

FACT-AUDIT: An Adaptive Multi-Agent Framework for Dynamic Fact-Checking Evaluation of Large Language Models

Hongzhan Lin^{1,2}, Yang Deng³, Yuxuan Gu^{4,2}, Wenxuan Zhang⁵,
Jing Ma^{1,*}, See-Kiong Ng², Tat-Seng Chua²

¹Hong Kong Baptist University, ²National University of Singapore,
³Singapore Management University, ⁴Harbin Institute of Technology,
⁵Singapore University of Technology and Design
{cshzlin,majing}@comp.hkbu.edu.hk, ydeng@smu.edu.sg

Abstract

Large Language Models (LLMs) have significantly advanced the fact-checking studies. However, existing automated fact-checking evaluation methods rely on static datasets and classification metrics, which fail to automatically evaluate the justification production and uncover the nuanced limitations of LLMs in fact-checking. In this work, we introduce FACT-AUDIT, an agent-driven framework that adaptively and dynamically assesses LLMs' fact-checking capabilities. Leveraging importance sampling principles and multi-agent collaboration, FACT-AUDIT generates adaptive and scalable datasets, performs iterative model-centric evaluations, and updates assessments based on model-specific responses. By incorporating justification production alongside verdict prediction, this framework provides a comprehensive and evolving audit of LLMs' factual reasoning capabilities, to investigate their trustworthiness. Extensive experiments demonstrate that FACT-AUDIT effectively differentiates among state-of-the-art LLMs, providing valuable insights into model strengths and limitations in model-centric fact-checking analysis.

1 Introduction

Large language models (LLMs) have transformed natural language processing (NLP), significantly enhancing performance in various tasks (Touvron et al., 2023a; OpenAI, 2023). Particularly, previous literature (Petroni et al., 2019; Jiang et al., 2020b) has shown that LLMs store factual knowledge and function as knowledge bases, which aids in fact-checking (Pan et al., 2023). However, LLMs still struggle with identifying factual errors and are prone to reasoning mistakes (Lin et al., 2022b; Bubeck et al., 2023). Errors in stored knowledge or deficiencies in fact reasoning capabilities may limit their credibility in fact-checking, impacting their utility (Elazar et al., 2021; Cao et al., 2021).

* Corresponding author.

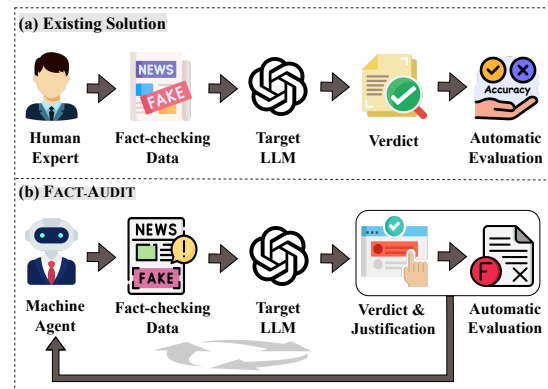


Figure 1: The pipelines of the existing solution and the proposed FACT-AUDIT in fact-checking evaluation.

Therefore, systematically revealing the boundaries of the fact-checking capacities in LLMs is essential to enhancing the trustworthiness of LLMs.

Auditing the fact-checking capacities of LLMs is challenging due to the complex and open-ended nature of real-world applications like complex claims, fake news, or rumors on social media. As illustrated in Figure 1(a), existing studies typically design intricate manual fact examinations to annotate check-worthy natural language scenarios (Yang et al., 2024b; Hu et al., 2024a; Wang et al., 2024). There are several limitations to such fact-checking evaluation methods: 1) The labor-intensive process restricts the scope of test scenarios, making it costly to scale. 2) All the static datasets for the fact-checking evaluation (Chen and Shu, 2024; Hu et al., 2024b) face risks like test data leakage and leaderboard swamping, failing to timely and adaptively reveal potential limitations of LLMs for understanding factuality. 3) Their problem settings often oversimplify the evaluation to a classification paradigm that focuses on accuracy, which may not adequately capture the other critical capabilities of fact-checking models, like justification production (Eldifrawi et al., 2024) for verdict prediction in the fact-checking process (Guo et al., 2022).

In response to these challenges, we introduce

a novel evaluation framework for systematically auditing the fact-checking capabilities of LLMs, called **FACT-AUDIT**. As illustrated in Figure 1(b), the core design philosophy of FACT-AUDIT centers on automating adaptive LLM auditing with two key features: (1) dynamically updated fact-checking test data and (2) in-depth evaluation of model-generated justifications. Theoretically, the creation of fact-checking test data can be framed as a Monte Carlo sampling process (Metropolis et al., 1953), where test cases are sampled from an oracle knowledge space. However, the inherent inefficiency of traditional Monte Carlo sampling limits its ability to generate comprehensive, scalable datasets for robustly assessing LLM fact-checking capabilities. To this end, we propose an importance sampling-based approach (Kahn and Marshall, 1953), which adaptively targets diverse weaknesses in LLM fact-checking by leveraging insights from the model-generated justifications.

In this work, FACT-AUDIT employs a multi-agent framework that leverages the exceptional capabilities of LLM-powered autonomous agents in experiential learning and complex reasoning (Park et al., 2023; Shen et al., 2023). Specifically, 1) FACT-AUDIT first establishes a detailed taxonomy to categorize different fact-checking scenarios, and then samples the prototype test data, with its quality validated through a tool-using module. 2) For each fact-checking test scenario, FACT-AUDIT evaluates the target LLM on both fact verification and justification production, by using the prototype test data as well as an iterative probing process to generate more diverse and unseen test cases under the scenario via importance sampling. 3) Upon completing evaluations for all the test scenarios in the taxonomy, FACT-AUDIT updates the test scenarios based on the model’s performance, enabling the auditing process to adaptively identify new and critical deficiencies in the LLM’s fact-checking capabilities. This updated taxonomy is then used to repeat the auditing process for creating a dynamic and model-centric evaluation loop in fact-checking.

Our contributions are summarized as follows:

- We introduce a novel and adaptive fact-checking evaluation framework, FACT-AUDIT, that utilizes multi-agent collaboration to dynamically unveil the limitations of the LLM’s fact-checking capabilities under diverse test scenarios.¹

¹The source code is released via <https://github.com/DanielLin97/FACT-AUDIT>.

- FACT-AUDIT addresses the restrictions of static fact-checking datasets by dynamically updating test scenarios and iteratively probing challenging cases. This framework ensures adaptability to real-world fact-checking complexity while maintaining diversity and scalability in LLM auditing.
- FACT-AUDIT goes beyond traditional accuracy-based automatic evaluations by integrating justification production with verdict prediction.
- We conduct extensive experiments on 13 state-of-the-art LLMs and detailed analyses of fact-checking performance, to provide insight into model strengths and areas for improvement.

2 Preliminary

In the context of assessing the fact-checking capacity in LLMs, we denote key components as follows:

Definition 2.1: Paradigm Definition

1. Oracle Knowledge Distribution: $p(x)$
The true distribution of factual knowledge.
2. Fact-Checking Limits of LLM α : $\mathcal{F}_\alpha(x)$
The function characterizing the LLM’s understanding limits of a given fact-checking test case x .

We formulate the automated auditing of the LLM α ’s fact-checking capabilities as a Monte Carlo sampling process (Metropolis and Ulam, 1949), *i.e.*, continuously sampling test cases x by humans from the oracle distribution $p(x)$ in the real world, and calculate its corresponding limits $\mathcal{F}_\alpha(x)$:

$$\mathbb{E}_{p(x)}[\mathcal{F}_\alpha(x)] = \int p(x)\mathcal{F}_\alpha(x)dx. \quad (1)$$

However, beyond the well-known inefficiency of Monte Carlo sampling with a convergence rate of $\mathcal{O}(1/\sqrt{N})$, the long-tail knowledge distribution further exacerbates the inefficiency of sampling from $p(x)$ for constructing fact-checking datasets.

Inspired by the classic Importance Sampling (Kahn and Marshall, 1953) method, which leverages a proposal distribution $q(x)$ to improve efficiency by allocating more densities to the regions where $\mathcal{F}_\alpha(x)$ is more likely to have higher values, we aim to adopt this concept for adaptively and efficiently sampling test data according to the fact-checking limits of the LLM. In the strategy of Importance Sampling, the process is adjusted as:

$$\begin{aligned} \mathbb{E}_{p(x)}[\mathcal{F}_\alpha(x)] &= \int q(x)\mathcal{F}_\alpha(x)\frac{p(x)}{q(x)}dx \\ &= \mathbb{E}_{q(x)}\left[\mathcal{F}_\alpha(x)\frac{p(x)}{q(x)}\right], \end{aligned} \quad (2)$$

where the importance weight $\frac{p(x)}{q(x)}$ compensates

for the discrepancy between the proposal distribution $q(x)$ and oracle distribution $p(x)$ to ensure an unbiased estimate. Besides, the efficiency of importance sampling critically depends on choosing $q(x) \propto p(x)\mathcal{F}_\alpha(x)$ as closely as possible. Therefore, our method is to find a well-designed $q(x)$ that minimizes the variance of the objective $p(x)\mathcal{F}_\alpha(x)$, thereby improving the reliability of the estimation.

3 FACT-AUDIT

3.1 Problem Definition

Given a source claim (SC), fact checking aims to predict the factuality and provide convincing justifications, to evaluate the claim as *Factual*, *Non-Factual*, or *Not Enough Information*, based on a knowledge source as auxiliary information (AI). Our objective is to develop a multi-agent evaluation framework, for modeling a new distribution $q(x)$ that tends to reveal fact-checking limitations, thus replacing the inefficient evaluation methods reliant on the sampling distribution of $p(x)$. Considering the difficulty of directly obtaining the optimal $q(x)$, we design an adaptive framework to iteratively converge to the desired distribution $q(x)$, automatically and dynamically evaluating the target LLM α 's capabilities across diverse fact-checking domains (e.g., complex claims, fake news, and rumors).

Following the definition in §2, our framework is formulated with three main stages:

Definition 3.1: Framework Formulation

1. Prototype Emulation: $x \sim q(x|\theta_i)$
Generate fact-checking test data for LLM auditing.
2. Fact Verification: $\mathbb{E}_{q_i} \left[\mathcal{F}_\alpha(x) \frac{p(x)}{q(x|\theta_i)} \right]$
Test the target LLM with the specific fact-checking questions x to verify fact and produce justification.
3. Adaptive Updating: $\pi(\Theta_i|\Theta_{i-1}, \mathcal{M})$
Explore more diverse and challenging test data.

As presented in Algorithm 1, FACT-AUDIT maintains a taxonomy of fact-checking scenarios Θ during iterations, where Θ_0 is initialized to be the foundational test scenarios that $\mathbb{E}_{\theta_0 \sim P(\Theta_0)} [q(x|\theta_0)] = p(x)$. During the loop, Θ_i will be updated to focus on the specific fact-checking scenarios that the target LLM α is likely to underperform. To audit the weaknesses of LLMs in fact checking, our process mainly involves three stages: 1) Generate the dynamic and check-worthy source claim dataset \mathbb{X} (§3.2); 2) Query the target LLM for veracity prediction and justification production (§3.3);

Algorithm 1 FACT-AUDIT Algorithm

- 1: **Initialize** fact-checking test scenarios Θ_0 and a memory pool $\mathcal{M} = \phi$
- 2: **for** $i := 0$ to n **do**
- 3: $\mathbb{X} := \phi$
- 4: **Stage 1: Prototype Emulation**
- 5: **while** $|\mathbb{X}| < k$ **do**
- 6: *Appraiser:* $\theta_i \sim P(\Theta_i)$
- 7: *Inquirer:* $x \sim q(x|\theta_i)$
- 8: **if** x satisfies *Quality Inspector* **then**
- 9: $\mathbb{X} := \mathbb{X} \cup \{x\}$
- 10: **end if**
- 11: **end while**
- 12: **Stage 2: Fact Verification with Justification**
- 13: $\mathcal{M} := \mathcal{F}_\alpha(\mathbb{X}) \frac{p(\mathbb{X})}{q(\mathbb{X}|\Theta_i)}$
- 14: **for** $j := 0$ to m **do**
- 15: *Prober:* $x \sim \rho(\mathcal{M})$
- 16: $\mathcal{M} := \mathcal{M} \cup \left\{ \mathcal{F}_\alpha(x) \frac{p(x)}{q(x|\theta_i)} \right\}$
- 17: **end for**
- 18: **Stage 3: Adaptive Updating**
- 19: $\Theta_{i+1} \sim \pi(\Theta_{i+1}|\Theta_i, \mathcal{M})$
- 20: **end for**
- 21: **Return** \mathcal{M}

3) Scrutinize the limitations of the target LLM in fact-checking adaptively based on specific model-generated justifications (§3.4). An overview of our FACT-AUDIT framework is shown in Figure 2.

