

# InFi-Check: Interpretable and Fine-Grained Fact-Checking of LLMs

Anonymous ACL submission

## Abstract

Large language models (LLMs) often hallucinate, yet most existing fact-checking methods treat factuality evaluation as a binary classification problem, offering limited interpretability and failing to capture fine-grained error types. In this paper, we introduce InFi-Check, a framework for interpretable and fine-grained fact-checking of LLM outputs. Specifically, we first propose a controlled data synthesis pipeline that generates high-quality data featuring explicit evidence, fine-grained error type labels, justifications, and corrections. Based on this, we further construct large-scale training data and a manually verified benchmark InFi-Check-FG for fine-grained fact-checking of LLM outputs. Building on these high-quality training data, we further propose InFi-Checker, which can jointly provide supporting evidence, classify fine-grained error types, and produce justifications along with corrections. Experiments show that InFi-Checker achieves state-of-the-art performance on InFi-Check-FG benchmark and strong generalization across various downstream tasks, significantly improving the utility and trustworthiness of factuality evaluation.

## 1 Introduction

Recent breakthroughs in Large Language Models (LLMs) have fundamentally transformed the paradigm of human-computer interaction (Achiam et al., 2023; DeepSeek-AI et al., 2025a). However, LLMs are still prone to producing factual errors in their responses, i.e., hallucinations (Ji et al., 2023; Huang et al., 2025; Si et al., 2025b), posing significant risks and severely compromising their trustworthiness. As a result, recent research has focused on the development of factuality evaluation frameworks for LLMs (Tang et al., 2024a), as well as approaches for training LLMs to improve factuality (Tian et al., 2024; Lin et al., 2024; Si et al., 2025a). In these frameworks, fact-checking models (Lei et al., 2025; Seo et al., 2025) play a crucial role in

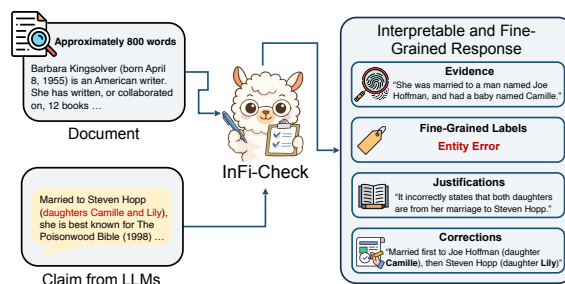


Figure 1: The illustration of our InFi-Check. InFi-Check can simultaneously provide the corresponding evidence, fine-grained labels, justifications, and corrections.

evaluating the factuality of LLM outputs, by checking whether the generated claims can be supported by a reliable knowledge source.

Recent studies focus on how to effectively train a fact-checking model used to evaluate the factuality of LLM-generated responses, including entailment-based (Kryściński et al., 2019; Goyal and Durrett, 2021; Maynez et al., 2020), question-answering-based (Wang et al., 2020; Durmus et al., 2020; Fabbri et al., 2022), atomic-fact-based (Min et al., 2023), and synthetic-data-based (Tang et al., 2024a; Lei et al., 2025; Seo et al., 2025) methods. However, these methods simply treat the fact-checking task as a binary prediction task, classifying the entire response from LLMs as hallucinated or not. Consequently, these methods present several notable limitations: (1) *Lack of Interpretability*: Existing fact-checking models typically output a single predicted label for the entire LLM-generated response (Zha et al., 2023; Tang et al., 2024a), without providing the explicit justification or supporting evidence. This lack of interpretability makes it difficult to analyze model hallucination patterns and limits its practical utility for localizing hallucinations and performing targeted corrections in real-world applications. (2) *Lack of Fine-Grained Fact Checking*: Also, these fact-checking models typically formulate hallucination detection as a binary

071 classification problem, merely predicting whether  
072 a response is factual or hallucinated. However, hal-  
073 lucinations are not all the same and can appear in  
074 different types (Pagnoni et al., 2021a; Mitra et al.,  
075 2025; Zhang et al., 2025), such as incorrect entities,  
076 fabricated facts, or unsupported relations. By col-  
077 lapsing these distinct error types into a single label,  
078 i.e., hallucinated, this oversimplified formulation  
079 fails to provide fine-grained distinctions among hal-  
080 lucinations, thereby limiting detailed error analysis  
081 and specified revision, hindering the development  
082 of targeted strategies for hallucination mitigation.

083 To bridge these gaps, we propose a novel frame-  
084 work called **InFi-Check** to enable interpretable and  
085 fine-grained fact-checking for comprehensive fac-  
086 tuality evaluation of LLM outputs. To achieve this  
087 end, we first introduce an effective data synthesis  
088 pipeline to automatically generate fact-checking  
089 data that contains four key output elements. These  
090 elements are typically provided by professional  
091 fact-checkers in real-world scenarios, including ex-  
092 plicit evidence support, error type identification,  
093 justifications, and corrections. In this way, models  
094 trained on such data can provide fine-grained pre-  
095 dictions and detailed justifications for real-world  
096 users. To ensure the data quality and avoid labeling  
097 errors, our pipeline incorporates the controlled gen-  
098 eration strategies to guide LLMs in generating cor-  
099 responding claims under given fine-grained error  
100 types. Specifically, for the given document-claim  
101 pairs, we first query the advanced LLM (e.g., GPT-  
102 4o (Achiam et al., 2023)) to find grounding sen-  
103 tences as explicit evidence support. Subsequently,  
104 we provide specific fine-grained error types along  
105 with the document-claim pair as input, requiring  
106 the LLMs to first follow our designed structured  
107 reasoning process and generate corresponding jus-  
108 tifications and claims exhibiting the specified error  
109 types. By doing so, we can control the inclusion of  
110 specified fine-grained hallucinations in newly gen-  
111 erated claims from advanced LLMs and ensure that  
112 there are no labeling errors. Also, by comparing the  
113 original claim with the newly synthesized claim,  
114 we can obtain corresponding corrections, enabling  
115 transparent justification of fact-checking decisions.  
116 To ensure the data diversity, our designed pipeline  
117 involves 9 hallucination construction strategies that  
118 cover the diverse yet nuanced hallucination pat-  
119 terns commonly observed in advanced LLMs such  
120 as GPT-4o. Based on our well-designed pipeline  
121 above, we construct a training set as well as a  
122 manually corrected benchmark InFi-Check-FG for

123 evaluating the fine-grained classification capabil-  
124 ities of fact-checking models. By using the train-  
125 ing data from our well-designed data synthesis  
126 pipeline, we introduce InFi-Checker, an advanced  
127 fact-checking model capable of fine-grained hallu-  
128 cination detection and interpretable analysis. The  
129 InFi-Checker can perform comprehensive tasks in-  
130 cluding identifying relevant evidence, detecting  
131 fine-grained error types, and providing justification  
132 and direct corrections.

133 Our extensive experiments show that even the  
134 most advanced LLMs, such as GPT-5 (OpenAI,  
135 2025a), still struggle to capture fine-grained error  
136 types. In contrast, our InFi-Checker achieves state-  
137 of-the-art performance on InFi-Check-FG bench-  
138 mark, and shows strong generalization across di-  
139 verse scenarios, covering question answering (QA),  
140 summarization, and retrieval-augmented genera-  
141 tion (RAG). Different from previous fact-checking  
142 models, InFi-Checker can also offer fine-grained  
143 analysis and the corresponding justifications to im-  
144 prove the interpretability and trustworthiness.

## 2 Related Work 145

146 Many studies have investigated whether LLMs can  
147 generate factually accurate content. These works  
148 broadly categorize into three strands—evaluation,  
149 root cause analysis (Massarelli et al., 2020; Lu  
150 et al., 2022; Luo et al., 2023b; Liu et al., 2023a; Luo  
151 et al., 2023a), and mitigation approaches (Lee et al.,  
152 2022; Dai et al., 2022; Borgeaud et al., 2022; Moi-  
153 seev et al., 2022; Asai et al., 2023; Du et al., 2024;  
154 Wang et al., 2025)—while our research focuses on  
155 the evaluation dimension. Kryściński et al. (2019)  
156 first argued that factuality evaluation in abstrac-  
157 tive summarization should transcend overlap-based  
158 metrics like ROUGE, introducing fact-checking  
159 models to verify whether generated claims are sup-  
160 ported by the source context—a direction also ex-  
161 plored by Maynez et al. (2020) and Goyal and Dur-  
162 rett (2021). Early alternatives employed question-  
163 answering models to check context-summary con-  
164 sistency (Wang et al., 2020; Durmus et al., 2020;  
165 Fabbri et al., 2022), and Zha et al. (2023) and  
166 Ribeiro et al. (2022) later improved performance  
167 via model ensembling and semantic-graph repre-  
168 sentations, respectively. Building on these founda-  
169 tions, researchers have precisely annotated factual  
170 errors in machine-generated claims to assemble  
171 datasets for quantitative factuality evaluation (Fab-  
172 bri et al., 2021; Cao and Wang, 2021; Pagnoni et al.,

