and Yi Chang. 2022b. A coarse-to-fine cascaded evidence-distillation neural network for explainable fake news detection. In *COLING*, pages 2608–2621.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. ReAct: Synergizing reasoning and acting in language models. In *ICLR*.
- Xi Ye and Greg Durrett. 2022. The unreliability of explanations in few-shot prompting for textual reasoning. *NeurIPS*.
- Fengzhu Zeng and Wei Gao. 2023. Prompt to be consistent is better than self-consistent? few-shot and zero-shot fact verification with pre-trained language models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4555–4569, Toronto, Canada. Association for Computational Linguistics.
- Fengzhu Zeng and Wei Gao. 2024. [Justilm: Few-shot justification generation for explainable fact-checking of real-world claims](#).
- Fanrui Zhang, Jiawei Liu, Qiang Zhang, Esther Sun, Jingyi Xie, and Zheng-Jun Zha. 2023. Ecenet: Explainable and context-enhanced network for multi-modal fact verification. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 1231–1240.
- Xuan Zhang and Wei Gao. 2023. Towards LLM-based fact verification on news claims with a hierarchical step-by-step prompting method. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 996–1011, Nusa Dua, Bali. Association for Computational Linguistics.
- Xuan Zhang and Wei Gao. 2024. Predicting viral rumors and vulnerable users with graph-based neural multi-task learning for infodemic surveillance. *Information Processing and Management*, 61(1):1.
- Ruo Chen Zhao, Xingxuan Li, Shafiq Joty, Chengwei Qin, and Lidong Bing. 2023. Verify-and-edit: A knowledge-enhanced chain-of-thought framework. In *ACL*, pages 5823–5840, Toronto, Canada.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. In *ICLR*.

A. Examples of Instruction Prompts

Table 4 and 5 display examples that serve as prompts for the LLM to make the final prediction and to generate intermediate questions, respectively.

<p>The following evidence is given: [DOC(, DOC, ...)]. Among [LABEL SET], the claim "Emerson Moser, who was Crayola's top crayon molder for almost 40 years, was colorblind." can be classified as <u>true</u>.</p> <p>The following evidence is given: [DOC(, DOC, ...)]. Among [LABEL SET], the claim "Bernie Sanders said 85 million Americans have no health insurance." can be classified as <u>half-true</u>.</p> <p>The following evidence is given: [DOC(, DOC, ...)]. Among [LABEL SET], the claim "Cheri Beasley "backs tax hikes — even on families making under \$75,000." can be classified as <u>False</u>.</p>
<p>The following evidence is given: [DOC(, DOC, ...)]. Among [LABEL SET], the claim "JAG charges Nancy Pelosi with treason and seditious conspiracy." can be classified as <u>False</u>.</p>

Table 4: An illustrative example of prompting LLM to make the final prediction. Examples used for demonstration are shaded in grey, with the claim undergoing verification highlighted in yellow and the LLM's prediction marked in green.

<p>Claim: <i>Emerson Moser, who was Crayola's top crayon molder for almost 40 years, was colorblind.</i></p> <p>To verify the claim, a fact-checker will go through a step-by-step process to ask and answer a series of questions relevant to its factuality. Here are the specific questions raised:</p> <p>Question: Is there any official record or documentation indicating that Emerson Moser worked as a crayon molder at Crayola?</p> <p><i>Question:</i> Are there any official records or documentation confirming Emerson Moser's length of employment at Crayola?</p> <p>Question: Are there credible sources or publications that mention Emerson Moser as Crayola's top crayon molder?</p> <p>Question: Are there any credible sources or records indicating that Emerson Moser was colorblind?</p> <p>Question: Was Emerson Moser's colorblindness only confusing for certain colors?</p>
<p>Claim: <i>"Bernie Sanders said 85 million Americans have no health insurance."</i></p> <p>To verify the claim, a fact-checker will go through a step-by-step process to ask and answer a series of questions relevant to its factuality. Here are the specific questions raised:</p> <p>Question: How many Americans did Bernie Sanders claim had no health insurance?</p> <p>Question: How did Bernie Sanders define "no health insurance"?</p> <p>Question: How many Americans were uninsured or under-insured according to the Commonwealth Fund survey?</p>
<p>Claim: <i>"Cheri Beasley "backs tax hikes — even on families making under \$75,000."</i></p> <p>To verify the claim, a fact-checker will go through a step-by-step process to ask and answer a series of questions relevant to its factuality. Here are the specific questions raised:</p> <p>Question: Does Cheri Beasley supports tax increases?</p> <p>Question: Does the ad accurately link Beasley's position on student loan debt forgiveness with her stance on tax hikes for families making under \$75,000 per year?</p>
<p>Claim: <i>"JAG charges Nancy Pelosi with treason and seditious conspiracy."</i></p> <p>To verify the claim, a fact-checker will go through a step-by-step process to ask and answer a series of questions relevant to its factuality. Here are the specific questions raised:</p> <p>Question: Is it true that JAG has made a claim or accusation against Nancy Pelosi?</p> <p>Question: Is it true that the specific charges or allegations made against Nancy Pelosi are treason and seditious conspiracy?</p>

Table 5: An illustrative example of prompting LLM to generate intermediate questions. Examples used for demonstration are shaded in grey, with the claim undergoing verification highlighted in yellow and the LLM's generation marked in green.