3.2 Prototype Emulation

The Stage 1 of Algorithm 1 is Prototype Emulation, which involves generating prototype test data for assessing the LLM's fact-checking capabilities. This stage is accomplished by three agents: 1) an *Appraiser* agent to develop the taxonomy of fact-checking scenarios for evaluation, 2) an *Inquirer* agent to generate prototype test data according to the taxonomy, and 3) a *Quality Inspector* agent to ensure the quality of the prototype test data.

Appraiser Given the fact-checking objects, the Appraiser agent first generates the detailed taxonomy $\theta_i \sim P(\Theta_i)$, which includes k fact-checking scenarios $\{\theta_i\}_k$ towards the specific fact-checking objects. As shown in Figure 3, the Appraiser initializes the taxonomy Θ_0 from the three classic fact-checking objects: complex claims (Jiang et al., 2020a; Aly et al., 2021), fake news (Hu et al., 2024a; Wang et al., 2024), and social rumors (Ma et al., 2015, 2017), drawing inspiration from previous literature (Hu et al., 2024b; Waldrop, 2017; Allport, 1947). Note that in the subsequent phase, Appraiser would excavate new fact-checking test scenarios to update the initial taxonomy by examining the intermediate evaluation feedback.

Inquirer According to each fact-checking sce-

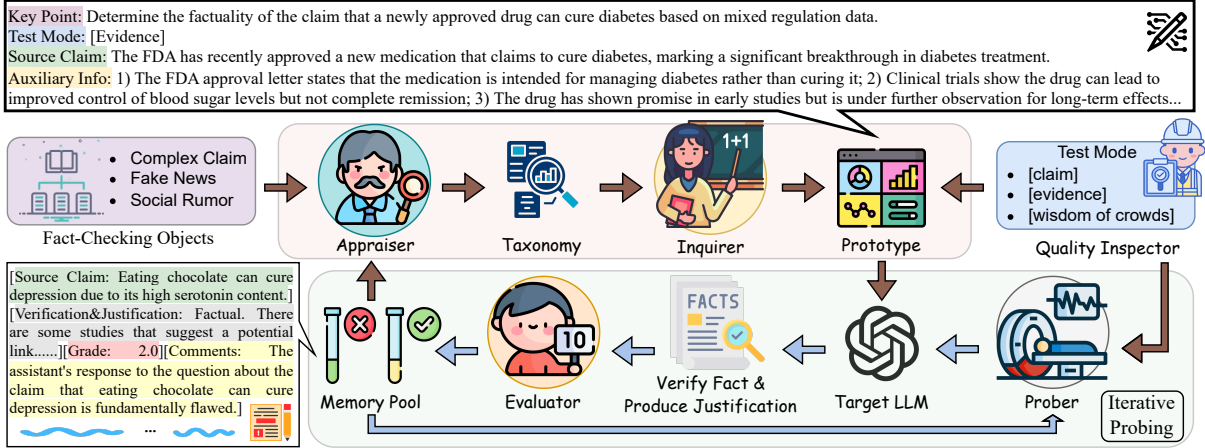


Figure 2: An overview of FACT-AUDIT, to adaptively unveil the limitations of fact-checking in LLMs.

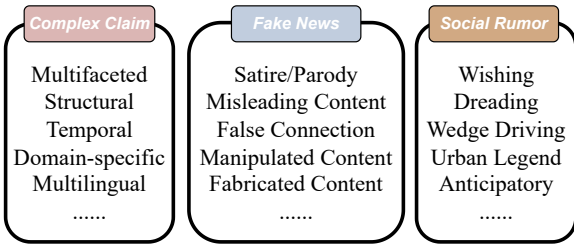


Figure 3: The taxonomy of fact-checking scenarios.

nario θ_i , the Inquirer agent generates the prototype test data: $x \sim q_i(x) = q(x|\theta_i)$, where $q_i(x)$ is the proposal distribution of importance sampling. As depicted in Figure 2, a prototype data sample x encompasses the following four components:

- **Key Point (KP)**: the specific task instruction for the test case.
- **Source Claim (SC)**: the claim to be verified.
- **Auxiliary Information (AI)**: the external knowledge source as the evidences for fact verification.
- **Test Mode (TM)**: the problem setting of the fact-checking task. Specifically, we consider three widely-studied settings, depending on the type of AI. 1) **[claim]**: The target LLM verifies SC without access to external knowledge sources (*i.e.*, AI remains empty), relying solely on the knowledge stored in its parameters. This setting is widely explored in studies utilizing LLMs for fact-checking (Lee et al., 2021). 2) **[evidence]**: AI is the set of gold evidence from Wiki knowledge that can support or refute the source claim. This setting is also called claim verification (Ma et al., 2019). 3) **[wisdom of crowds]**: AI is simulated as the conversation thread on social media towards the source claim. This setting is usually used to verify fake news or rumors by collecting the user interaction as the fact-indicative signal (Shu et al., 2019; Lin et al., 2021).

Quality Inspector Multiple levels of measures are implemented to guarantee the reliability of the fact-checking questions. To check against the generator role like the Inquirer, we further employ a Quality Inspector agent as the judge role to ensure the diversity of the fact-checking topics and the quality of AI. Especially, in the setting of [claim], AI is checked to be empty. In the setting of [evidence], the Quality Inspector would first integrate external tools to coarsely validate whether the evidence set is more likely from Wiki knowledge via Wikipedia web API, then leverage the rich knowledge embedded in the dominant LLM to finely scrutinize the pieces of evidence. In the setting of [wisdom of crowds], the Quality Inspector would make sure that the simulated user comments towards the source claim should be valuable enough as the wisdom of crowds for fact verification.

3.3 Fact Verification with Justification

After obtaining the prototype set of fact-checking test data, we evaluate target LLMs automatically on both fact verification and justification production.

Evaluator The Evaluator agent scores the target LLM’s predicted verdict and generated justification via LLM-as-a-Judge (Zheng et al., 2023), *i.e.*, $\mathcal{F}_\alpha(x)$. In this way, we can assess the fact-checking limits of the target LLM under the hypothesized oracle distribution, *i.e.*, the objective $p(x)\mathcal{F}_\alpha(x)$.

Specifically, given a specific test case x , the target LLM α generates its response r consisting of predicted verdict and derived justification. Then the output of the Evaluator agent includes an integer rating grade $s \in [1, 10] \cap \mathbb{Z}$ and a natural language assessment comment c . As a higher score s indicates stronger capabilities, the corresponding fact-checking limitations can be denoted as

$\mathcal{F}_\alpha(x) \propto 1/s$. As illustrated in Figure 2, we formally define a FACT-AUDIT’s memory pool as $\mathcal{M} = \{x, r, s, c\}$, which stores the test cases along with their evaluation results. The Evaluator is additionally instructed to distinguish the poorly-performing test cases $\{x|s < \epsilon\}$ based on a predefined threshold ϵ . Note that an invalid justification could still get a relatively low grade even if the predicted verdict is correct for fact verification.

Prober While collaboration among the four agent roles (*i.e.*, Appraiser, Inquirer, Quality Inspector, and Evaluator) ensures comprehensive fact-checking coverage and model-specific tailoring of our framework, a key challenge lies in effectively identifying areas where the target model underperforms. Although fact-checking prototype test cases provide an intuitive yet superficial assessment of the target LLM’s fact-checking capabilities, we argue that they are insufficient to fully reveal the fact-checking limitations and knowledge boundaries due to the inherent constraints of fixed seeds.

To craft more diverse and unseen test data about each test scenario, we propose iteratively probing for a more comprehensive fact-checking evaluation. Specifically, given the memory pool \mathcal{M} for the current test scenario, the Prober ρ generates new test data by learning from the model behaviors of the past auditing records stored in \mathcal{M} as the environmental feedback, $x \sim \rho(\mathcal{M})$. Then the Evaluator agent assesses the target LLM on the new test data, and the results are subsequently added to the memory pool. Through this iterative probing process, we can effectively identify test data with poor performance under each fact-checking test scenario, pinpointing comprehensive insights into the target LLM at the adaptive and different test scenarios.

3.4 Adaptive Updating

After going through all the existing test scenarios in the fact-checking taxonomy, the Appraiser appeals to new valuable test scenarios, by conducting a critical analysis of instances where the target LLM underperformed in each fact-checking scenario, as indicated by low rating grades in the memory pool \mathcal{M} , to unveil potential fact-checking limitations. Theoretically, the transition probability $\pi(\Theta_{i+1}|\Theta_i, \mathcal{M})$ is estimated, where Θ_{i+1} is more likely to contain the new test scenario beyond the fact-checking capacities of the target LLM. This insight prompts the Appraiser to adaptively refine the taxonomy, ensuring our framework remains

relevant and effective in identifying new deficiencies. The cyclical interaction among the Appraiser, Inquirer, and Evaluator establishes a continuous improvement loop, making our auditing framework both comprehensive and responsive to the evolving fact-checking capabilities of different target LLMs.

Finally, after the adaptive updating, the expectation of $\mathcal{F}_\alpha(x)$ in Equation (2) for importance sampling can be computed approximately as:

$$\begin{aligned} \mathbb{E}_{q(x)} \left[\mathcal{F}_\alpha(x) \frac{p(x)}{q(x)} \right] &\leq \mathbb{E}_{q(x)} [\mathcal{F}_\alpha(x)] \\ &\propto \frac{1}{|\mathcal{M}|} \sum_{s \in \mathcal{M}} \frac{1}{s}, \end{aligned} \quad (3)$$

where the distributions $q(x)$ and $p(x)$ are intractable in practice. Therefore, since the whole process can only perform limited sampling within the high-probability region ($p(x)/q(x) < 1$) of $q(x)$, we compute an upper bound of the target LLM’s limitations to effectively reflect its utility.

Overall, this framework enables the adaptive sampling of more targeted and representative fact-checking data, facilitating a comprehensive evaluation of the target LLM’s fact-checking capabilities.

4 Experiments and Results

4.1 Experimental Setup

Data Different from existing static data work, the data within the FACT-AUDIT agentic framework is dynamically updated to alleviate sampling bias and fairness issues from a fresh perspective. We consider common fact-checking objects such as complex claims, fake news, and social rumors, simulating a diverse real-world data environment.