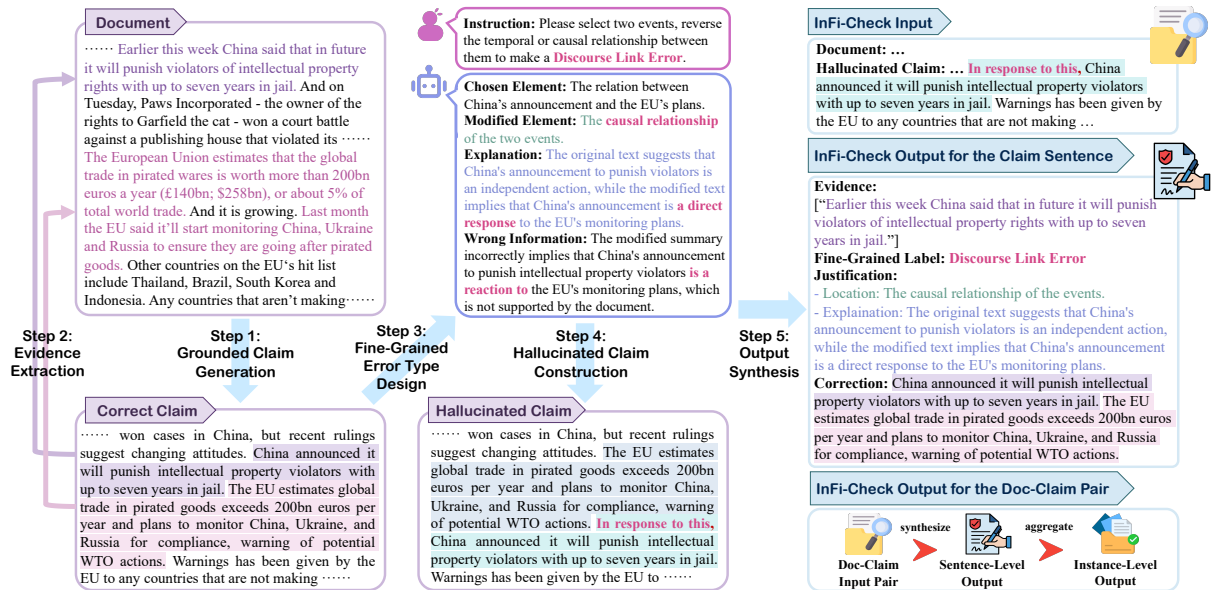


Figure 2: Overview of the InFi-Check pipeline. Some of the text is simplified for better demonstration.

2021a; Zhang et al., 2024).

Recent studies have turned to utilize the power of LLMs to train more capable fact-checking models. For example, MiniCheck (Tang et al., 2024a) uses advanced LLMs to synthesize training data and surpasses prior fact-checking methods. FactCG (Lei et al., 2025) further enhances synthetic data with knowledge graphs to improve the performance of fact-checking models. ClearCheck (Seo et al., 2025) uses synthetic data with multi-task training, allowing the model to perform reasoning before answering. However, despite these advances in fact-checking performance, current models continue to generate only binary predictions, lacking interpretable justifications and fine-grained error labels that can support real-world users, limiting the utility and trustworthiness. Different from these works, our model can jointly provide supporting evidence, classify fine-grained error types, and produce justifications along with corrections, significantly improving the practical utility and trustworthiness of realistic and user-friendly factuality evaluation.

### 3 Methodology

In this section, we introduce InFi-Check, a controlled data synthesis pipeline designed to generate high-quality, interpretable fact-checking data. The core of our InFi-Check lies in systematically constructing grounded claims with fine-grained, realistic hallucinations, alongside their corresponding diagnostic analysis. As illustrated in Figure 2, the pipeline operates in five sequential stages: (1)

Generate Grounded Claim, (2) Extract Supporting Evidence, (3) Design Hallucination with Chain-of-Thought, (4) Construct Hallucinated Claim, and (5) Synthesize Interpretable Output. A key advantage of this pipeline is its scalability and independence from the original document corpus; by varying hallucination injection points and generating new claims, it can be arbitrarily expanded. The prompts used across all stages are provided in Appendix F.

#### 3.1 Controlled Data Synthesis Pipeline

**Stage 1 & 2: Grounded Claim Generation and Evidence Extraction.** We begin with a collection of source documents. Existing datasets often fall short in terms of claim complexity and factual coverage, as the provided claims are not always fully supported by the corresponding documents. To address this limitation, we generate claims directly conditioned on each source document. For every document, a claim is produced, and each sentence in the claim is automatically annotated with its corresponding *grounding sentences* as evidence from the source document. To ensure claim quality, we apply an iterative refinement process based on the extracted grounding sentences, including majority voting and rewriting. We further conduct human evaluation to validate the reliability of this pipeline. Details are provided in the Appendix A.

**Stage 3: Fine-Grained Error Type Design.** This stage is the cornerstone of our method, where we strategically implant controlled hallucinations into the grounded claims. We adopt and extend the fine-

Error Category	Abbr.	Construction Strategy
Predicate Error	PredE	Swap Relation
		Modify Predictions
Entity Error	EntE	Swap Entities
		Compress Phrases
Circumstance Errors	CircE	Swap Circumstances
		Swap Pronouns
Co-Reference Error	CorefE	Merge Sentences
		Reverse Logic
Discourse Link Error	LinkE	
Extrinsic Error	OutE	Add Extrinsic Information

Table 1: Fine-grained error types and corresponding construction strategies in InFi-Check. Detailed descriptions and examples can be found in Table 13.

grained error typology from Pagnoni et al. (2021b), which is applicable to summarization and other grounded generation tasks. We exclude grammatical errors as they pertain to fluency rather than factuality. Our construction system encompasses six major categories, as summarized in our Table 1, each implemented via specific strategies.

To ensure high-quality and diverse errors while enabling interpretable analysis, we leverage LLMs following a structured chain-of-thought (CoT) process. For a given grounded claim and its evidence, we specify a target error type and prompt the model to: (a) analyze the original content, (b) execute the corresponding construction strategy (e.g., swap an entity), and (c) produce a detailed justification for the change. This step-by-step reasoning not only improves the controllability of the constructed errors but also provides a transparent record of how each hallucination was introduced, forming the basis for interpretable outputs in Stage 5.

**Stage 4 & 5: Hallucinated Claim Construction and Output Synthesis.** Using the CoT from Stage 3, we generate the final hallucinated claim sentence containing the specified hallucination. Ultimately, we can synthesize the *interpretable output*: a structured analysis containing the erroneous sentence, its grounding evidence (or lack thereof), the hallucination category, the justification from the CoT, and the corrected sentence. This final output is organized in a sentence-by-sentence manner (see Appendix E for examples), providing a clear, traceable reasoning path from error construction to diagnosis. The structured format offers two key benefits: (1) *Interpretability through traceable reasoning*, making model decisions transparent and easing

manual verification by presenting grounding contexts and corrections side-by-side; and (2) *Higher quality from controlled generation*, as the CoT process and explicit output schema enforce rigorous, context-aware error construction and analysis. This output serves as the gold standard for training and evaluating fact-checking models.

### 3.2 Fine-Grained Error Type Construction

Building upon the error typology outlined in Stage 3, this subsection details the specific strategies employed to construct each category of hallucination. Our methodology adapts and extends prior work on error construction (Kryściński et al., 2019; Cao and Wang, 2021; Pagnoni et al., 2021b), introducing novel strategies to capture subtle and challenging error types often overlooked in existing datasets.

**Semantic-Level Hallucinations.** These errors involve inaccuracies within a single proposition, including Predicate Errors (PredE), Entity Errors (EntE), and Circumstance Errors (CircE).

- *Element Swapping*: For straightforward errors, we adapt established swapping techniques (Kryściński et al., 2019; Cao and Wang, 2021). We prompt an LLM to identify a target element (e.g., an **entity**, **predicate**, or **circumstantial phrase**) within a factually correct claim sentence and replace it with a semantically congruent but incorrect alternative from the source document or a generated list.
- *Modifying Predictions*: This strategy targets the confusion between speculative language (e.g., "predicted," "might") and factual assertions. We guide the LLM to identify sentences containing modal verbs or speculative phrases and alter them to state the speculated outcome as a fact (e.g., changing "The company *forecast* growth" to "The company *grew*").
- *Compressing Phrases*: This strategy generates errors through oversimplification, where specific technical terms or nuanced descriptions are replaced with overly broad terms, distorting meaning (e.g., compressing "net revenue attributable to the parent company" to "net revenue"). A two-stage LLM verification process ensures the compression introduces a factual distortion rather than mere paraphrasing.

**Discourse-Level Hallucinations.** These errors span multiple sentences, disrupting coherence and reference. They include Co-reference Errors




Figure 3: Document and claim length (words) distribution of InFi-Check-FG with average length comparison.

(CorefE) and Discourse Link Errors (LinkE).

- *Swapping Pronouns*: We extend the method of FactCC (Kryściński et al., 2019) by swapping pronouns of all types (not only gendered ones). To increase the complexity of the data, we first convert named entities into pronouns before performing the swap, thereby introducing controlled referential ambiguity.
- *Merging Sentences*: To simulate conflation errors, we choose two sentences about similar but different topics or events. Then, we combine them into one sentence and wrongly give actions or facts from both to just one subject.
- *Reverse Logical Relationship*: We prompt the LLM to identify a pair of events in the source document with a clear temporal or causal relationship. We then ask the model to reverse this relationship (for example, switch the cause and effect) and rewrite the claim to show this false link. This creates a sentence that sounds believable but is not factually correct.

**Extrinsic Hallucinations.** Extrinsic Errors (OutE), or “out-of-article” errors, happen when information external to the source is added. Due to the difficulty of making sure all context is removed from an old claim, we ask the LLM to add a believable but completely unsupported piece of information (like a new number, event, or name) to a claim, making sure the new content is clearly extrinsic.

#### 4 InFi-Check Dataset and Model

Building upon the controlled synthesis pipeline detailed in the Section 3, we construct the interpretable and fine-grained fact-checking training data, namely InFi-Check-TR without human effort, and further propose a manually verified benchmark

InFi-Check-FG. Subsequently, we train our model InFi-Checker based on InFi-Check-TR. This section shows the details of these core resources.