Metric To audit the fact-checking capacities of LLMs, we introduce three automatic evaluation metrics for quantitative analysis: Insight Mastery Rate (*IMR*), Justification Flaw Rate (*JFR*), and *Grade*. Specifically, *IMR* represents the proportion of low-scoring fact-checking responses relative to the total number of questions, where a *Grade* of three or below (on a ten-point scale) indicates errors in the target LLM’s response, as the Evaluator agent was additionally instructed not to assign a grade higher than three if the target LLM underperformed in either the verdict prediction or justification production stages. *JFR* denotes the percentage of cases where the target LLM conducted correct verdict predictions yet had poor justification, based on the conditions set by *IMR*. *Grade* is assigned by the FACT-AUDIT framework with employing the scoring prompt inspired by Zheng et al. (2023).

| Model (Target LLM) | Complex Claim | | | Fake News | | | Social Rumor | | | Overall | | |
|--------------------|---------------|--------------|----------------|--------------|--------------|----------------|--------------|--------------|----------------|--------------|--------------|----------------|
| | <i>IMR</i> ↓ | <i>JFR</i> ↓ | <i>Grade</i> ↑ | <i>IMR</i> ↓ | <i>JFR</i> ↓ | <i>Grade</i> ↑ | <i>IMR</i> ↓ | <i>JFR</i> ↓ | <i>Grade</i> ↑ | <i>IMR</i> ↓ | <i>JFR</i> ↓ | <i>Grade</i> ↑ |
| Mistral-7B | 60.21 | 25.62 | 3.98 | 47.50 | 19.58 | 4.87 | 59.05 | 39.52 | 3.97 | 54.79 | 23.34 | 4.34 |
| Llama2-7B | 46.67 | 19.79 | 4.85 | 32.73 | 18.18 | 5.54 | 62.86 | 26.67 | 3.89 | 45.49 | 20.68 | 4.88 |
| Llama2-13B | 65.67 | 21.66 | 3.71 | 55.33 | 16.67 | 4.42 | 48.10 | 20.48 | 4.78 | 57.28 | 19.50 | 4.25 |
| Llama3-8B | 39.79 | 12.09 | 5.19 | 33.75 | 17.28 | 5.51 | 46.25 | 19.18 | 4.83 | 38.67 | 15.60 | 5.25 |
| Llama3.1-8B | 55.83 | 21.46 | 4.36 | 36.39 | 12.78 | 5.60 | 47.62 | 12.86 | 5.00 | 47.52 | 16.77 | 4.91 |
| Llama3.1-70B | 41.56 | 14.22 | 5.34 | 25.00 | 11.88 | 6.42 | 38.33 | 10.00 | 5.55 | 34.10 | 12.38 | 5.83 |
| Qwen2.5-7B | 38.97 | 9.74 | 5.38 | 21.54 | 8.20 | 6.58 | 36.67 | 5.42 | 5.68 | 31.76 | 8.14 | 5.91 |
| Qwen2.5-72B | <u>22.08</u> | <u>5.41</u> | <u>6.62</u> | 10.42 | 1.46 | 7.67 | <u>15.00</u> | 3.75 | 7.28 | <u>16.00</u> | 3.50 | <u>7.17</u> |
| GLM4-9B | 52.73 | 16.36 | 4.76 | 51.67 | 14.00 | 4.93 | 50.00 | 15.24 | 5.00 | 51.67 | 15.24 | 4.88 |
| Gemma2-9B | 41.67 | 28.00 | 4.84 | 35.48 | 28.11 | 5.13 | 44.07 | 23.31 | 4.74 | 39.70 | 26.78 | 4.94 |
| Gemini-Pro | 30.21 | 11.87 | 5.98 | 19.39 | 5.76 | 6.59 | 32.86 | 5.72 | 5.78 | 27.25 | 8.62 | 6.14 |
| Claude3.5-Sonnet | 32.71 | 9.37 | 6.16 | 15.00 | <u>2.33</u> | <u>7.41</u> | 18.57 | <u>3.33</u> | <u>7.31</u> | 24.34 | 5.96 | 6.78 |
| GPT-4o | 14.05 | 4.34 | 6.78 | <u>10.56</u> | 4.93 | 7.26 | 10.48 | 1.41 | 7.62 | 12.02 | <u>3.55</u> | 7.21 |

Table 1: The fact-checking performance of different LLMs audited by FACT-AUDIT. Metrics include *IMR* (%), *JFR* (%), and *Grade*, where *IMR* indicates the insight mastery rate of fact-checking limitations, *JFR* means the flaw rate of the justifications provided by LLMs. The best and second performances are in bold and underlined, respectively.

Overall, *IMR* is the dominant evaluation metric.

Target LLMs To provide a comprehensive LLM auditing, we select 13 representative models as the target LLMs to perform zero-shot inference in FACT-AUDIT. We adopt ten open-source models: Mistral (7B) (Jiang et al., 2023), Llama2 (7B, 13B) (Touvron et al., 2023b), Llama3 (8B) (Dubey et al., 2024), Llama3.1 (8B, 70B), Qwen2.5 (7B, 72B) (Yang et al., 2024a), GLM4 (9B) (GLM et al., 2024), Gemma2 (9B) (Team et al., 2024); and three proprietary models: Gemini-Pro (Team et al., 2023), Claude3.5-Sonnet, GPT-4o, as our target LLMs. To ensure results are reproducible, the temperature is set as 0 without any sampling mechanism. More implementation details and baseline descriptions are provided in Appendix §A - §C.

4.2 Main Results

Table 1 presents the auditing results of various LLMs in FACT-AUDIT, offering a new perspective on fact-checking by incorporating automatic justification production evaluation alongside verdict prediction. Key observations include:

- *GPT-4o*, *Qwen2.5-72B*, *Claude3.5-Sonnet*, and *Gemini-Pro* form the leading tier. Note that *GPT-4o*, *Claude3.5-Sonnet*, and *Gemini-Pro* are proprietary closed-source models, while *Qwen2.5-72B* is an open-source model that demonstrates comparable performance in fact-checking evaluation. Besides, *GPT-4o* achieves the best performance 12.02% on the dominant metric, *IMR*.
- *The LLaMA series exhibits relatively poorer performance*, spanning the second and third tiers. *Llama3-8B* and *Llama3.1-70B* belong to the second tier, alongside *Qwen2.5-7B* and *Gemma2-*

9B, while other LLaMA models fall into the third tier, showing greater fact-checking limitations on both *IMR* and *Grade* performance.

- The auxiliary metric *JFR* of the strong LLM *GPT-4o* is not the best among all target LLMs, as most low-scoring cases are more likely to be poor justifications when a model excels in factual verdict prediction. This implies FACT-AUDIT *could elicit the fact-checking limitation of individual target LLMs in accordance with their aptitude*.
- LLMs perform relatively well on fake news but struggle with complex claims. This discrepancy may stem from the advanced reasoning capabilities required for complex claims compared to the more factually explicit nature of fake news. The fluctuating performance on social rumors is primarily attributed to their contextual dependence and linguistic complexity, which increase the difficulty of fact-checking for target LLMs.

Overall, the automatic model-centric evaluation, considering justifications beyond verdicts, aligned with intuitive expectations of LLM capabilities and introduced additional fresh dimensions for auditing fact-checking performance and limitations.

4.3 Analysis of Reliability

To verify the robustness and fairness of the LLM-generated prototypes, we further conducted the ablative study by adding a setting based on the human-generated prototype seed questions. Specifically, we sampled the same amount of prototypes from the Pinocchio dataset (Hu et al., 2024b) as the fixed seed data in FACT-AUDIT. As shown in Table 2, it can be observed that the performance of the ‘LLM-Generated’ setting is comparable to that of

| Prototype | LLM-Generated | | | Human-Generated | | |
|-------------|---------------|--------------|--------------|-----------------|--------------|--------------|
| | <i>IMR</i> % | <i>JFR</i> % | <i>Grade</i> | <i>IMR</i> % | <i>JFR</i> % | <i>Grade</i> |
| Llama3.1-8B | 55.24 | 21.46 | 4.34 | 55.83 | 21.21 | 4.36 |
| Qwen2.5-7B | 38.97 | 9.74 | 5.38 | 39.62 | 9.93 | 5.25 |
| GPT-4o | 14.05 | 4.34 | 6.78 | 14.24 | 5.02 | 6.59 |

Table 2: The comparison of LLM performance based on LLM-generated and human-generated prototypes.

| Target LLM | Test Mode | <i>IMR</i> % | <i>JFR</i> % | <i>Grade</i> |
|-------------|--------------------|--------------|--------------|--------------|
| Llama3.1-8B | [claim] | 68.80 | 22.87 | 3.56 |
| | [evidence] | 38.16 | 13.33 | 5.50 |
| | [wisdom of crowds] | 45.29 | 16.08 | 4.96 |
| Qwen2.5-7B | [claim] | 48.86 | 12.76 | 4.74 |
| | [evidence] | 20.83 | 7.31 | 6.45 |
| | [wisdom of crowds] | 39.58 | 7.40 | 5.43 |
| GPT-4o | [claim] | 23.05 | 16.67 | 6.11 |
| | [evidence] | 10.61 | 8.77 | 7.00 |
| | [wisdom of crowds] | 15.40 | 8.51 | 6.67 |

Table 3: The fact-checking performance of three representative LLMs under three fixed test modes.

the ‘Human-Generated’ setting, which highlights the fairness of the LLM auditing in FACT-AUDIT. We further provided comprehensive human subject studies for quality assurance in Appendix §D - §E.

4.4 Performance by Test Modes

To thoroughly examine the impact of different test modes on model performance, we evaluate three representative LLMs (Llama3.1-8B, Qwen2.5-7B, and GPT-4o) in the context of fact-checking. As shown in Table 3, we can observe that: 1) [claim] mode is the most challenging for LLMs, as they must rely solely on their parametric knowledge to verify factuality in a closed-book setting. 2) [evidence] mode is the easiest, as all evidence provided is factual and facilitates reasoning, even when conflicting viewpoints are present. 3) [wisdom of crowds] mode falls in the middle. Unlike [claim] mode, it does not depend entirely on the LLM’s internal knowledge, and unlike [evidence] mode, it does not explicitly provide guiding signals. Instead, the model must extract valuable insights from the simulated conversation thread to reason effectively. More detailed results are shown in Appendix §F.

4.5 Challenging Test Scenarios

As shown in Figure 4, we conduct an analysis to discuss the challenging test scenarios in FACT-AUDIT, by taking the *IMR* performance of the well-performed open-source LLM Qwen2.5-72B as an example. We can see that: (1) ‘‘Multi-Step Reasoning’’ (MSR) and ‘‘Aggregated Statistical Reasoning’’ (ASR), (2) News content with ‘‘Mismatched




Figure 4: *IMR* of two most challenging test scenarios in each fact-checking object, with the averaged *IMR*.




Figure 5: The effect of different iterations in scaling.

Headline or Caption’’ (MHC) and ‘‘ManiPulated Content’’ (MPC), and (3) ‘‘Wishing Rumor’’ (WR) and ‘‘Dreading Rumor’’ (DR), are the two most challenging scenarios for Complex Claim, Fake News, and Social Rumor, respectively. Besides, although the Fake News averaged (FN_avg) *IMR* is lower than the Complex Claim averaged (CC_avg) and Social Rumor averaged (SR_avg), the detailed MHC scenario of Fake News is the most difficult than those scenarios of Complex Claim and Social Rumor. We additionally provided more discussion about fact-checking topics in Appendix §G - §H.

4.6 Iterative Probing

We analyze the effect of multi-turn iterative probing in FACT-AUDIT by examining the *IMR* performance across different iterations, as illustrated in Figure 5. The results show that the *IMR* metric decreases as the number of iterations increases, eventually converging. As the test data expands, the model’s performance is more comprehensively evaluated, allowing the identification of truly problematic cases that represent the model’s inherent weaknesses. This approach enables a deeper exploration of less obvious model limitations while reinforcing the flexibility of our evaluation framework to scale seamlessly with the size of the assessment.

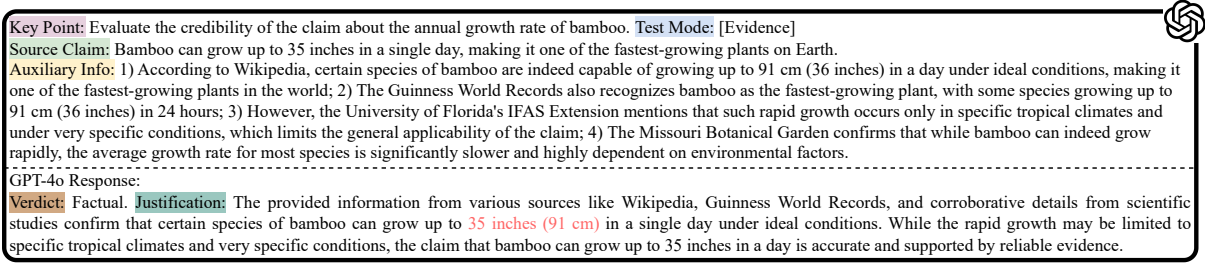


Figure 6: Example of cases with correct verdict prediction yet poor justification for fact-checking the claim.