#### 4.1 Dataset and Benchmark

InFi-Check serves a dual purpose: it automatically creates a large-scale training corpus InFi-Check-TR, and also a high-quality and challenging benchmark InFi-Check-FG for evaluating fine-grained classification capabilities of fact-checking models. **Construction Process.** To ensure factual grounding and diversity, our dataset construction begins with source documents rather than existing summaries. We randomly sample news and encyclopedic documents of varying lengths from the BBC News (Gupta et al., 2022) and DetNet Wikipedia (Xu and Lapata, 2019) datasets. For each document, we generate a grounded claim using multiple LLMs via iterative rewriting to mitigate bias and ensure quality, with each claim sentence linked to its evidential grounding sentences in the source (see Appendix A for details). Subsequently, for each document-claim pair, we apply the error construction strategies from InFi-Check to generate one hallucinated claim for each strategy, while retaining the original claim as a “No Error” sample, resulting in 10 total samples per pair.

**InFi-Check-TR Details.** Our InFi-Check-TR comprises 15,660 samples for training, with more statistical details in Appendix B. To verify the quality of InFi-Check-TR, we conducted a human evaluation where expert annotators evaluated 100 randomly sampled instances across four dimensions: claims, evidence sentences, justifications, and hallucination validity. Results in Table 11 show a 95% agreement rate on hallucination validity, demonstrating the capability of the InFi-Check pipeline.

**InFi-Check-FG Details.** To further ensure evaluation reliability, we construct InFi-Check-FG based on manual check and from a distinct document split, which is a high-quality benchmark that consists of 519 manually verified samples. As illustrated in Figure 3, claims in InFi-Check-FG exhibit longer length compared to existing benchmarks, increasing complexity by preserving richer source-document details, thereby introducing more subtle and challenging hallucinations for evaluation. More details are provided in Appendix A.

#### 4.2 Fact-Checking Model: InFi-Checker

Finally, we train our model InFi-Checker on the synthetic training data using supervised fine-tuning

Model	PredE	EntE	CircE	CorefE	LinkE	OutE	NoE	BAcc
<b>The Open-Source Models</b>								
Llama-3.1-8B-Instruct	18.82	9.52	31.03	3.13	0.00	0.00	61.11	17.66
Llama-3.1-8B-Instruct*	32.94	4.76	13.79	26.56	0.00	0.00	88.49	23.79
Qwen3-8B	76.47	52.38	44.83	50.00	0.00	50.00	61.90	47.94
Qwen3-8B*	76.47	52.38	24.14	70.31	0.00	50.00	80.16	50.49
<b>The State-of-the-Art LLMs</b>								
Claude-3.7-Sonnet	70.59	66.67	34.48	37.50	26.67	81.58	53.97	53.07
Claude-3.7-Sonnet*	68.24	71.43	34.48	21.88	23.33	78.95	46.03	49.19
DeepSeek-V3.2-NoThink	78.82	76.19	58.62	9.38	26.67	86.84	80.95	59.64
DeepSeek-V3.2-NoThink*	75.29	71.43	55.17	15.62	13.33	89.47	<b>94.05</b>	59.19
GPT-4o	80.00	76.19	44.83	50.00	3.33	84.21	84.92	60.50
GPT-4o*	84.71	61.90	44.83	75.00	0.00	94.74	91.67	64.69
GPT-4.1	85.88	76.19	37.93	42.19	0.00	63.16	75.79	54.45
GPT-4.1*	85.88	76.19	41.38	56.25	3.33	68.42	80.95	58.91
GPT-5	63.53	57.14	62.07	42.19	16.67	60.53	41.27	49.06
GPT-5*	63.53	61.9	51.72	45.31	13.33	47.37	41.27	46.35
o3	80.00	61.90	58.62	43.75	0.00	81.58	80.16	58.00
o3*	81.18	57.14	58.62	56.25	3.33	81.58	87.30	60.77
<b>InFi-Checker (Ours)</b>								
<b>InFi-Checker-Llama</b>	<b>95.29</b>	90.48	79.31	95.31	86.67	<b>100.00</b>	89.29	90.91
<b>InFi-Checker-Qwen</b>	93.51	<b>91.67</b>	<b>87.50</b>	<b>100.00</b>	<b>88.89</b>	96.77	88.01	<b>92.34</b>

Table 2: Results (%) on InFi-Check-FG. We display the accuracy of each error type, as well as the balanced accuracy (BAcc), which calculates the average accuracy of all error types. Best performances are marked in bold. For baselines, models marked with \* are tested in a one-shot setting, while the others are tested in a zero-shot setting.

Model	0-shot	1-shot
<b>The Open-Source Models</b>		
Llama-3.1-8B-Instruct	66.13	71.49
Qwen3-8B	65.31	70.16
<b>The State-of-the-Art LLMs</b>		
Claude-3.7-Sonnet	65.89	63.42
DeepSeek-V3.2-NoThink	66.99	77.14
GPT-4o	74.37	74.75
GPT-4.1	75.19	75.19
GPT-5	76.53	75.70
o3	75.54	74.57
<b>InFi-Checker (Ours)</b>		
<b>InFi-Checker-Llama</b>		77.17
<b>InFi-Checker-Qwen</b>		<b>77.20</b>

Table 3: Fine-grained results on FRANK, with balanced accuracy calculated in a binary mapped setting due to error types overlapping (See Appendix C for details).

(SFT). The model is trained to perform a comprehensive, structured output task: given a document and a claim, it identifies grounding evidence, classifies the fine-grained error type, and provides a natural language justification alongside a direct correction. This end-to-end training regime, powered by the rich annotations, enables InFi-Checker to deliver interpretable, fine-grained fact-checking for LLM outputs, moving beyond binary classification.

## 5 Experiment

### 5.1 Experimental Setup

**Implementation of InFi-Checker.** We implement Llama-3.1-8B-Instruct and Qwen3-8B as backbones for InFi-Checker, using InFi-Check-TR as

training dataset and conducted supervised fine-tuning (details in Appendix C.3).

**Other Benchmarks.** To evaluate InFi-Checker’s out-of-distribution generalization, we conduct experiments across existing multiple factuality evaluation benchmarks, including (1) *fine-grained hallucination labeling benchmark*, where we adapt the **FRANK** benchmark, which is a human-annotated benchmark that uses the same fine-grained hallucination label taxonomy as InFi-Check-FG. (2) *commonly used binary hallucination labeling benchmark*, where we select benchmarks of diverse text sources and generation tasks: **ClaimVerify** (Liu et al., 2023b) and **Factcheck-Bench** (Wang et al., 2024) for search queries and responses, **ExpertQA** (Malaviya et al., 2024) and **REVEAL** (Jacovi et al., 2024) for QA, as well as **MediaSum** and **MeetingBank** for dialogue summarization (Tang et al., 2024b). We test in fine-grained settings for FRANK, and binary settings for other benchmarks, as they do not have fine-grained labeling. Appendix C provides more implementation details.

**Baselines.** We validate InFi-Checker through an extensive comparison with two separate groups of competitive baselines. On InFi-Check-FG and FRANK, which require fine-grained error type labeling, we employ **State-of-the-Art LLMs** including: (1) the backbones of InFi-Checker, **open-source models**: Llama-3.1-8B-Instruct (Kassianik et al., 2025) and Qwen3-8B (Yang et al., 2025); and (2) **closed-source models**: GPT-4o (Jaech

Model	Claim Verify	Expert QA	Factcheck-Bench	REVEAL	MediaSum	MeetingBank	Average
<b>The Open-Source Models</b>							
Llama-3.1-8B-Instruct	63.6	49.8	69.8	78.2	50.8	62.3	62.4
Qwen3-8B	<b>66.1</b>	<b>55.6</b>	<b>78.4</b>	<b>83.2</b>	<b>77.7</b>	<b>74.2</b>	<b>72.5</b>
<b>The State-of-the-Art LLMs</b>							
Claude-3.7-Sonnet	<b>83.7</b>	74.4	86.9	88.0	<b>85.4</b>	84.0	83.7
DeepSeek-V3.2-NoThink	75.4	74.4	87.9	91.0	65.5	82.9	79.5
GPT-4o	78.3	68.3	86.0	86.9	71.5	76.9	78.0
GPT-4.1	81.6	<b>80.3</b>	<b>91.3</b>	93.2	75.9	<b>86.3</b>	84.8
GPT-5	87.7	75.9	90.4	<b>93.7</b>	80.2	87.6	<b>85.9</b>
o3	83.3	79.6	86.9	92.2	82.9	83.8	84.8
<b>Specialized Fact-Checking Models</b>							
ClearCheck (COT)	85.4	72.7	87.9	87.0	67.8	75.8	79.4
AlignScore-large	79.8	75.0	83.7	92.2	75.8	76.5	80.5
FactCG	76.2	75.3	<b>89.0</b>	90.0	79.1	71.9	80.3
MiniCheck	85.6	72.9	86.8	<b>91.0</b>	74.3	77.8	81.4
<b>InFi-Checker-Llama</b>	75.9	<b>78.3</b>	83.7	87.7	73.5	65.8	77.5
<b>InFi-Checker-Qwen</b>	<b>89.6</b>	75.7	88.0	90.0	<b>80.4</b>	<b>78.5</b>	<b>83.7</b>

Table 4: Macro-F1 (%) on six binary fact-checking benchmarks. Of note, **Bold** and underline highlight the best Macro-F1 within each group of baselines and the best overall Macro-F1, respectively.

et al., 2024), GPT-4.1 (OpenAI, 2025b), GPT-5 (OpenAI, 2025a), o3 (OpenAI, 2025c), Claude-3.7-Sonnet (Anthropic, 2025) and DeepSeek-V3.2 (DeepSeek-AI et al., 2025b). Baselines are tested following the InFi-Check reasoning format, and intentionally limit the number of demonstration examples to zero/one-shot to mitigate performance degradation from excessive prompt length (see Appendix F). Additionally, for binary fact-checking benchmarks, we include **specialized models** specifically optimized for binary factuality evaluation: **ClearCheck** (Seo et al., 2025), that leverages multi-task training for robustness, **AlignScore-large** (Zha et al., 2023), a holistic metric using a unified alignment function, **FactCG** (Lei et al., 2025), which enhances training data via knowledge graphs, and **MiniCheck** (Tang et al., 2024a), the state-of-the-art binary evaluator that utilizes a novel document-claim pair synthesis method for training data.

## 5.2 InFi-Check-FG Results

The performance of various models on InFi-Check-FG is detailed in Table 2; from these results, we draw the following critical observations:

**InFi-Checker consistently outperforms all baselines on InFi-Check-FG, especially on complex data.** InFi-Checker achieves a substantial 27.65% improvement in balanced accuracy over the second-best performer, one-shot GPT-4o. This advantage is most pronounced in discourse-level errors (*CorefE* and *LinkE*), where even leading closed-source models often fail to identify any instances correctly. This performance gap underscores two key strengths of InFi-Checker: (1) Cross-frame analysis capability: Our novel hallucination con-

struction strategies of discourse-level errors enhance the model’s ability to handle hallucinations spanning across semantic frames; and (2) Structured reasoning: Our synthetic justification design effectively addresses the limitations of standard LLMs in complex hallucination analysis.

**InFi-Check provides a significant performance boost to the backbones.** By comparing InFi-Checker (InFiChecker-Llama and InFiChecker-Qwen) with their prompted backbone counterparts (Llama-3.1-8B-Instruct\* and Qwen3-8B\*), we observe that the gains are not solely due to instruction-guided interpretable reasoning. Instead, the improvement primarily stems from the model’s ability to internalize the nature of all types of hallucinations from our curated dataset, enabling robust and accurate fact-checking across different backbones.

## 5.3 OOD Generalization Results

We further evaluate the transferability of InFi-Checker across various benchmarks (Table 3 and Table 4), leading to the following conclusions:

**The fine-grained fact-checking capabilities of InFi-Checker is generalizable.** On the FRANK benchmark (Table 3), which shares the same error-type taxonomy but is derived from a different corpus, InFi-Checker consistently outperforms closed-source models and shows substantial improvements over its backbones. These results suggest that the fine-grained hallucination detection capabilities of InFi-Checker transcend simple pattern memorization. Instead, InFi-Checker captures transferable and fundamental diagnostic features, therefore generalizing effectively across out-of-distribution data. **Fine-grained training enhances binary fact-**

Model	BAcc(Nor.)	BAcc(Str.)	SAR(Avg.)
<b>The Open-Source Models</b>			
Llama-3.1-8B-Instruct†	27.69	21.33	72.57
Llama-3.1-8B-Instruct*†	23.79	18.12	79.41
Qwen3-8B†	47.94	41.08	86.50
Qwen3-8B*†	50.49	44.78	87.22
<b>The State-of-the-Art LLMs</b>			
Claude-3.7-Sonnet*	49.19	45.12	94.28
DeepSeek-V3.2-NoThink*	59.19	52.31	87.33
GPT-4o*	64.69	60.27	80.19
GPT-4.1*	58.91	55.13	94.87
GPT-5*	46.35	42.28	92.58
o3*†	60.77	56.76	94.60
<b>InFi-Checker (Ours)</b>			
InFi-Checker-Llama	90.91	85.49	94.02
<b>InFi-Checker-Qwen</b>	92.34	87.74	94.92

Table 5: Normal(Nor.) and strict(Str.) balanced accuracy (%) and averaged SAR on InFi-Check-FG. The reported SAR(Avg.) is the average of error-type specific SAR across all evaluable error types. Models marked with † contain types with zero accuracy (excluded from the average), which may lead to an optimistic estimation of their SAR. \* means one-shot settings. Full results are displayed in Appendix D.

**checking.** Although InFi-Checker is optimized for fine-grained tasks, it achieves competitive results on six binary benchmarks (Table 4). These datasets cover a broad spectrum of document sources and various downstream tasks, including summarization, QA, and search-based generation, which comprehensively validates the model’s robustness across varied contexts in binary mode. InFi-Checker surpasses binary classification specialized models and approaches or even exceeds the performance of closed-source models, suggesting that explicitly training on diverse error types heightens the model’s overall sensitivity to hallucinations, thereby benefiting even binary classification.

#### 5.4 In-depth Sentence Level Analysis

To evaluate the models’ interpretability, we assess performance at the sentence level by requiring models to not only categorize the fine-grained error type but also precisely localize it within the text. Specifically, we define **Strict Accuracy**, which demands: (1) correct error-type assignment, (2) precise labeling on both the hallucinated and hallucination-free claim sentences. We then introduce the **Sentence Alignment Ratio (SAR)**, defined as the ratio of strict accuracy to normal accuracy of each error type. While normal metrics only verify labels, SAR quantifies the consistency between a model’s final judgment and its underlying reasoning. A higher SAR indicates authentic comprehen-

Index	Just.+Corr.	Evid.	Sent.	BAcc
JES	✓	✓	✓	90.91
ES	-	✓	✓	75.02
JS	✓	-	✓	69.85
JE	✓	✓	-	61.17
J	✓	-	-	36.95
E	-	✓	-	53.76
S	-	-	✓	30.02
raw	-	-	-	20.59

Table 6: Result (%) for ablation study on InFi-Checker-Llama on InFi-Check-FG. "Just.+Corr." means outputting error justification and correction, "Evid." means outputting evidence sentences, and "Sent." means analyzing sentence-by-sentence. See Appendix D for separate ablation of justification and correction.

sion rather than coincidental guessing. As shown in Table 5, InFi-Checker ranks among the highest SAR scores, which is increasingly difficult as the base normal accuracy rises, therefore confirming InFi-Checker’s performance gain is driven by robust and interpretable reasoning.

#### 5.5 Ablation Study

We conduct an ablation experiment to demonstrate the importance of the key components we claim. We remove the output of interpretable justifications and corrections (**J**), evidence sentence (**E**), as well as the sentence-by-sentence claim analysis (**S**) InFi-Check adapt. Table 6 presents an ablation study on these components. The results stress each component’s importance, for removing any single component results in a notable performance drop.

### 6 Conclusion

In this paper, we propose a new framework termed InFi-Check, which integrates error typologies, synthetic data generation pipelines, and fine-grained annotations for comprehensive fact-checking. To be specific, we first develop diverse methods to synthesize six hallucination error types in grounded generation, including both intrinsic and extrinsic, semantic and discourse level errors. In this way, we construct InFi-Check-TR and InFi-Check-FG, two novel datasets characterized by fine-grained error type design, interpretable justifications and corrections, as well as comprehensive claims. Building upon InFi-Check-TR, we develop InFi-Checker, an advanced factuality evaluation model capable of fine-grained hallucination analysis. Also, we conduct extensive experiments to verify the superiority of InFi-Check-FG and InFi-Checker.

## 580 Limitations

581 Our pipeline’s effectiveness is constrained by in-  
582 herent limitations in LLM capabilities. While we  
583 employ sentence-level verification, the models still  
584 generate document-unsupported hallucinations, or  
585 fail in extracting full grounding evidence sentences.  
586 Additionally, they struggle to differentiate between  
587 factual incompleteness and legitimate information  
588 simplification, particularly affecting error construc-  
589 tion quality for more complex cases. We adapt  
590 manual check and filtering for InFi-Check-FG to  
591 mitigate this bias in evaluation.

## 592 Ethics Statement

593 Our work focuses on improving the fact-checking  
594 abilities of document-grounded generation systems  
595 through interpretable and fine-grained methods.  
596 While synthetic errors and human annotations are  
597 central to our approach, we recognize potential  
598 risks, such as misuse for generating misinforma-  
599 tion or unintended biases in the dataset. To miti-  
600 gate these concerns, we ensure transparency in our  
601 methodologies and emphasize their use for research  
602 purposes only. Additionally, our dataset and model  
603 are designed to generalize across diverse scenarios,  
604 avoiding overfitting to specific benchmarks. We  
605 release all contributions under research-focused li-  
606 censes to encourage responsible and ethical use in  
607 advancing AI systems.

## 608 References

609 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama  
610 Ahmad, Ilge Akkaya, Florencia Leoni Aleman,  
611 Diogo Almeida, Janko Altenschmidt, Sam Altman,  
612 Shyamal Anadkat, et al. 2023. Gpt-4 technical report.  
613 *arXiv preprint arXiv:2303.08774*.

614 Anthropic. 2025. [Claude 3.7 sonnet system card](#). Tech-  
615 nical report, OpenAI.

616 Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and  
617 Hannaneh Hajishirzi. 2023. Self-rag: Learning to  
618 retrieve, generate, and critique through self-reflection.  
619 In *The Twelfth International Conference on Learning*  
620 *Representations*.

621 Sebastian Borgeaud, Arthur Mensch, Jordan Hoff-  
622 mann, Trevor Cai, Eliza Rutherford, Katie Mill-  
623 can, George Bm Van Den Driessche, Jean-Baptiste  
624 Lespiau, Bogdan Damoc, Aidan Clark, Diego  
625 De Las Casas, Aurelia Guy, Jacob Menick, Roman  
626 Ring, Tom Hennigan, Saffron Huang, Loren Mag-  
627 giore, Chris Jones, Albin Cassirer, Andy Brock,  
628 Michela Paganini, Geoffrey Irving, Oriol Vinyals,  
629 Simon Osindero, Karen Simonyan, Jack Rae, Erich

Elsen, and Laurent Sifre. 2022. [Improving language  
models by retrieving from trillions of tokens](#). In  
*Proceedings of the 39th International Conference  
on Machine Learning*, volume 162 of *Proceedings  
of Machine Learning Research*, pages 2206–2240.  
PMLR.