4.7 Case Study

One key advantage of FACT-AUDIT is that the limitations of fact-checking in LLMs where poor justification is provided for the correct verification could be well probed. To better understand how FACT-AUDIT proactively unveil such limitations, we conduct a case study on the GPT-4o’s response to the test data sample, as exemplified in Figure 6. We can observe that there is a factual error in the provided justification. The target LLM states that bamboo can grow “up to 35 inches (91 cm)”, which conflicts with the unit conversion knowledge that 35 inches is equivalent to 88.9 cm. Even though the related knowledge is provided in the auxiliary information, the target LLM still failed to provide precise justification for fact-checking the claim. This reaffirms the importance of incorporating justification into automatic evaluations, enabling more comprehensive auditing beyond merely assessing accuracy. More cases are shown in Appendix §I.

5 Related Work

Fact-Checking Evaluation Automated fact-checking has gained significant attention in the NLP research community in recent years as a means of combating misinformation and disinformation. Various datasets have been proposed that enable the development and evaluation of systems for automatic fact-checking, the most popular ones being based on human-crafted claims from Wikipedia content (Thorne et al., 2018; Sathe et al., 2020; Schuster et al., 2021), claims in fake news published by a news outlet (Buntain and Golbeck, 2017; Shu et al., 2020; Nakov et al., 2022), rumorous claims on social media (Ma et al., 2015, 2017; Lin et al., 2022a), complex claims that require multi-step reasoning (Jiang et al., 2020a; Aly et al., 2021), naturally occurring claims in specific domains (Gupta and Srikumar, 2021; Wadden et al., 2022; Lin et al., 2023), and LLM-generated misinformation (Chen and Shu, 2024), etc. To understand the factual knowledge of LLMs, Hu et al.

(2024b) curated a new fact-checking benchmark by organizing previous representative datasets, aiming to identify weaknesses in LLM fact verification. However, besides the inevitable issue of test set leakage, this static evaluation approach relied primarily on expert-designed, specialized tasks from existing datasets, overlooking emerging LLM-generated content and lacking adaptability to the complex, open-ended nature of real-world applications. Different from previous work on static accuracy evaluation, leveraging the derived justification (Atanasova et al., 2020; Guo et al., 2022) from LLMs, our work aims to explore the dynamic auditing beyond the veracity prediction, to dynamically elicit the limitations of fact-checking in LLMs.

LLM Agent The integration of LLMs as agents spans various domains, such as code generation and game-playing, demonstrating their robust planning and reasoning capabilities across diverse contexts (Wang et al., 2023; Yao et al., 2022; Shen et al., 2023; Mu et al., 2023; Hong et al., 2023; Liu et al., 2023; Sun et al., 2023; Qian et al., 2023). These advancements highlight the ability of LLMs to handle complex tasks with minimal supervision. In parallel, self-improvement methodologies (Chen et al., 2022, 2023; Shinn et al., 2023; Madaan et al., 2023) have emerged, utilizing feedback-driven processes to iteratively enhance output quality. Building on these insights, we develop a novel agentic framework for systematical LLM auditing in fact-checking complex claims, fake news or rumors.

6 Conclusion and Future Work

We introduced FACT-AUDIT, an adaptive multi-agent evaluation framework that dynamically elicits the fact-checking limitations of LLMs. By automatically evaluating the justification production beyond the verdict prediction, FACT-AUDIT enables scalable, model-centric LLM auditing for fact-checking tasks. Experiments on a dozen mainstream LLMs reveal a notable performance gap between closed and open-source models with dif-

ferent sizes. In future work, we plan to further exploit the reliability of the proposed framework.

Limitations

There are multiple ways for further improvement of this work to alleviate the following limitations:

- Firstly, despite implementing various measures, such as error-correction mechanisms and human evaluations, to enhance the stability and transparency of the agent controller and reduce bias and errors, we argue that the potential biases in fact-checking (much like those inherent to humans) remain unavoidable. Even human beings or the most advanced models have knowledge bias. In future research, we will continue updating the evaluation framework to a more robust and reliable evaluation framework. This would constitute another targeted area of research.
- Secondly, despite its vast knowledge reserves, the agent controller is constrained by its limited ability to acquire and integrate new information dynamically. This limitation hinders its capacity to adapt to evolving knowledge landscapes. In future work, we aim to incorporate advanced techniques such as Retrieval-Augmented Generation (RAG) to enhance the agent’s decision-making capabilities, enabling it to access up-to-date information and provide more accurate, context-aware responses.
- Lastly, while our multi-agent evaluation framework adaptively and dynamically identifies specific deficiencies in target LLMs related to fact verification and justification production, it currently lacks an effective mechanism for model improvement. In future work, we aim to integrate preference optimization methodologies, enabling the framework not only to audit the fact-checking capabilities of LLMs and generate actionable insights for performance refinement but also to provide high-quality training data to facilitate effective model improvement.

Ethics Statement

This research involved human subject studies to evaluate the quality and reliability of FACT-AUDIT. The following considerations were adhered to ensure the protection and ethical treatment of participants: 1) Voluntary Participation: All participants

were informed about the nature of the research and their role in it. Participation was entirely voluntary, with participants having the right to withdraw at any time without any consequences. 2) Informed Consent: Written informed consent was obtained from all participants. This consent form detailed the purpose of the research, the procedures involved, potential risks, and measures taken to safeguard participant data. 3) Data Anonymity and Confidentiality: All data collected during the study were anonymized. Personal identifiers were removed to maintain confidentiality and data were stored securely to prevent unauthorized access. 4) Minimal Risk: The study involved minimal risk to participants. The tasks performed were similar to everyday activities, and no sensitive personal information was requested or recorded.

Research indicates that evaluating content related to misinformation can have negative effects. To protect our human evaluators, we establish three guidelines: 1) ensuring their acknowledgment of viewing potentially misleading content, 2) limiting weekly evaluations and encouraging a lighter daily workload, and 3) advising them to stop if they feel overwhelmed. Finally, we regularly check in with evaluators to ensure their well-being.

The purpose of this work is to prevent the spread of misinformation/disinformation and to ensure that people are not subjected to non-factual information. Nevertheless, we are aware of the potential for malicious users to reverse-engineer and create misinformation guided by FACT-AUDIT. This is strongly discouraged and condemned. Furthermore, all the fact-checking test data generated by the agents does not contain any personal information.

Acknowledgments

This work is partially supported by Tencent Rhino-Bird Focused Research Program (Value-aligned Credible Large Language Model), RMGS project (Artificial Intelligence and Big Data Analytics for Social Good), and the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant (No. MSS24C012).

References

- Gordon W Allport. 1947. The psychology of rumor. *Holt, Rinehart & Winston*.
- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos

- Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. Feverous: Fact extraction and verification over unstructured and structured information. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. Generating fact checking explanations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7352–7364.
- 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*.
- Cody Buntain and Jennifer Golbeck. 2017. Automatically identifying fake news in popular twitter threads. In *2017 IEEE international conference on smart cloud (smartCloud)*, pages 208–215. IEEE.
- Boxi Cao, Hongyu Lin, Xianpei Han, Le Sun, Lingyong Yan, Meng Liao, Tong Xue, and Jin Xu. 2021. Knowledgeable or educated guess? revisiting language models as knowledge bases. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1860–1874.
- Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. 2022. Codet: Code generation with generated tests. In *The Eleventh International Conference on Learning Representations*.
- Canyu Chen and Kai Shu. 2024. Can llm-generated misinformation be detected? In *The Twelfth International Conference on Learning Representations*.
- Xinyun Chen, Maxwell Lin, Nathanael Schaerli, and Denny Zhou. 2023. Teaching large language models to self-debug. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Jiale Cheng, Yida Lu, Xiaotao Gu, Pei Ke, Xiao Liu, Yuxiao Dong, Hongning Wang, Jie Tang, and Minlie Huang. 2024. Autodetect: Towards a unified framework for automated weakness detection in large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 6786–6803.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. 2024. Length-controlled alpaca-eval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhishava Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. Measuring and improving consistency in pretrained language models. *Transactions of the Association for Computational Linguistics*, 9:1012–1031.
- Islam Eldifrawi, Shengrui Wang, and Amine Trabelsi. 2024. Automated justification production for claim veracity in fact checking: A survey on architectures and approaches. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6679–6692.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, et al. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Ashim Gupta and Vivek Srikumar. 2021. X-factor: A new benchmark dataset for multilingual fact checking. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 675–682.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, et al. 2023. Metagpt: Meta programming for multi-agent collaborative framework. In *The Twelfth International Conference on Learning Representations*.
- Beizhe Hu, Qiang Sheng, Juan Cao, Yuhui Shi, Yang Li, Danding Wang, and Peng Qi. 2024a. Bad actor, good advisor: Exploring the role of large language models in fake news detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 22105–22113.
- Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, S Yu Philip, and Zhijiang Guo. 2024b. Do large language models know about facts? In *The Twelfth International Conference on Learning Representations*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Yichen Jiang, Shikha Bordia, Zheng Zhong, Charles Dognin, Maneesh Singh, and Mohit Bansal. 2020a. Hover: A dataset for many-hop fact extraction and claim verification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3441–3460.

- Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. 2020b. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Herman Kahn and Andy W Marshall. 1953. Methods of reducing sample size in monte carlo computations. *Journal of the Operations Research Society of America*, 1(5):263–278.
- Nayeon Lee, Yejin Bang, Andrea Madotto, and Pascale Fung. 2021. Towards few-shot fact-checking via perplexity. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1971–1981.
- Hongzhan Lin, Jing Ma, Liangliang Chen, Zhiwei Yang, Mingfei Cheng, and Chen Guang. 2022a. Detect rumors in microblog posts for low-resource domains via adversarial contrastive learning. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2543–2556.
- Hongzhan Lin, Jing Ma, Mingfei Cheng, Zhiwei Yang, Liangliang Chen, and Guang Chen. 2021. Rumor detection on twitter with claim-guided hierarchical graph attention networks. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10035–10047.
- Hongzhan Lin, Pengyao Yi, Jing Ma, Haiyun Jiang, Ziyang Luo, Shuming Shi, and Ruifang Liu. 2023. Zero-shot rumor detection with propagation structure via prompt learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 5213–5221.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022b. Teaching models to express their uncertainty in words. *arXiv preprint arXiv:2205.14334*.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. 2023. Agentbench: Evaluating llms as agents. In *The Twelfth International Conference on Learning Representations*.
- Jing Ma, Wei Gao, Shafiq Joty, and Kam-Fai Wong. 2019. Sentence-level evidence embedding for claim verification with hierarchical attention networks. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2561–2571.
- Jing Ma, Wei Gao, Zhongyu Wei, Yueming Lu, and Kam-Fai Wong. 2015. Detect rumors using time series of social context information on microblogging websites. In *Proceedings of the 24th ACM international on conference on information and knowledge management*, pages 1751–1754.
- Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect rumors in microblog posts using propagation structure via kernel learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 708–717.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. 1953. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092.
- Nicholas Metropolis and Stanislaw Ulam. 1949. The monte carlo method. *Journal of the American statistical association*, 44(247):335–341.
- Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. 2023. Embodiedgpt: Vision-language pre-training via embodied chain of thought. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Preslav Nakov, Alberto Barrón-Cedeño, Giovanni Da San Martino, Firoj Alam, Julia Maria Struß, Thomas Mandl, Rubén Míguez, Tommaso Caselli, Mucahid Kutlu, Wajdi Zaghouni, et al. 2022. The clef-2022 checkthat! lab on fighting the covid-19 infodemic and fake news detection. In *European Conference on Information Retrieval*, pages 416–428. Springer.
- OpenAI. 2023. [Gpt-4 technical report](https://arxiv.org/abs/2303.08774). *ArXiv*, abs/2303.08774.
- 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. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6981–7004.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.
- 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)*, pages 2463–2473.
- Chen Qian, Xin Cong, Cheng Yang, Weize Chen, Yusheng Su, Juyuan Xu, Zhiyuan Liu, and Maosong Sun. 2023. Communicative agents for software development. *arXiv preprint arXiv:2307.07924*.