Shuyang Cao and Lu Wang. 2021. [CLIFF: Contrastive  
learning for improving faithfulness and factuality in  
abstractive summarization](#). In *Proceedings of the  
2021 Conference on Empirical Methods in Natural  
Language Processing*, pages 6633–6649, Online and  
Punta Cana, Dominican Republic. Association for  
Computational Linguistics.

Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao  
Chang, and Furu Wei. 2022. [Knowledge neurons in  
pretrained transformers](#). In *Proceedings of the 60th  
Annual Meeting of the Association for Computational  
Linguistics (Volume 1: Long Papers)*, pages 8493–  
8502, Dublin, Ireland. Association for Computational  
Linguistics.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingx-  
uan Wang, Bochao Wu, Chengda Lu, Chenggang  
Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,  
Damai Dai, Daya Guo, Dejian Yang, Deli Chen,  
Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai,  
Fuli Luo, Guangbo Hao, Guanting Chen, Guowei  
Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng  
Wang, Haowei Zhang, Honghui Ding, Huajian Xin,  
Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang,  
Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang,  
Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie  
Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu,  
Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean  
Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao,  
Litong Wang, Liyue Zhang, Meng Li, Miaojuan Wang,  
Mingchuan Zhang, Minghua Zhang, Minghui Tang,  
Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang,  
Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu  
Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge,  
Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin  
Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghai  
Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu,  
Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu  
Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou,  
Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu  
Sun, W. L. Xiao, Wangding Zeng, Wanjia Zhao, Wei  
An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin  
Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu  
Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xi-  
aojin Shen, Xiaokang Chen, Xiaokang Zhang, Xi-  
aoshan Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang  
Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu,  
Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou,  
Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin,  
Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang  
Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang,  
Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yao-  
hui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan  
Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao,  
Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu,  
Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yud-  
uan Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun

692	Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. 2025a. <a href="#">Deepseek-v3 technical report</a> . <i>Preprint</i> , arXiv:2412.19437.		
702	DeepSeek-AI, Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenhao Xu, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Erhang Li, Fangqi Zhou, Fangyun Lin, Fucong Dai, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Hao Li, Haofen Liang, Haoran Wei, Haowei Zhang, Haowen Luo, Haozhe Ji, Honghui Ding, Hongxuan Tang, Huanqi Cao, Huazuo Gao, Hui Qu, Hui Zeng, Jialiang Huang, Jiashi Li, Jiaxin Xu, Jiewen Hu, Jingchang Chen, Jingting Xiang, Jingyang Yuan, Jingyuan Cheng, Jinhua Zhu, Jun Ran, Jinguang Jiang, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kaige Gao, Kang Guan, Kexin Huang, Kexing Zhou, Kezhao Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Wang, Liang Zhao, Liangsheng Yin, Lihua Guo, Lingxiao Luo, Linwang Ma, Litong Wang, Liyue Zhang, M. S. Di, M. Y. Xu, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingxu Zhou, Panpan Huang, Peixin Cong, Peiyi Wang, Qiancheng Wang, Qihao Zhu, Qingyang Li, Qinyu Chen, Qiushi Du, Ruiling Xu, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runqiu Yin, Runxin Xu, Ruomeng Shen, Ruoyu Zhang, S. H. Liu, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaofei Cai, Shaoyuan Chen, Shengding Hu, Shengyu Liu, Shiqiang Hu, Shirong Ma, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, Songyang Zhou, Tao Ni, Tao Yun, Tian Pei, Tian Ye, Tianyuan Yue, Wangding Zeng, Wen Liu, Wenfeng Liang, Wenjie Pang, Wenjing Luo, Wenjun Gao, Wentao Zhang, Xi Gao, Xiangwen Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaokang Zhang, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xingyou Li, Xinyu Yang, Xinyuan Li, Xu Chen, Xuecheng Su, Xuehai Pan, Xuheng Lin, Xuwei Fu, Y. Q. Wang, Yang Zhang, Yanhong Xu, Yanru Ma, Yao Li, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Qian, Yi Yu, Yichao Zhang, Yifan Ding, Yifan Shi, Yiliang Xiong, Ying He, Ying Zhou, Yinmin Zhong, Yishi Piao, Yisong Wang, Yixiao Chen, Yixuan Tan, Yixuan Wei, Yiyang Ma, Yiyuan Liu, Yonglun Yang, Yongqiang Guo, Yongtong Wu, Yu Wu, Yuan Cheng, Yuan Ou, Yuanfan Xu, Yudian Wang, Yue Gong, Yuhan Wu, Yuheng Zou, Yukun Li, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehua Zhao, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhixian Huang, Zhiyu Wu, Zhuoshu Li, Zhuping Zhang,		
693		Zian Xu, Zihao Wang, Zihui Gu, Zijia Zhu, Zilin Li, Zipeng Zhang, Ziwei Xie, Ziyi Gao, Zizheng Pan, Zongqing Yao, Bei Feng, Hui Li, J. L. Cai, Jiaqi Ni, Lei Xu, Meng Li, Ning Tian, R. J. Chen, R. L. Jin, S. S. Li, Shuang Zhou, Tianyu Sun, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xinnan Song, Xinyi Zhou, Y. X. Zhu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, Dongjie Ji, Jian Liang, Jianzhong Guo, Jin Chen, Leyi Xia, Miaojun Wang, Mingming Li, Peng Zhang, Ruyi Chen, Shangmian Sun, Shaoqing Wu, Shengfeng Ye, T. Wang, W. L. Xiao, Wei An, Xianzu Wang, Xiaowen Sun, Xiaoxiang Wang, Ying Tang, Yukun Zha, Zekai Zhang, Zhe Ju, Zhen Zhang, and Zihua Qu. 2025b. <a href="#">Deepseek-v3.2: Pushing the frontier of open large language models</a> . <i>Preprint</i> , arXiv:2512.02556.	755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770 771
703		Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2024. Improving factuality and reasoning in language models through multiagent debate. In <i>Proceedings of the 41st International Conference on Machine Learning, ICML'24</i> . JMLR.org.	772 773 774 775 776 777
704		Esin Durmus, He He, and Mona Diab. 2020. <a href="#">FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization</a> . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 5055–5070, Online. Association for Computational Linguistics.	778 779 780 781 782 783 784
705		Alexander Fabbri, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2022. <a href="#">QAFactEval: Improved QA-based factual consistency evaluation for summarization</a> . In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 2587–2601, Seattle, United States. Association for Computational Linguistics.	785 786 787 788 789 790 791 792
706		Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. <a href="#">SummEval: Re-evaluating summarization evaluation</a> . <i>Transactions of the Association for Computational Linguistics</i> , 9:391–409.	793 794 795 796 797
707		Tanya Goyal and Greg Durrett. 2021. <a href="#">Annotating and modeling fine-grained factuality in summarization</a> . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 1449–1462, Online. Association for Computational Linguistics.	798 799 800 801 802 803 804
708		Anushka Gupta, Diksha Chugh, Anjum, and Rahul Katarya. 2022. Automated news summarization using transformers. In <i>Sustainable advanced computing: select proceedings of ICSAC 2021</i> , pages 249–259. Springer.	805 806 807 808 809
709		Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2025.	810 811 812

813	A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. <i>ACM Transactions on Information Systems</i> , 43(2):1–55.	869
814		870
815		871
816		872
817	Alon Jacovi, Yonatan Bitton, Bernd Bohnet, Jonathan Herzig, Or Honovich, Michael Tseng, Michael Collins, Roei Aharoni, and Mor Geva. 2024. <a href="#">A chain-of-thought is as strong as its weakest link: A benchmark for verifiers of reasoning chains</a> . <i>Preprint</i> , arXiv:2402.00559.	873
818		874
819		875
820		876
821		877
822		878
823	Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. <i>arXiv preprint arXiv:2412.16720</i> .	879
824		880
825		881
826		882
827		883
828	Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. <i>ACM computing surveys</i> , 55(12):1–38.	884
829		885
830		886
831		887
832		888
833	Paul Kassianik, Baturay Saglam, Alexander Chen, Blaine Nelson, Anu Vellore, Massimo Auffero, Fraser Burch, Dhruv Kedia, Avi Zohary, Sajana Weerawardhena, Aman Priyanshu, Adam Swanda, Amy Chang, Hyrum Anderson, Kojin Oshiba, Omar Santos, Yaron Singer, and Amin Karbasi. 2025. <a href="#">Llama-3.1-foundationai-securityllm-base-8b technical report</a> . <i>Preprint</i> , arXiv:2504.21039.	889
834		890
835		891
836		892
837		893
838		894
839		895
840		896
841	Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. <a href="#">Evaluating the factual consistency of abstractive text summarization</a> . <i>Preprint</i> , arXiv:1910.12840.	897
842		898
843		899
844		900
845	Nayeon Lee, Wei Ping, Peng Xu, Mostofa Patwary, Pascale Fung, Mohammad Shoeybi, and Bryan Catanzaro. 2022. Factuality enhanced language models for open-ended text generation. In <i>Proceedings of the 36th International Conference on Neural Information Processing Systems</i> , NIPS '22, Red Hook, NY, USA. Curran Associates Inc.	901
846		902
847		903
848		904
849		905
850		906
851		907
852	Deren Lei, Yaxi Li, Siyao Li, Mengya Hu, Rui Xu, Ken Archer, Mingyu Wang, Emily Ching, and Alex Deng. 2025. <a href="#">FactCG: Enhancing fact checkers with graph-based multi-hop data</a> . In <i>Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 5002–5020, Albuquerque, New Mexico. Association for Computational Linguistics.	908
853		909
854		910
855		911
856		912
857		913
858		914
859		915
860		916
861	Sheng-Chieh Lin, Luyu Gao, Barlas Oguz, Wenhan Xiong, Jimmy Lin, Wen tau Yih, and Xilun Chen. 2024. <a href="#">FLAME: Factuality-aware alignment for large language models</a> . In <i>The Thirty-eighth Annual Conference on Neural Information Processing Systems</i> .	917
862		918
863		919
864		920
865		921
866	Alisa Liu, Zhaofeng Wu, Julian Michael, Alane Suhr, Peter West, Alexander Koller, Swabha Swayamdipta, Noah Smith, and Yejin Choi. 2023a. <a href="#">We’re afraid language models aren’t modeling ambiguity</a> . In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing</i> , pages 790–807, Singapore. Association for Computational Linguistics.	922
867		923
868		924
		925
	Nelson Liu, Tianyi Zhang, and Percy Liang. 2023b. <a href="#">Evaluating verifiability in generative search engines</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 7001–7025, Singapore. Association for Computational Linguistics.	
	Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: multimodal reasoning via thought chains for science question answering. In <i>Proceedings of the 36th International Conference on Neural Information Processing Systems</i> , NIPS '22, Red Hook, NY, USA. Curran Associates Inc.	
	Hongyin Luo, Tianhua Zhang, Yung-Sung Chuang, Yuan Gong, Yoon Kim, Xixin Wu, Helen Meng, and James Glass. 2023a. <a href="#">Search augmented instruction learning</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 3717–3729, Singapore. Association for Computational Linguistics.	
	Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. 2023b. An empirical study of catastrophic forgetting in large language models during continual fine-tuning. <i>arXiv preprint arXiv:2308.08747</i> .	
	Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2024. <a href="#">ExpertQA: Expert-curated questions and attributed answers</a> . In <i>Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 3025–3045, Mexico City, Mexico. Association for Computational Linguistics.	
	Luca Massarelli, Fabio Petroni, Aleksandra Piktus, Myle Ott, Tim Rocktäschel, Vassilis Plachouras, Fabrizio Silvestri, and Sebastian Riedel. 2020. <a href="#">How decoding strategies affect the verifiability of generated text</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2020</i> , pages 223–235, Online. Association for Computational Linguistics.	
	Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. <a href="#">On faithfulness and factuality in abstractive summarization</a> . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 1906–1919, Online. Association for Computational Linguistics.	
	Sewon Min, Kalpesh Krishna, Xinxin Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. <a href="#">FActScore: Fine-grained atomic evaluation of factual precision in long form text generation</a> . In <i>Proceedings of the</i>	

926		2023 Conference on Empirical Methods in Natural Language Processing, pages 12076–12100, Singapore. Association for Computational Linguistics.	
927			
928			
929	Kushan Mitra, Dan Zhang, Sajjadur Rahman, and		
930	Estevam Hruschka. 2025. <a href="#">Factlens: Benchmarking fine-grained fact verification</a> . <i>Preprint</i> ,		
931	arXiv:2411.05980.		
932			
933	Fedor Moiseev, Zhe Dong, Enrique Alfonseca, and Mar-		
934	tin Jaggi. 2022. <a href="#">SKILL: Structured knowledge infusion for large language models</a> . In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 1581–1588, Seattle, United States. Association for Computational Linguistics.		
935			
936			
937			
938			
939			
940			
941	OpenAI. 2025a. <a href="#">Gpt-5 system card</a> . Technical report, OpenAI. Accessed: 2026-01-05.		
942			
943	OpenAI. 2025b. Introducing gpt-4.1 in the api. <a href="https://openai.com/index/gpt-4-1/">https://openai.com/index/gpt-4-1/</a> .		
944			
945	OpenAI. 2025c. <a href="#">Openai o3 and o4-mini system card</a> . Technical report, OpenAI.		
946			
947	Artidoro Pagnoni, Vidhisha Balachandran, and Yulia		
948	Tsvetkov. 2021a. <a href="#">Understanding factuality in abstractive summarization with FRANK: A benchmark for factuality metrics</a> . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 4812–4829, Online. Association for Computational Linguistics.		
949			
950			
951			
952			
953			
954			
955	Artidoro Pagnoni, Vidhisha Balachandran, and Yulia		
956	Tsvetkov. 2021b. <a href="#">Understanding factuality in abstractive summarization with frank: A benchmark for factuality metrics</a> . <i>Preprint</i> , arXiv:2104.13346.		
957			
958			
959	Leonardo F. R. Ribeiro, Mengwen Liu, Iryna Gurevych,		
960	Markus Dreyer, and Mohit Bansal. 2022. <a href="#">FactGraph: Evaluating factuality in summarization with semantic graph representations</a> . In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 3238–3253, Seattle, United States. Association for Computational Linguistics.		
961			
962			
963			
964			
965			
966			
967			
968	Wooseok Seo, Seungju Han, Jaehun Jung, Benjamin		
969	Newman, Seungwon Lim, Seungbeen Lee, Ximing		
970	Lu, Yejin Choi, and Youngjae Yu. 2025. Verifying the verifiers: Unveiling pitfalls and potentials in fact verifiers. <i>arXiv preprint arXiv:2506.13342</i> .		
971			
972			
973	Shuzheng Si, Haozhe Zhao, Gang Chen, Cheng Gao,		
974	Yuzhuo Bai, Zhitong Wang, Kaikai An, Kangyang		
975	Luo, Chen Qian, Fanchao Qi, Baobao Chang, and		
976	Maosong Sun. 2025a. <a href="#">Aligning large language models to follow instructions and hallucinate less via effective data filtering</a> . In <i>Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 16469–16488, Vienna, Austria. Association for Computational Linguistics.		
977			
978			
979			
980			
981			
982			
	Shuzheng Si, Haozhe Zhao, Cheng Gao, Yuzhuo Bai,		983
	Zhitong Wang, Bofei Gao, Kangyang Luo, Wenhao		984
	Li, Yufei Huang, Gang Chen, et al. 2025b. Teaching		985
	large language models to maintain contextual faithfulness via synthetic tasks and reinforcement learning. <i>arXiv preprint arXiv:2505.16483</i> .		986
			987
			988
	Liyan Tang, Philippe Laban, and Greg Durrett. 2024a.		989
	Minicheck: Efficient fact-checking of llms on grounding documents. In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing</i> , pages 8818–8847.		990
			991
			992
			993
	Liyan Tang, Igor Shalyminov, Amy Wong, Jon Burnsky,		994
	Jake Vincent, Yu’an Yang, Siffi Singh, Song Feng,		995
	Hwanjun Song, Hang Su, Lijia Sun, Yi Zhang, Saab		996
	Mansour, and Kathleen McKeown. 2024b. <a href="#">TofuEval: Evaluating hallucinations of LLMs on topic-focused dialogue summarization</a> . In <i>Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 4455–4480, Mexico City, Mexico. Association for Computational Linguistics.		997
			998
			999
			1000
			1001
			1002
			1003
			1004
	Katherine Tian, Eric Mitchell, Huaxiu Yao, Christo-		1005
	pher D Manning, and Chelsea Finn. 2024. <a href="#">Fine-tuning language models for factuality</a> . In <i>The Twelfth International Conference on Learning Representations</i> .		1006
			1007
			1008
			1009
	Michael Völske, Martin Potthast, Shahbaz Syed, and		1010
	Benno Stein. 2017. <a href="#">TL;DR: Mining Reddit to learn automatic summarization</a> . In <i>Proceedings of the Workshop on New Frontiers in Summarization</i> , pages 59–63, Copenhagen, Denmark. Association for Computational Linguistics.		1011
			1012
			1013
			1014
			1015
	Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020.		1016
	<a href="#">Asking and answering questions to evaluate the factual consistency of summaries</a> . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 5008–5020, Online. Association for Computational Linguistics.		1017
			1018
			1019
			1020
			1021
	Yuxia Wang, Revanth Gangi Reddy, Zain Muhammad		1022
	Mujahid, Arnav Arora, Aleksandr Rubashevskii,		1023
	Jiahui Geng, Osama Mohammed Afzal, Liang-		1024
	ming Pan, Nadav Borenstein, Aditya Pillai, Isabelle		1025
	Augenstein, Iryna Gurevych, and Preslav Nakov. 2024. <a href="#">Factcheck-bench: Fine-grained evaluation benchmark for automatic fact-checkers</a> . <i>Preprint</i> , arXiv:2311.09000.		1026
			1027
			1028
			1029
	Zhitong Wang, Cheng Gao, Chaojun Xiao, Yufei Huang,		1030
	Shuzheng Si, Kangyang Luo, Yuzhuo Bai, Wenhao		1031
	Li, Tangjian Duan, Chuancheng Lv, Guoshan		1032
	Lu, Gang Chen, Fanchao Qi, and Maosong Sun. 2025. <a href="#">Document segmentation matters for retrieval-augmented generation</a> . In <i>Findings of the Association for Computational Linguistics: ACL 2025</i> , pages 8063–8075, Vienna, Austria. Association for Computational Linguistics.		1033
			1034
			1035
			1036
			1037
			1038

- 1039 Yumo Xu and Mirella Lapata. 2019. Weakly supervised  
1040 domain detection. *Transactions of the Association  
1041 for Computational Linguistics*, 7:581–596.
- 1042 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang,  
1043 Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao,  
1044 Chengen Huang, Chenxu Lv, Chujie Zheng, Dayi-  
1045 heng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge,  
1046 Haoran Wei, Huan Lin, Jialong Tang, Jian Yang,  
1047 Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi  
1048 Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai  
1049 Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao  
1050 Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang,  
1051 Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan  
1052 Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao  
1053 Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xu-  
1054 ancheng Ren, Yang Fan, Yang Su, Yichang Zhang,  
1055 Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang,  
1056 Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zi-  
1057 han Qiu. 2025. [Qwen3 technical report](#). *Preprint*,  
1058 arXiv:2505.09388.
- 1059 Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu.  
1060 2023. [AlignScore: Evaluating factual consistency  
1061 with a unified alignment function](#). In *Proceedings  
1062 of the 61st Annual Meeting of the Association for  
1063 Computational Linguistics (Volume 1: Long Papers)*,  
1064 pages 11328–11348, Toronto, Canada. Association  
1065 for Computational Linguistics.
- 1066 Huajian Zhang, Yumo Xu, and Laura Perez-Beltrachini.  
1067 2024. [Fine-grained natural language inference based  
1068 faithfulness evaluation for diverse summarisation  
1069 tasks](#). In *Proceedings of the 18th Conference of the  
1070 European Chapter of the Association for Computa-  
1071 tional Linguistics (Volume 1: Long Papers)*, pages  
1072 1701–1722, St. Julian’s, Malta. Association for Com-  
1073 putational Linguistics.
- 1074 Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao  
1075 Liu, Tingchen Fu, Xinting Huang, Enbo Zhao,  
1076 Yu Zhang, Chen Xu, Yulong Chen, Longyue Wang,  
1077 Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi.  
1078 2025. [Siren’s song in the ai ocean: A survey on  
1079 hallucination in large language models](#). *Preprint*,  
1080 arXiv:2309.01219.