- Aalok Sathe, Salar Ather, Tuan Manh Le, Nathan Perry, and Joonsuk Park. 2020. Automated fact-checking of claims from wikipedia. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6874–6882.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. Get your vitamin c! robust fact verification with contrastive evidence. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugging-gpt: Solving ai tasks with chatgpt and its friends in hugging face. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: language agents with verbal reinforcement learning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 8634–8652.
- Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. defend: Explainable fake news detection. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 395–405.
- Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2020. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big data*, 8(3):171–188.
- Haotian Sun, Yuchen Zhuang, Ling kai Kong, Bo Dai, and Chao Zhang. 2023. Adaplaner: Adaptive planning from feedback with language models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819.
- 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 arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Iz Beltagy, Lucy Lu Wang, and Hannaneh Hajishirzi. 2022. Scifact-open: Towards open-domain scientific claim verification. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4719–4734.
- M Mitchell Waldrop. 2017. The genuine problem of fake news. *Proceedings of the National Academy of Sciences*, 114(48):12631–12634.
- Bo Wang, Jing Ma, Hongzhan Lin, Zhiwei Yang, Ruichao Yang, Yuan Tian, and Yi Chang. 2024. Explainable fake news detection with large language model via defense among competing wisdom. In *Proceedings of the ACM on Web Conference 2024*, pages 2452–2463.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An open-ended embodied agent with large language models. In *Intrinsically-Motivated and Open-Ended Learning Workshop@ NeurIPS2023*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024a. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Ruichao Yang, Wei Gao, Jing Ma, Hongzhan Lin, and Bo Wang. 2024b. Reinforcement tuning for detecting stances and debunking rumors jointly with large language models. *arXiv preprint arXiv:2406.02143*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 46595–46623.

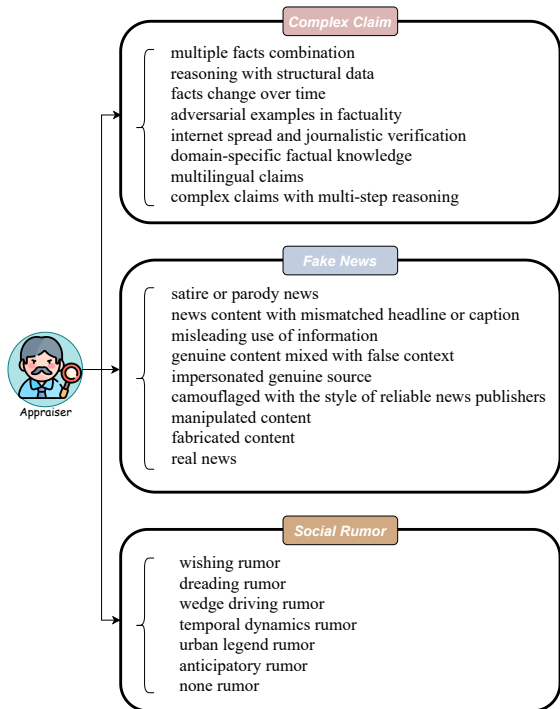


Figure 7: Illustration of the initial taxonomy on the fact-checking objects: Complex Claim, Fake News, Social Rumor.

A Taxonomy

We provide the initial taxonomy in Figure 7. The detailed taxonomy of the three fact-checking objects draws the practice of the previous fact-checking literature: 1) Complex Claim (Pan et al., 2023; Hu et al., 2024b) involves assertions that require detailed analysis and support from multiple sources and are common in scientific or academic discussions, 2) Fake News (Waldrop, 2017) refers to deliberately fabricated or distorted information aimed at misleading audiences, often seen on social media to influence public opinion or for economic gain, and 3) Social Rumor (Allport, 1947) is a piece of information that spreads quickly and remains to be verified, usually through word of mouth or social media, and can lead to misunderstandings or unnecessary panic. The choice of complex claims, fake news, and rumors as fact-checking objects stems from their prominent impact on public discourse and their prevalence in today’s information landscape. In FACT-AUDIT, the final test scenarios of each fact-checking object would be evolved and updated according to the model-specific performance. Due to the dynamic nature, we provide the averaged statistics of the data in FACT-AUDIT as shown in Table 4.

The taxonomy of fact-checking objects is systematically designed to address the diverse forms

of misinformation based on their intrinsic characteristics, verification challenges, and real-world impact. Specifically, the following principles are used to guide the categorization: 1) Complexity: The level of reasoning and factual knowledge required to validate the claim. 2) Intent and Structure: Whether the content aims to mislead, parody, or inform and how it is presented. 3) Propagation Dynamics: The nature and speed at which rumors or misinformation spread within social contexts. The proposed taxonomy serves as the foundation for systematically evaluating fact-checking capacities in LLMs. By dividing fact-checking objects into Complex Claims, Fake News, and Social Rumors, the framework achieves the following objectives: 1) Targeted Evaluation: Addressing the unique verification challenges posed by each category. 2) Comprehensive Coverage: Ensuring that the taxonomy encompasses a wide range of misinformation types prevalent in real-world scenarios. 3) Practical Utility: Facilitating the generation of more targeted and representative fact-checking datasets to evaluate model performance. This taxonomy is designed to systematically uncover fact-checking limitations in LLMs by segmenting diverse fact-checking objects into detailed test scenarios. Each category reflects the nature, complexity, and propagation style of the potential true or false information, enabling a more structured and comprehensive evaluation framework.

B Implementation Details

For all experiments, we adopt GPT-4o as the core model for FACT-AUDIT. For importance sampling, we formalize the probability density of each fact-checking data as a uniform distribution, to mitigate potential long-tail issues. Compared results ($p < 0.05$ under t-test) are averaged over three random runs. The maximum number of iterations is set to 30 for evaluations on each fact-checking test scenario. The threshold ϵ for the poorly-performing test cases is set as 4.0. The cost for evaluating one target LLM is about 25 dollars and 6 hours. All experiments were conducted using two NVIDIA A100 80GiB GPUs. In the following, the details of agent implementation would be depicted.

Appraiser. For the Appraiser agent, the taxonomy is initialized as shown in Figure 7. We set the temperature of the Appraiser agent as the default setting of 1.0. To update the taxonomy, the instruction prompt is used as shown in Figure 8. If the

| Model | Complex Claim | Fake News | Social Rumor | Overall |
|------------------|---------------|-----------|--------------|---------|
| Mistral-7B | 480 | 480 | 210 | 1170 |
| Llama2-7B | 480 | 330 | 210 | 1020 |
| Llama2-13B | 300 | 300 | 210 | 810 |
| Llama3-8B | 480 | 480 | 240 | 1200 |
| Llama3.1-8B | 480 | 360 | 210 | 1050 |
| Llama3.1-70B | 450 | 480 | 240 | 1170 |
| Qwen2.5-7B | 390 | 390 | 240 | 1020 |
| Qwen2.5-72B | 480 | 480 | 240 | 1200 |
| GLM4-9B | 330 | 300 | 210 | 840 |
| Gemma2-9B | 300 | 420 | 270 | 990 |
| Gemini-Pro | 480 | 330 | 210 | 1020 |
| Claude3.5-Sonnet | 480 | 300 | 210 | 990 |
| GPT-4o | 420 | 360 | 210 | 990 |

Table 4: The averaged data statistics of the dynamically-updated auditing framework corresponding to each specific target LLM.

Fact-checking is an important capability of LLMs, where the LLM should analyze textual information to identify the factuality of the source claim. Here, the LLM must be tested to accurately assess the factuality of the information presented within the source claim according to the claim itself or the auxiliary information.

Here is a subtask's taxonomy as well as the averaged score on these tasks (lower means worse performance):
`{taxonomy_with_grade_on_each_test_scenario}`

And here are some bad cases:
`{bad_cases}`

Based on the given information, please judge if the taxonomy is comprehensive, if so please just output `[[Stop]]`.

If not, please give me a new possible issue you inferred from the present taxonomy and bad cases. Please focus on `{fact-checking_object}`. Ensure the new task is text-only (no multimodal). Also give a brief explanation of how you find the issue. Please output in a JSON format:
`{"task_name": ..., "explanation": ...}`

Figure 8: Instruction for the Appraiser agent to generate new valuable test scenarios.

Appraiser outputs the “[stop]” tokens in three times, the adaptive updating process would be terminated.

Note that as the new subject task cannot be always added, we additionally apply a judge agent to check the quality of the new proposed test scenario. The prompt for judging the new test scenario is shown in Figure 9.

Inquirer. The role of the Inquirer is to generate the prototype fact-checking data. To ensure the fairness, we set the temperature of the Inquirer agent as 0.0 without any sampling mechanism. The number of prototype seed questions for each test scenario is set as 10. The instruction prompt is designed as shown in Figure 10.

Quality Inspector. To ensure the quality of the LLM-generated fact-checking data, the Quality Inspector agent is deployed to use external tools and inspect whether the generated data conforms to the basic requirements. The parameter temperature is set as 0.0 since here we do need the most reliable

Fact-checking is an important capability of LLMs, where the LLM should analyze textual information to identify the factuality of the source claim. Here, the LLM must be tested to accurately assess the factuality of the information presented within the source claim according to the claim itself or the auxiliary information.

Here is a subtask's taxonomy on the task "`{fact-checking_object}`":
`{taxonomy}`

Based on the given taxonomy, please judge whether the new test point "`{new_scenario}`" is suitable as a subtask on the task "`{fact-checking_object}`". The judge criteria are as follows:

1. The new test point should precisely cover an important and meaningful part of the main task.
2. The new test point should be sufficiently different from the existing test points.
3. The new test point should be text-only (no multimodal).

If the new test point "`{new_scenario}`" is suitable as a subtask on the task "`{fact-checking_object}`", please ONLY output `[[Yes]]`. If not, please first output `[[No]]`, and then provide the reason why it's not suitable as a subtask on the task "`{fact-checking_object}`".

Figure 9: Instruction for the Appraiser agent to check whether the new proposed test scenario is suitable to be added into the current taxonomy.

content instead of the generation model’s creativity. First of all, the Wikipedia search API is called to coarsely check the credibility of the auxiliary information if the test scenario is under the [evidence] mode. Then the specific prompt is curated finely as shown in Figure 11.

Evaluator. Some recent pioneering benchmark work (Zheng et al., 2023; Dubois et al., 2024; Cheng et al., 2024) used LLM-as-a-judge to ask the strong LLM to compare model responses to a static dataset of questions. The prompt template of the question $\{SC, AI\}$ for the reference answer of the Evaluator and the answer of the target LLM is shown in Figure 12. The model’s judgments achieved over 80% agreement with human preferences, proving the usability of using LLMs to evaluate response quality. Inspired by the previous

Fact-checking is an important capability of LLMs, where the LLM should analyze textual information to identify the factuality of the source claim. Here, the LLM must be tested to accurately assess the factuality of the information presented within the source claim according to the claim itself or the auxiliary information.

Here is a taxonomy for the fact-checking task:
{taxonomy}

Step 1: Please read the provided initial taxonomy carefully. Based on this, please generate 10 test cases of "{test_scenario}" category, to test if language models can accurately identify facts or misinformation in the source claim on task "{test_scenario}".

Step 2: When generating each test case, consider which one of the three optional test modes is the most suitable: [claim], [evidence], and [wisdom of crowds] for each case. For [claim], the factuality can be verified according to the source claim itself; For [evidence], the factuality of the source claim needs to be verified according to the attached evidence set; For [wisdom of crowds], the factuality of the source claim needs to be assessed from the simulated conversation tree of user comments on social media.

Step 3: Based on the selected test mode in Step 2, if not the [claim] mode is selected, generate the auxiliary information "auxiliary_info" for the source claim. If else, "auxiliary_info" is empty. For "auxiliary_info" of [evidence], please ensure that: 1) more than three pieces of evidence are in "auxiliary_info", and 2) the provided pieces of detailed evidence in "auxiliary_info" must only be ground truth quoted directly and solely from Wikipedia word for word (without any personal insight), where different amounts of supported, refuted, and neutral evidence to the source claim should be included; For "auxiliary_info" of [wisdom of crowds], please ensure that: 1) the depth of the conversation tree in "auxiliary_info" must be more than two, and 2) the hierarchical conversation tree in "auxiliary_info" can be noisy but valuable to help verify the source claim.

Step 4: "Key_point" is a short sentence that summarizes the key point you want to test the language model, clearly stating the target content to be fact-checked. The constraints on "{test_scenario}" should be explicitly expressed. Besides, your test cases should cover common topics in fact-checking and different test modes mentioned before, to increase prompt diversity. Please be as diverse as you can but focus on "{test_scenario}" and ensure the prompt is text-only (no multimodal).

Step 5: Repeat Step 1-4 for each test case and then form all the test cases into a JSON format. The "test_mode" of the test cases should include [claim], [evidence], and [wisdom of crowds].

Please reply strictly in the following format:

Step 1 "source_claim":

Step 2 "test_mode":

Step 3 "auxiliary_info":

Step 4 "key_point":

Step 5 Repeat Step 1-4 for each test case and then output one final JSON format: {"test_case1": {"key_point": string(...), "test_mode": string(...), "prompt": {"source_claim": string(...), "auxiliary_info": string(...)}}, "test_case2": {...}, ...}.

Figure 10: Instruction for the Inquirer agent to generate the prototype fact-checking data.

Fact-checking is an important capability of LLMs, where the LLM should analyze textual information to identify the factuality of the source claim. Here, the LLM must be tested to accurately assess the factuality of the information presented within the source claim according to the claim itself or the auxiliary information.

Please judge whether the new test cases "{new_data}" are suitable as diverse and comprehensive exam questions on the subtask "{test_scenario}". The judgment criteria are as follows:

1. Each claim of the new test cases should be important and meaningful to the task "{test_scenario}", avoiding unnecessary ambiguity in the key point.
2. If "auxiliary_info" is not empty in each of the new test cases, it can be noisy but must be helpful to the fact verification process; If "auxiliary_info" is empty, just keep it empty.
3. If "test_mode" is [claim], "auxiliary_info" must be empty.
4. If "test_mode" is [wisdom of crowds], please check "auxiliary_info" that: a) the user comments in "auxiliary_info" should be valuable enough as the wisdom of crowds for fact verification and b) the depth of the propagation conversation tree composed of the user response in "auxiliary_info" must be a random integer more than two.
5. If "test_mode" is [evidence], please check "auxiliary_info" that: a) four or more random pieces of evidence are in "auxiliary_info", and b) the provided pieces of detailed evidence in "auxiliary_info" must be ONLY ground truth based on Wikipedia or other authority, where all supported, refuted, and neutral evidence to the source claim should be included.
6. The fact-checking topic in each test case should be diverse enough and sufficiently different from each other.

If the new test cases are judged suitable as the exam questions on the subtask "{test_scenario}" by checking the judgment criteria, please ONLY keep the original content "{new_data}" as output in a JSON format: [json]; If there is one test case not conforming to the judgment criteria, you have to revise and improve the original content "{new_data}" to conform to the aforementioned judgment criteria, and ONLY output the improved test cases in a JSON format: [json].

Figure 11: Instruction for the Quality Inspector agent to improve the quality of the fact-checking data.

This is a fact-checking task. Please directly verify the factual accuracy of the statement provided below. Your response should conclude whether the statement in the question is Factual, Non-Factual, or Not Enough Information based on the claim itself, or an attached evidence set, or a given conversation thread of user replies.

Question: {question}

Answer: [Factual, Non-Factual, or Not Enough Information]

Justification:

Figure 12: Instruction for generating the reference answer of the Evaluator agent and the answer of the target LLM for fact-checking data.

literature, we employ an Evaluator agent to evaluate the response of the target LLM in a scoring and comparison-based manner. First, we employ three agent controllers (temperature = 1.0) with the currently dominant LLM GPT-4o, to vote a relatively perfect answer in a self-reflection manner. Nevertheless, to further ensure the quality of the reference answer, we made another judgment agent role (temperature = 0.0) to check the content of the reference answer generated by the Evaluator agent, where the prompt is shown in Figure 13.

In this way, we can alleviate the potential mistakes in the reference answer. Then, we employ the scoring prompt (Zheng et al., 2023), to elicit an evaluation output consisting of a grade and a com-

Fact-checking is an important capability of LLMs, where the LLM should analyze textual information to identify the factuality of the source claim. Here, the LLM must be tested to accurately assess the factuality of the information presented within the source claim according to the claim itself or the auxiliary information.

Please judge whether the answer "{refer_answer}" is accurate and reasonable for the question "{fact-checking_data}" with the subtask "{test_scenario}" to be examined.

If the answer is correct and reasonable, please ONLY output the original answer "{refer_answer}" in a string format; If not, please revise and improve the original answer "{refer_answer}" to be more precise, and ONLY output the improved answer in a string format.

Figure 13: Instruction for the Evaluator agent to ensure the quality of the reference answers for fact-checking data.

```
[Instruction]\n Please act as an impartial judge and evaluate the quality of the response provided by an AI assistant to the fact-checking question displayed below. Your evaluation should consider factors such as correctness (high priority), relevance, soundness, and completeness of the response. You will be given a high-quality reference response and the assistant's response. Begin your evaluation by comparing the assistant's response with the reference response. Identify and correct any mistakes in the answer and its justification. Be as objective as possible. After providing your explanation, you must rate the response on a scale of 1 to 10 by strictly following this format: "[[rating]]", for example: "Rating: [[5]]". Please do not score higher than 3.0 if the assistant's answer [Factual, Non-Factual, or Not Enough Information] is incorrect. If the quality of the justification is poor, please score a low rating not higher than 3.0, even if the answer is correct.\n\n [Question]\n {fact-checking_data}\n\n [Key Point]\n {key_point}\n\n [The Start of Reference Response]\n {refer_answer}\n\n [The End of Reference Response]\n\n [The Start of Assistant's Response]\n {Target LLM_answer}\n\n [The End of Assistant's Response]", "output_format": "[[rating]]"
```

Figure 14: Instruction for the Evaluator agent to conduct the fact-checking evaluation in an LLM-as-a-Judge manner.

ment on the response of the target LLM. Specifically, the prompt is devised as shown in Figure 14.

Note that in our designed evaluation setup, a *Grade* of three or below (on a ten-point scale) was selected to represent errors in the target LLM’s response, based on our analysis of quality differentiation and the practice of previous literature (Cheng et al., 2024). This threshold effectively captures significant issues in either the verdict prediction or justification production stages, ensuring that only responses demonstrating an adequately reliable quality receive a higher score. As we employ the ten-point scale in the evaluation, considering different levels of uniform division, a grade under 4.0 naturally represents a low level, while 7.0 is the dividing line between medium and high levels, which is a reasonable setting in our evaluation system. This threshold was determined through careful consideration of maintaining strict reliability and consistency across the subsequent introduced evaluation metrics.

Prober. All the evaluation output and fact-

checking data would be recorded in a memory pool. Based on the collection of the evaluation history in the memory pool, we deploy a Prober agent (temperature = 1.0) to further explore more comprehensive fact-checking data to query the target LLM. The concrete instruction prompt is designed as shown in Figure 15.

For the baselines, we conduct extensive experiments in FACT-AUDIT to evaluate a total of 13 representative target LLMs:

- **Mistral-7B:** A highly efficient 7-billion parameter open-source large language model optimized for performance, offering state-of-the-art results in various natural language processing tasks while maintaining lightweight computational requirements. We specifically utilize the “Mistral-7B-Instruct-v0.2” version.
- **Llama2-7B:** An advanced 7-billion parameter open-source large language model developed by Meta, designed to deliver strong performance in natural language understanding and generation tasks, with fine-tuning options for specialized applications. We specifically utilize the “Llama-2-7b-hf” version.
- **Llama2-13B:** A 13-billion version of LLaMA 2 series. We specifically utilize the “Llama-2-13b-hf” version.
- **Llama3-8B:** An 8-billion parameter large language model released by Meta in April 2024 as part of the LLaMA 3 series, optimized for dialogue and conversational tasks with the ability to generate natural language text. We specifically utilize the “Meta-Llama-3-8B-Instruct” version.
- **Llama3.1-8B:** An 8-billion parameter large language model released by Meta in July 2024, supporting multilingual dialogue with a 128,000-token context window, designed for various natural language processing tasks. We specifically utilize the “Meta-Llama-3.1-8B-Instruct” version.
- **Llama3.1-70B:** A 70-billion version of LLaMA 3.1 series. We specifically utilize the “Meta-Llama-3.1-70B-Instruct” version.
- **Qwen2.5-7B:** A 7-billion parameter large language model from the Qwen2.5 series, offering enhanced capabilities in coding, mathe-

This task involves generating test cases for the fact-checking task. Fact-checking is an important capability of LLMs, where the LLM should analyze textual information to identify the factuality of the source claim. Here, you need to ask the LLM to be tested to accurately assess the factuality of the information presented with the claim itself, or auxiliary information.

Previous Prompts:

Prompt: {fact-checking_data_1}
 Key Point: {key_point_1}
 Score: {grade_1}
 ...
 Prompt: {fact-checking_data_10}
 Key Point: {key_point_10}
 Score: {grade_10}

The objective is to create new prompts that are challenging for the language model, with a focus on diverse types of instructions about "{test_scenario}". Each prompt should be solvable by a language model completely, and aimed at achieving a lower score (indicating difficulty and complexity).

Guidelines for Creating New Prompts:

1. Each prompt to be fact-checked should be solvable by a language model (no visual task) and should contain all necessary information.
2. Understand how factuality in the Prompt is discerned by the tested model as shown in its Answer and the comment on the score.
3. Aim for prompts that would result in a low score (less than 3.0), indicating a high level of complexity and difficulty of the question that requires more effort to identify the factuality or misinformation.
4. Do not repeat topics across different instructions in the previous prompts to maximize diversity.
5. The "key_point" should be no more than 15 words and summarize the key points of the prompt to clearly state the target content to be fact-checked.
6. The "test_mode" should be one of the three options: 1) [claim], (i.e., only the source claim), or 2) [evidence], (i.e., additional contrastive evidence based on Wikipedia), or 3) [wisdom of crowds], (i.e., user comments on social media).
7. The "auxiliary_info" should be provided according to the "test_mode": if not the [claim] mode is selected, generate the auxiliary information "auxiliary_info" for the source claim. If else, "auxiliary_info" is empty. For "auxiliary_info" of [evidence], please ensure that: 1) more than three pieces of evidence are in "auxiliary_info", and 2) the provided pieces of detailed evidence in "auxiliary_info" must only be ground truth quoted directly and solely from Wikipedia word for word (without any personal insight), where different amounts of supported, refuted, and neutral evidence to the source claim should be included; For "auxiliary_info" of [wisdom of crowds], please ensure that: 1) the depth of the conversation tree in "auxiliary_info" must be more than two, and 2) the hierarchical conversation tree in "auxiliary_info" can be noisy but valuable to help verify the source claim.
8. Please focus on "{test_scenario}" constraints, and ensure that upon careful consideration a human fact-checker with commonsense can identify the factuality of the new prompt.
9. The new prompt should be STRICTLY within 512 words and should not be too long.