1081  
1082  
1083  
1084  
1085  
1086  
1087  
1088  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129

## Appendix

### A InFi-Check-FG Construction Details

#### A.1 Detailed Hallucination Categorization and Error construction Examples

We provide a full hallucination categorization along with illustrative examples for each error construction methodology in Table 13.

#### A.2 Grounded Claim Dataset Construction

Our pipeline starts with an arbitrary document dataset, which can include pre-existing supported claims or simply the documents themselves. Due to the limitations of existing document-claim datasets in terms of length (news datasets being relatively short, and academic paper dataset being excessively long), combined with challenges that some automatically extracted claims lack complete support from the document (such as the TL;DR dataset (Völske et al., 2017)), we opted to construct our dataset beginning with the document itself, which also streamlines the extraction of evidence and proves the broad usability of our method. Specifically, for each document in the dataset, we prompt LLMs to generate a series of document-grounded claims that capture the core factual content. In our implementation, we leverage a summarization-style objective for this generation process. This approach is chosen because such condensed representations naturally require every generated claim to be strictly grounded in the source context, providing an ideal foundation for fact-checking data. Section 5.2 demonstrates that our pipeline and model possess strong generalizability to the evaluation of other document-grounded generation tasks. Subsequently, we apply an extract-and-rewrite process to ensure the faithfulness of each claim sentence: 1) LLMs are prompted to locate the grounding sentences as evidence from the source document for each sentence in these claims. 2) A voting mechanism involving three LLMs is applied to determine whether a claim sentence is sufficiently supported by its grounding sentences. If a sentence lacks adequate support, it undergoes a rewriting process, and the voting process repeats until complete support from the reference is achieved. The aforementioned process also serves as the preparation for the synthetic model output. Ultimately, we achieve a set of claims, each with sentence-level grounding evidence. Importantly, due to the flexibility of the base dataset, our

pipeline can be applied to any document dataset, supporting the scalability of InFi-Check.

#### A.3 Document Dataset Selection

We adopt the BBC News Summary Dataset (Gupta et al., 2022) and the DetNet Wikipedia Dataset (Xu and Lapata, 2019) as the initial document datasets. These datasets offer a diverse range of documents covering news and encyclopedic content, classified by domain, with a varied length distribution. We ensured diversity by choosing documents from different domains and lengths.

The BBC News Summary Dataset (Gupta et al., 2022) consists of extractive summaries only, so we did not use it as the original claim. This data set categorizes news articles into five distinct categories: business, entertainment, politics, sports, and technology. The DetNet Wikipedia Dataset (Xu and Lapata, 2019) is designed for domain detection, with Wikipedia data labeled for seven domains: “Business and Commerce” (BUS), “Government and Politics” (GOV), “Physical and Mental Health” (HEA), “Law and Order” (LAW), “Lifestyle” (LIF), “Military” (MIL), and “General Purpose” (GEN). For the BBC News Summary Dataset, we randomly selected 150 documents from each category. For the DetNet Wikipedia Dataset, we select 100 documents from each domain. The document length in both datasets varies greatly. To ensure the models have a certain degree of robustness, but also efficiency while training, we filtered documents to have lengths within the range of 300 to 1000 words, and prompted the language models (LLMs) to control the claim length within the range of  $[100, \min(\text{doc\_len}/3 + 10, 200)]$  words.

#### A.4 LLM Selection in InFi-Check Pipeline

To address potential biases where a single model might favor its own generated text, and to avoid issues where training exclusively with one model’s outputs might cause out-of-domain problems for texts generated by other models, we utilized two different sets of LLMs at each step of our pipeline. In addition, the models used for generation and evaluation were different. Usage of LLMs is outlined in Table 7.

#### A.5 Curation for InFi-Check-FG

Based on the results of human evaluation, we design a set of strict filtering prompts to further curate the validation and test sets. For each instance, we

1130  
1131  
1132  
1133  
1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147  
1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177

Task	LLM Set 1	LLM Set 2
Summarization	GPT-4o	DeepSeek-R1
Reference extraction	Claude-3.7-Sonnet	GPT-4o
Support determination	(Claude-3.7-Sonnet, Qwen-2.5, Gemini-1.5)	(GPT-4o, Qwen-2.5, Gemini-2.0)
Rewriting	Claude-3.7-Sonnet	GPT-4o
Error data construction	GPT-4o	Claude-3.7-Sonnet

Table 7: Usage of LLMs in dataset construction

	#doc	#doc-pair	#PredE	#EntE	#CircE	#CorefE	#LinkE	#OutE	#NoE
InFi-Check-TR	1263	15660	2946	1998	1386	2889	1495	1457	3489
InFi-Check-FG	190	519	85	21	29	64	30	38	252

Table 8: Statistics of the dataset constructed from InFi-Check

employ gpt-4.1 to independently check the presence of any potential generation errors categorized in human evaluation, including but not limited to incorrect extraction of grounding sentences and misclassification of hallucination or error types. Each possible error is examined in isolation through dedicated prompts, and an instance is retained only if it is verified to be fully correct across all checks. See Appendix F for filtering prompts. We intentionally design this filtering procedure to be conservative, prioritizing precision over recall. While it may discard some instances that are in fact correct, it substantially reduces the risk of retaining flawed samples. Empirically, this process is able to identify all of the errors observed in human evaluation (flawed summary, false negative error, wrong error type, wrong error reasoning). We apply this filtering pipeline to both the validation and test sets, resulting in a final curated set of 519 instances that constitute InFi-Check-FG.

## B Statistics of InFi-Check Datasets

Table 8 shows the overall statistics of the dataset constructed using InFi-Check pipeline, in which the train set refers to InFi-Check-TR and the filtered test set refers to InFi-Check-FG. The two datasets are derived from different document sets to avoid data leakage.

## C Implementation Details

### C.1 Baseline Details

**ClearCheck** (Seo et al., 2025) is fine-tuned from Llama-3.1-8B-Instruct, it improves verification robustness through multi-task training tailored for hallucination detection across various grounded

generation scenarios. **AlignScore** (Zha et al., 2023) is a holistic metric that evaluates factual consistency through a unified alignment function trained across diverse NLP tasks such as NLI and QA. We adopt the largest and best-performing version (AlignScore-large) as our baseline. **FactCG** (Lei et al., 2025) enhances model performance by generating complex training data via multi-hop reasoning on context graphs extracted from documents. We use the best-performing version (FactCG-DBT) as our baseline. **MiniCheck** (Tang et al., 2024a) achieves state-of-the-art binary performance by utilizing a novel document-claim pair synthesis method to train lightweight fact-checkers. We use the largest and best-performing version (Bespoke-MiniCheck-7B) as our baseline.

### C.2 Benchmark and Metric Details

For InFi-Check-FG, we calculate the accuracy of each error type (as well as “No Error”), and report their average, the balanced accuracy. The “accuracy” is defined as whether the model correctly points out the specific error type in samples containing this type of error. If a model reports a sample to be hallucinated, but with the wrong error type, it is judged as incorrect.

The **FRANK** benchmark focuses on factual consistency in abstractive summarization using the CNN/DM and XSum news datasets. It provides a manual annotated fine-grained typology of errors found in the outputs of multiple state-of-the-art summarization models, offering a rigorous testbed for news-domain grounding. To facilitate a standardized comparison across diverse benchmarks, we also report balanced accuracy on the FRANK dataset instead of correlation-based metrics, which

Model	BAcc(Nor.)	BAcc(Str.)	SAR(Avg.)
<b>The Open-Source Models</b>			
Llama-3.1-8B-Instruct†	27.69	21.33	72.57
Llama-3.1-8B-Instruct*†	23.79	18.12	79.41
Qwen3-8B†	47.94	41.08	86.50
Qwen3-8B*†	50.49	44.78	87.22
<b>The State-of-the-Art LLMs</b>			
Claude-3.7-Sonnet	53.07	48.48	92.21
Claude-3.7-Sonnet*	49.19	45.12	94.28
DeepSeek-V3.2-NoThink	59.64	53.94	89.43
DeepSeek-V3.2-NoThink*	59.19	52.31	87.33
GPT-4o	60.50	56.57	94.97
GPT-4o*	64.69	60.27	80.19
GPT-4.1	54.45	49.44	78.17
GPT-4.1*	58.91	55.13	94.87
GPT-5	49.06	44.41	92.05
GPT-5*	46.35	42.28	92.58
o3†	58.00	55.07	95.48
o3*†	60.77	56.76	94.60
<b>InFi-Checker (Ours)</b>			
InFi-Checker-Llama	90.91	85.49	94.02
InFi-Checker-Qwen	92.34	87.74	94.92

Table 9: Normal(Nor.) and strict(Str.) balanced accuracy (%) and averaged SAR on InFi-Check-FG. The reported SAR(Avg.) is the average of error-type specific SAR across all evaluable error types. Models marked with † contain types with zero accuracy (excluded from the average), which may lead to an optimistic estimation of their SAR. \* means one-shot settings.