Please generate a new test case. Output in a json format: {"key_point": string(...), "test_mode": string(...), "prompt": {"source_claim": string(...), "auxiliary_info": string(...)}}

Figure 15: Instruction for the Prober agent to generate more diverse and unseen fact-checking data.

matics, and instruction following, with support for over 29 languages and a context length of up to 131,072 tokens. We specifically utilize the “Qwen2.5-7B-Instruct” version.

- Qwen2.5-72B: A 72-billion version of Qwen 2.5 series. We specifically utilize the “Qwen2.5-72B-Instruct” version.
- GLM4-9B: An open-source large language model developed by Zhipu AI, featuring 9 billion parameters and supporting 26 languages, including Japanese, Korean, and German. We specifically utilize the “glm-4-9b-chat” version.
- Gemma2-9B: An open-source language model developed by Google, trained on 8 trillion tokens of diverse data, including web documents, code, and mathematical texts, to deliver state-of-the-art performance across various natural language processing tasks. We specifically utilize the “gemma-2-9b-it” version.
- Gemini-Pro: A multimodal large language model developed by Google DeepMind, designed to process and generate text, audio, and visual data, offering advanced reasoning and planning capabilities across various tasks. We specifically utilize the “Gemini-Pro” version.
- Claude3.5-Sonnet: Anthropic’s most advanced large language model, offering enhanced reasoning, coding proficiency, and vision capabilities, with a 200,000-token context window and improved speed and cost efficiency. We specifically utilize the “claude-3-5-sonnet-20240620” version.
- GPT-4o: The latest flagship model developed by OpenAI, designed for real-time reasoning across audio, visual, and textual inputs. We specifically utilize the “gpt-4o-2024-05-13” version.

To ensure the reproducibility, we set the temperature as 0.0 without any sampling mechanism.

C Evaluation Metrics

To audit the fact-checking capacities of LLMs, we introduce three automatic evaluation metrics for quantitative analysis: Insight Mastery Rate (*IMR*),

Justification Flaw Rate (*JFR*), and *Grade*. Specifically, *IMR* represents the proportion of low-scoring fact-checking responses relative to the total number of questions, where a *Grade* of three or below (on a ten-point scale) indicates errors in the target LLM’s response, as the Evaluator agent was additionally instructed not to assign a grade higher than three if the target LLM underperformed in either the verdict prediction or justification production stages. Specifically, the *IMR* metric can be formulated as follows:

$$IMR = \frac{\text{Number of Tests with Grade} \leq 3.0}{\text{Total Number of Tests}}, \quad (4)$$

where *IMR* represents the degree of mastering the fact-checking limitation insight of target LLMs.

JFR denotes the percentage of cases where the target LLM conducted correct verdict prediction yet had poor justification, based on the conditions set by *IMR*. Specifically, the *JFR* metric can be formulated as follows:

$$JFR = \frac{\text{Number of Tests with CVPJ}}{\text{Total Number of Tests}}, \quad (5)$$

where CVPJ denotes the case that the target LLM predicted Correct Verdict, but provided a relatively Poor Justification.

Grade is assigned by the FACT-AUDIT framework with employing the scoring prompt inspired by Zheng et al. (2023). Overall, *IMR* is the dominant evaluation metric.

D Quality Assurance

To guarantee the reliability of the fact-checking data generated by FACT-AUDIT agents, 3 professional fact-checking annotators (between the ages of 26 and 29) were asked to judge whether the quality of fact-checking data on each sample was up to standard or not, including the source claim, the auxiliary information in the test mode [evidence] or [wisdom of crowds], and the reference answer. Thus we randomly sampled 600 pieces, with 200 from each fact-checking object, across all evaluations of target LLMs.

Specifically, as shown in Table 5, the annotators, with the averaged intra-class agreement score 0.669, need to evaluate: 1) whether the factual knowledge taxonomy in the categorization is suitable as the test scenarios in the fact-checking task, for the quality judgment of the detailed categorization; 2) whether the claim is check-worthy in the fact-checking process, for the quality judgment

| Fact-checking Data | Judgment [†] | Agreement [†] |
|--------------------|-----------------------|------------------------|
| Detailed Taxonomy | 98.86 | 0.810 |
| Source Claim | 97.17 | 0.795 |
| Relevant Evidence | 87.00 | 0.619 |
| Wisdom of Crowds | 81.83 | 0.581 |
| Reference Answer | 90.33 | 0.762 |
| Output Evaluation | 89.02 | 0.658 |

Table 5: Human subject study on the reliability of the FACT-AUDIT framework. The Judgment (%) means the proportion of fact-checking data that meets the criteria, and the Agreement denotes the average Cohen’s Kappa between any two expert annotators.

of the source claim; 3) whether the supported, refuted, and neutral snippets included in the auxiliary information to the source claim are all ground truth, for the quality judgment of the relevant evidence as auxiliary information; 4) whether the simulated conversation thread is valuable to conduct fact verification, for the quality judgment of the wisdom of crowds as auxiliary information; 5) whether the reference answer is reasonable and correct, for the quality judgment of the reference answer; 6) whether the grade and comment scored by the agent are acceptable for the auditing basis, for the quality judgment of the output evaluation.

We can observe that: 1) The highest judgment indicates that the taxonomy used for fact-checking tasks is highly suitable, demonstrating the reliability of Fact-Audit in designing test scenarios. 2) Most source claims were judged as check-worthy, reflecting the high quality of the claims generated for fact-checking tasks. 3) For the provided relevant evidence, while still high, the quality score here is slightly lower. This may be due to the complexity of the supporting, refuting, or neutral snippets provided as auxiliary evidence. 4) For the simulated conversation thread as wisdom of crowds, the lowest judgment (81.83%) suggests that extracting valuable information from simulated conversations poses significant challenges. The unstructured nature or semantic ambiguity in dialogues may contribute to this difficulty. 5) Over 90% of the reference answers being correct reflects their good quality. 6) Furthermore, the acceptability of the Evaluator agent’s grades and comments is robust, with a high-quality judgment (89.02%) indicating reliability in the auditing process.

Detailed Taxonomy (0.810) and Source Claim (0.795) demonstrate very high agreement, surpassing or approaching the 0.8 threshold, indicating the objectivity and reliability of these evaluations.

| Benchmark | Pinocchio | LLMFake | FACT-AUDIT |
|--------------|-----------|---------|------------|
| Redundancy↓ | 2.03 | 2.31 | 1.22 |
| Diversity↑ | 1.94 | 2.17 | 2.62 |
| Readability↑ | 2.86 | 2.43 | 2.91 |
| Coverage↑ | 2.14 | 1.65 | 2.58 |
| Fairness↑ | 2.57 | 2.53 | 2.56 |
| Suitability↑ | 2.79 | 2.78 | 2.81 |

Table 6: Human evaluation of the benchmark quality.

Relevant Evidence (0.619) and Wisdom of Crowds (0.581) have lower agreement, especially Wisdom of Crowds (0.581), which falls within the moderate range. This suggests that these tasks involve more subjective judgments and are more challenging to evaluate consistently. Reference Answer (0.762) and Output Evaluation (0.658) show reasonable agreement, although slightly less consistent compared to the detailed taxonomy and source claims.

E Comparison with Traditional Benchmarks

We conduct the human subject study to compare the benchmark quality of our proposed FACT-AUDIT and two other well-known benchmarks Pinocchio (Hu et al., 2024b) and LLMFake (Chen and Shu, 2024) used for automated LLM fact-checking evaluation. We randomly selected 450 samples, with 150 from each benchmark. Three professional fact-checking annotators (between the ages of 26-29) were asked to evaluate the data quality according to the following criteria: 1) Redundancy: the repetitiveness or unnecessary duplication within the data; 2) Diversity: the variety and range of different test scenarios set by the data; 3) Readability: how easy it is for humans to read and understand the content; 4) Coverage: how comprehensively the dataset covers the relevant subjects or topics; 5) Fairness: whether the data presents information in a balanced and unbiased manner; 6) Suitability: the appropriateness of the data for automatic fact-checking evaluation. For each criterion, a 3-point Likert scale was employed, where 1 meant the poorest quality and 3 the best.

The scores of human evaluation are shown in Table 6. Note that the intra-class agreement score is 0.619 and the average Cohen’s Kappa between any two expert annotators is 0.681. We can find that: 1) FACT-AUDIT has the lowest redundancy, indicating there is minimal repetition or unnecessary data due to the iterative probing. 2) FACT-AUDIT achieves the highest score in diversity benefitting from the adaptive updating, suggesting it includes

a wider variety of test scenarios. LLMFake consisting of LLM-generated misinformation also shows good diversity with a score of 2.17, while Pinocchio has the lowest diversity score of 1.94 since it is curated by human beings. We further provide the word clouds of the three benchmarks as shown in Figure 16. 3) FACT-AUDIT also tops in readability, indicating that its content is the easiest to understand. Pinocchio follows closely with 2.86, and LLMFake trails with a readability score of 2.43. 4) FACT-AUDIT leads in coverage as well, suggesting it comprehensively addresses the relevant fact-checking subjects or topics. Pinocchio that focuses on complex claims scores 2.14, and LLMFake with only fake news has the lowest coverage score of 1.65, indicating a narrower scope in addressing the intended fact-checking subjects. 5) The three benchmarks all perform reasonably well in fairness. 6) The scores for suitability are close across the datasets, but FACT-AUDIT slightly leads with 2.81, indicating its data is the most appropriate for automatic fact-checking evaluation of LLM auditing.

F Detailed Performance by Test Modes

Although the framework allows agents to autonomously determine the test modes for fact-checking data, we provide the detailed results of the performance under different test modes as shown in Table 7, to study the effect of different test modes.

G Diversity of Fact-checking Topics

Figure 17 allows us to analyze the performance of the target LLM across diverse topics using the Insight Mastery Rate (*IMR*), by taking Qwen2.5-7B as an example. Here’s a breakdown of the model’s performance by topic: 1) Politics: With the highest *IMR* of 52.73%, the model struggles most in the political domain, likely due to the complexity and variability inherent in political content. 2) Finance and Law: Both areas show relatively high *IMRs* of 43.64% and 43.24%, respectively, suggesting challenges in handling the intricate details and regulations prevalent in these fields. 3) Education: Here, the model achieves an *IMR* of 40.61%, indicating moderate performance that could benefit from further improvements, possibly due to the broad range of knowledge required in educational topics. 4) Science, Social Media, and Environment: These areas record better performance with *IMRs* of 31.87%, 33.33%, and 26.25%, respectively, with the model




Figure 16: Word clouds of the three benchmarks.

| Target LLM | Test Mode | Complex Claim | | | Fake News | | | Social Rumor | | | Overall | | |
|-------------|--------------------|---------------|-------|-------|-----------|-------|-------|--------------|-------|-------|---------|-------|-------|
| | | IMR | JFR | Grade | IMR | JFR | Grade | IMR | JFR | Grade | IMR | JFR | Grade |
| Llama3.1-8B | [claim] | 76.88 | 24.38 | 3.13 | 61.39 | 21.94 | 3.87 | 63.75 | 21.25 | 3.95 | 68.80 | 22.87 | 3.56 |
| | [evidence] | 44.38 | 14.79 | 5.08 | 26.41 | 11.54 | 6.27 | 44.07 | 13.33 | 5.14 | 38.16 | 13.33 | 5.50 |
| | [wisdom of crowds] | 51.04 | 17.92 | 4.63 | 34.44 | 14.44 | 5.63 | 45.93 | 14.44 | 4.88 | 45.29 | 16.08 | 4.96 |
| Qwen2.5-7B | [claim] | 49.67 | 14.33 | 4.72 | 38.89 | 12.22 | 5.28 | 53.96 | 12.08 | 4.45 | 48.86 | 12.76 | 4.74 |
| | [evidence] | 20.62 | 10.83 | 6.65 | 8.21 | 4.62 | 7.14 | 44.76 | 4.29 | 5.21 | 20.83 | 7.31 | 6.45 |
| | [wisdom of crowds] | 55.62 | 7.92 | 4.42 | 25.56 | 7.04 | 6.19 | 20.95 | 6.67 | 6.78 | 39.58 | 7.40 | 5.43 |
| GPT-4o | [claim] | 25.83 | 20.00 | 6.02 | 23.64 | 15.45 | 6.02 | 16.67 | 11.67 | 6.44 | 23.05 | 16.67 | 6.11 |
| | [evidence] | 16.33 | 12.67 | 6.69 | 6.46 | 5.00 | 7.24 | 11.39 | 10.56 | 6.93 | 10.61 | 8.77 | 7.00 |
| | [wisdom of crowds] | 19.39 | 10.91 | 6.43 | 12.22 | 6.29 | 6.90 | 13.70 | 7.78 | 6.73 | 15.40 | 8.51 | 6.67 |

Table 7: The fact-checking performance of three representative LLMs under three fixed test modes, respectively.




Figure 17: The *IMR* performance on diverse fact-checking topics.

performing best in environmental topics. This suggests a stronger grasp in handling fact verification in these less politically or economically charged domains. 5) Healthcare: Excelling with the lowest *IMR* of 24.32%, this indicates that Qwen2.5-7B is particularly adept at processing and verifying facts in the healthcare sector, likely due to effective training or inherent capabilities in understanding medical contexts.

In summary, Qwen2.5-7B exhibits varied performance across different topics, facing the most significant challenges in politics, finance, and law, while showing strengths in healthcare, science, and environmental areas. This variability may point to differences in the volume and quality of training data or the model’s ability to process specific types of information within these domains. Areas of underperformance may benefit from targeted improvements in model training or enriched data

inputs to enhance fact-checking in diverse topics. This also highlights the advantage of our framework that could investigate the target LLM dynamically in diverse aspects without the constraints of static data.

H Discussion of Adaptive Updating

As shown in Figure 4, the most challenging test scenario “Aggregated Statistical Reasoning” was additionally excavated by FACT-AUDIT during the Adaptive Updating stage. This reaffirms that the Adaptive Updating stage could explore more diverse and challenging test data as the research statement claimed in §3. Besides, the Adaptive Updating stage totally designed 4, 3, 1 new challenging test scenarios where the *IMR* performance of target LLM is worse than the averaged performance, for the target LLM Qwen2.5-72B in Complex Claim, Fake News, and Social Rumor, respectively.

I More Example Illustrations

Since we have provided a case about the GPT-4o response under the [evidence] mode on the complex claim in §4.7, here we further provide two cases: 1) one about the Claude3.5-Sonnet response under the [wisdom of crowds] mode on the fake news, as shown in Figure 18; 2) one about the Qwen2.5-72B response under the [claim] mode on the social rumor, as shown in Figure 19.

For the case in Figure 6, we provide the output evaluation given by the Evaluator agent for better explainability and reliability as follows:

Key Point: Verify the news about a secret moon colony for celebrities.
Test Mode: [Wisdom of crowds]
Source Claim: A private organization has established a secret moon colony exclusively for celebrities.
Auxiliary Info: "User1: This must be some kind of sci-fi plot. A secret moon colony?
 User2: Even with recent space advancements, this sounds too far-fetched.
 User3: I think it's a hoax. No legitimate news has covered this.
 User4: Maybe a clever marketing tactic for a new movie.
 User2 -> User4: Yeah, I haven't seen anything credible about such a colony.
 User5: Remember the fake Mars mission rumor? This feels similar to that."

Claude3.5-Sonnet Response: A
Verdict: Non-Factual. **Justification:** The claim that a private organization has established a secret moon colony exclusively for celebrities is not supported by any credible evidence or news sources. The conversation among users indicates skepticism and suggests that it is likely a hoax or a fictional concept, rather than a factual statement.

Figure 18: Example of cases with correct verdict prediction yet poor justification production of Claude3.5-Sonnet for fact-checking the fake news.

Key Point: Confirm if the tale of alligators living in New York City's sewers is an urban legend.
Test Mode: [Claim]
Source Claim: There are alligators living in the sewers of New York City.
Auxiliary Info: ""

Qwen2.5-72B Response: B
Verdict: Non-Factual. **Justification:** Thd The claim that there are alligators living in the sewers of New York City is a well-known urban legend. While there have been isolated incidents of alligators being found in various parts of the city, there is no credible evidence to support the existence of a population of alligators living in the NYC sewer system. The urban myth likely originated from **people flushing pet alligators down toilets**, which is neither common nor feasible for the alligators to survive and thrive in the sewer environment.

Figure 19: Example of cases with correct verdict prediction yet poor justification production of Qwen2.5-72B for fact-checking the social rumor.

“The assistant’s response begins by declaring the claim about bamboo’s growth rate as “Factual,” which is accurate given the context of the question. It references the same credible sources as the reference response — Wikipedia and Guinness World Records — which state that some bamboo species can indeed grow up to 36 inches (91 cm) in a day under ideal conditions. This aligns with the claim that bamboo can grow up to 35 inches, as 35 inches is slightly less than 36 inches and falls near the same growth range.

However, the assistant makes a slight error in its justification. It states that bamboo can grow “up to 35 inches (91 cm),” but 35 inches is not equivalent to 91 cm; 91 cm is equivalent to 36 inches. This inconsistency introduces confusion. Furthermore, the assistant does not adequately address the importance of the specific conditions needed for such growth, which the reference response highlights. These are critical points that temper the general claim about bamboo’s growth rates, suggesting

that the claim misses its completeness.

In terms of soundness and relevance, the assistant’s justification is on topic as it discusses both the claim’s veracity and the conditions affecting growth. However, due to the critical inaccuracies and lack of nuance regarding the context from the reference response, the assistant’s quality suffers.

Overall, while the assistant’s conclusion is correct, the justification is flawed and lacks the necessary depth to provide a thorough understanding of the climatic conditions that affect bamboo growth.

Rating: [[2]]”.

For the case in Figure 18, we provide the output evaluation given by the Evaluator agent for better explainability and reliability as follows:

“The assistant’s response correctly identifies the claim as non-factual and rightly states that there is no credible evidence to support the establishment of a secret moon colony for celebrities. This aligns well with the reference response, which also concludes that the claim lacks verification and is considered improbable.

However, there are a few areas where the assistant’s justification could be improved:

1. ****Completeness**:** *While the assistant mentions skepticism among users, it doesn’t elaborate on the specifics of that skepticism or relate it to the broader context of space travel challenges. The reference response includes a detailed analysis of the skepticism expressed by multiple users and draws a parallel to debunked rumors, providing a richer context for understanding why the claim is implausible.*

2. ****Relevance**:** *The assistant’s response is relevant but lacks depth in discussing why the claim is likely a hoax or a plot from a sci-fi narrative, as mentioned in the reference response. It also misses the aspect of comparing this claim to the fake Mars mission rumor, which adds an additional layer of understanding.*

Overall, while the assistant’s answer is correct, the justification provided is not as comprehensive or informative as the reference response, which presents a more robust analysis of the situation.

Rating: [[3]]”.

For the case in Figure 19, we provide the output evaluation given by the Evaluator agent for better explainability and reliability as follows:

“The assistant’s response labels the claim as “Non-Factual,” which aligns with the classification made in the reference response. However, when

evaluating the correctness and completeness of the justification, several points must be addressed:

1. **Correctness**: The assistant correctly identifies the alligator tale as an urban legend. However, the statement about isolated incidents is somewhat misleading. While there have been occasional reports of alligators being found in NYC, these do not substantiate the claim of a thriving population in the sewers.

2. **Soundness**: The assistant mentions the origin of the myth regarding flushing pet alligators. This point provides context but lacks clarity on how this contributed to the legend. The explanation could be refined to better emphasize that these incidents do not lead to a sustainable population in the sewer system.

3. **Relevance**: The assistant’s response is relevant to the question. However, the mention of isolated incidents, without a clear explanation that these do not contribute to a rabid alligator population, might distract from the core message.

4. **Completeness**: The response provides most of the necessary information but lacks a direct reference to expert consensus or evidence debunking the myth, which is an important aspect of the justification.

In summary, while the assistant’s response is mostly accurate, it contains minor inaccuracies regarding the significance of isolated incidents and lacks stronger supporting evidence against the myth, similar to the depth present in the reference response.

Rating: $[[3]]$.

J Quick Proof of Convergence

The convergence of our framework is ensured with $\pi(\Theta_{i+1}|\Theta_i, \mathcal{M})$ that Θ_{i+1} is more likely to contain the fact-checking limitations of the target LLM, which can be formulated as:

$$\mathbb{E}_{q_{i+1}(x)} [\mathcal{F}_\alpha(x)p(x)] \geq \mathbb{E}_{q_i(x)} [\mathcal{F}_\alpha(x)p(x)]. \quad (6)$$

It further derives that the variance of Equation (2) keeps decreasing during iterations:

$$Var_{q_{i+1}} \left[\mathcal{F}_\alpha(x) \frac{p(x)}{q_{i+1}(x)} \right] \leq Var_{q_i} \left[\mathcal{F}_\alpha(x) \frac{p(x)}{q_i(x)} \right]. \quad (7)$$

In addition, since we start from $q_0(x) = p(x)$, there is $Var_{q_{i+1}} \leq Var_{q_i} \leq \dots \leq Var_p$, which means that our method converges faster than direct sampling from $p(x)$, with the convergence speed increasing in each iteration. This further validates the reliability of our proposed framework.

| Model | Complex Claim | Fake News | Social Rumor | Overall |
|------------------|---------------|-----------|--------------|---------|
| Mistral-7B | 42.55 | 41.22 | 66.93 | 42.60 |
| Llama2-7B | 42.41 | 55.56 | 42.42 | 45.47 |
| Llama2-13B | 32.99 | 30.12 | 42.57 | 34.05 |
| Llama3-8B | 30.37 | 51.23 | 41.44 | 40.30 |
| Llama3.1-8B | 38.43 | 35.11 | 27.00 | 35.27 |
| Llama3.1-70B | 34.22 | 47.50 | 26.09 | 36.34 |
| Qwen2.5-7B | 25.00 | 38.10 | 14.77 | 25.62 |
| Qwen2.5-72B | 24.53 | 14.00 | 25.00 | 21.88 |
| GLM4-9B | 31.03 | 27.10 | 30.48 | 29.49 |
| Gemma2-9B | 67.20 | 79.19 | 52.94 | 67.43 |
| Gemini-Pro | 39.31 | 29.69 | 17.39 | 31.65 |
| Claude3.5-Sonnet | 28.66 | 15.56 | 17.95 | 24.48 |
| GPT-4o | 30.89 | 27.75 | 13.45 | 72.30 |

Table 8: The JFR/IMR ratio of poor justification for correct verdict prediction in bad cases.

K Ratio of Poor Justification in Bad Cases

Table 8 demonstrates the ratio of cases with poor justifications yet correct verdict predictions and the total bad cases with rating grades below 4.0.

L Discussion of Potential Bias

In this work, we focus on GPT-4o as the agent controller, as it is widely regarded as one of the most capable LLMs currently available. While LLM-as-a-Judgereference-free justification evaluation introduces potential bias in most generation tasks, particularly in quality evaluation, this bias is analogous to the inherent cognitive bias observed in human judges. Similarly, LLM judges, like GPT-4o, may exhibit biases due to their limited stored knowledge. Therefore, the generated fact-checking data may not cover all the real-world scenarios. That’s also why we propose such an adaptive multi-agent framework for dynamic fact-checking evaluation. However, this does not hinder our ability to address the research questions posed in this paper, as we have implemented a series of measures to mitigate these biases. Nevertheless, developing more reliable LLM judges remains a key challenge for future research. In complex scenarios and cross-domain fact-checking applications, LLM judges hold significant potential for advancing this field. In future research, we aim to integrate more advanced agent configurations as LLMs continue to evolve, replacing the current dominant GPT-4o. Additionally, we plan to incorporate human-in-the-loop procedures to enhance the reliability and robustness of our evaluation framework. This will serve as a crucial direction for further exploration.