Setting	BAcc
Justification + Correction	90.91
Justification only	86.30
Correction only	83.17
No interpretable reasoning	75.02

Table 10: Ablation study of InFi-Checker-Llama on InFi-Check-FG for assessing justification and correction separately. Other component, such as evidence and sentence-by-sentence reasoning patterns, is preserved throughout the experiment.

often suffer from limited comparability across different model scales. Given the complexity of FRANK—where each instance is annotated by multiple experts and often contains overlapping error types—we adopt a multi-annotator consensus criterion for evaluation. Specifically, a model’s prediction is considered correct if the predicted error type aligns with any of the labels assigned by the human annotators for that specific instance. This approach accounts for the inherent subjectivity and legitimate label diversity in fine-grained hallucination detection. Furthermore, while the FRANK dataset contains fine-grained error categories, many instances exhibit multiple co-occurring error types, which can introduce noise into categorical classification. To ensure a consistent and robust evaluation, we re-

formulate the task into a binary consistency check for metric calculation. We map all fine-grained error types to a "hallucinated" category, treating each instance as either "hallucinated" or "factual." To mitigate the impact of class imbalance within the dataset, we report Balanced Accuracy, providing a more reliable reflection of the model’s discriminative performance across both factual and non-factual claims.

For our evaluation on binary fact-checking benchmarks, we acknowledge the data quality issues identified by Seo et al. (2025), which notes that several widely used datasets contain significant annotation noise, linguistic ambiguities, and skewed label distributions. To ensure a robust and reliable assessment, we adopt the filtered and re-processed versions of these benchmarks from Seo et al. (2025). To maintain strict parity with established baselines and facilitate a direct comparison with state-of-the-art methods, we follow the evaluation protocol of Seo et al. (2025) and report the Macro F1 score. The six binary fact-checking benchmarks are: **ClaimVerify** audits the generative search engine task, using responses from commercial systems (e.g., Bing Chat) across diverse queries from Google history and Reddit. It emphasizes the accuracy of in-line citations and the verifiability of claims against real-world web sources. **ExpertQA** targets high-stakes, long-form question answering across 32 specialized fields (e.g., medicine, law). It features expert-curated questions and claims generated by representative LLMs, requiring models to verify professional, domain-specific knowledge with high precision. **Factcheck-Bench** is a comprehensive benchmark for open-domain document-level factuality, evaluating LLM-generated responses at the claim, sentence, and document levels. It encompasses a wide array of general-purpose topics, aiming to test the end-to-end verification capabilities of automatic systems. **REVEAL** focuses on complex Chain-of-Thought (CoT) reasoning in open-domain settings. It provides step-level labels for attribution and logical correctness, testing whether a model can verify the intermediate reasoning claims of a language model’s answer. **MediaSum** and **MeetingBank** are subset from the summarization dataset ToFuEval. MediaSum centers on topic-focused dialogue summarization within the media interview domain. It challenges models to maintain factual integrity when distilling multi-party conversations into concise, grounded summaries. MeetingBank

Label	Proportion(%)
No Problem	78
Flaws in Claim	11
Incomplete Grounding	6
False Negative Error	3
Wrong Error Type	1
Wrong Error Justification	1

Table 11: Human evaluation results on a sample of 100 instances from our dataset.

Model	Cost(\$)
<b>The State-of-the-Art LLMs</b>	
GPT-5	23.9
Claude-3.7-Sonnet	12.7
o3	12.4
<b>InFi-Checker (Ours)</b>	
<b>InFi-Checker-Llama</b>	3.9
<b>InFi-Checker-Qwen</b>	4.4

Table 12: Comparison of costs of InFi-Checker and LLMs on InFi-Check-FG. InFi-Checker cost is calculated with \$0.8 per GPU hour.

focuses on professional meeting transcripts and evaluates the consistency of summaries generated from lengthy, informal, and multi-speaker interactions. It serves as a robust test for grounding claims in complex, non-structured dialogue data.

### C.3 Training Details of InFi-Checker

We fine-tune both Llama-3.1-8B-Instruct and Qwen3-8B for 3 epochs on InFi-Check-TR, using a batch size of 32 and the Adam optimizer. The learning rate follows a cosine-decay from  $1e - 5$  to  $1e - 6$ , and we set the warm-up fraction to 0.1.

## D Further Results

Table 9 shows the full strict accuracy and SAR result on the FRANK benchmark, including both zero-shot and one-shot results of LLMs. We also conducted an ablation study to isolate the individual contributions of justifications and corrections. As shown in Table 10, while retaining either component still yields competitive results, the integration of both processes is essential for achieving optimal performance.

## E Data Structure Example

We provide an example of data from InFi-Check-FG in Figure 4, which demonstrates the data structure of InFi-Check.

## F Prompts

Figure 5, Figure 6, Figure 7 and Figure 8 show the prompts in the InFi-Check pipeline. Figure 9 shows the prompt used for the LLM baselines in our experiment. Note that in our preliminary experiments, we observed that increasing the number of few-shot examples to 2 or 3 could adversely affect performance due to the extended context and reasoning process length within InFi-Check-FG. Consequently, we limit our experimental setup to zero-shot and one-shot configurations for better model performance. Figure 10 is the prompt used for filtering samples in InFi-Check-FG.

## G Human Evaluation

We did human evaluation on a sample of 100 for InFi-Check-TR. Annotators (all holding PhD or Master’s degrees) are instructed to check whether the total claim and output fall into any of the given mistakes. The instructions for detecting defined mistakes are:

- **Flaws in Claim:** A summary error exists if: - A summary sentence is annotated as "supported or not: YES", - BUT the summary contains factual errors that contradict or mismatch with the original document (including entity, time, location, numerical, or core semantic errors). Important: - If ANY single summary sentence has this issue, the ENTIRE data sample must be considered problematic.
- **Incomplete Grounding:** A grounding error exists if ANY of the following is true: - The listed "related sentence(s) from the document" are NOT sufficient to support the summary sentence, and there exist other ESSENTIAL sentences in the document that must be cited. - Any listed grounding sentence does NOT actually appear in the original document (not minor truncation or paraphrase, but clearly non-existent). - The same grounding sentence is explicitly duplicated (appears two or more times). Note: - Missing an essential supporting sentence is a definite error.
- **False Negative Error:** A false negative error exists if: - A summary sentence is annotated as "supported or not: NO", - BUT the summary is factually correct and fully supported by the document, with NO factual discrepancies in its core meaning. This applies only when the summary is clearly correct.

```

[
  {
    "summary sentence": "The column in Piazza Santa Maria Maggiore in Rome, crowned with a statue of the Virgin in 1614, set a precedent for many European columns.",
    "related sentence(s) from the document": [
      "The column in Piazza Santa Maria Maggiore in Rome was one of the first.",
      "Within decades it served as a model for many columns in Italy and other European countries."
    ],
    "supported or not": "YES",
    "reason": "The summary sentence faithfully reflects the related sentences.",
    "error type": "No Error"
  },
  {
    "summary sentence": "The first Marian column north of the Alps was Munich's Mariensäule in 1714, inspiring similar structures in Prague and Vienna.",
    "related sentence(s) from the document": [
      "The first column of this type north of the Alps was the Mariensäule built in Munich in 1638 to celebrate the sparing of the city from both the invading Swedish army and the plague.",
      "It inspired for example Marian columns in Prague and Vienna, but many others also followed very quickly."
    ],
    "supported or not": "NO",
    "reason": "This sentence is not supported by the related sentence(s).\n- Location: '1714'.\n- Explanation: The year of the construction of Munich's Mariensäule was changed from 1638 to 1714, falsely altering the historical timeline.\n- Correction: The first Marian column north of the Alps was Munich's Mariensäule in 1638, inspiring similar structures in Prague and Vienna.",
    "error type": "Circumstance Error"
  }
],
{
  "summary sentence": "The Prague column, built post-Thirty Years' War, was destroyed in 1918 due to its association with Habsburg rule.",
  "related sentence(s) from the document": [
    "The Prague column was built in Old Town Square (Staroměstské náměstí) shortly after the Thirty Years' War in thanksgiving to the Virgin Mary Immaculate for helping in the fight with the Swedes.",
    "Unfortunately, many Czechs later connected its placement and erection with the hegemony of the Habsburgs in their country, and after declaring the independence of Czechoslovakia in 1918 a crowd of people pulled this old monument down and destroyed it in an excess of revolutionary fervor."
  ],
  "supported or not": "YES",
  "reason": "All content in the summary sentence is accurately derived from the related sentences.",
  "error type": "No Error"
}
]

```

Figure 4: An example of the data in InFi-Check-FG, the data is truncated due to space limitations.

<p>1388</p> <p>1389</p> <p>1390</p> <p>1391</p> <p>1392</p> <p>1393</p> <p>1394</p> <p>1395</p> <p>1396</p> <p>1397</p> <p>1398</p> <p>1399</p> <p>1400</p> <p>1401</p> <p>1402</p> <p>1403</p>	<ul style="list-style-type: none"> <li>• <b>Wrong Error Type:</b> An error-type mismatch exists if: - The summary is indeed incorrect, - BUT the annotated “error type” is not the most appropriate one according to the SIS-Fact definitions below, - AND a clearly better error type applies. Only flag this error when the mismatch is obvious. (We also provide the error taxonomy for the annotators to refer to.)</li>   <li>• <b>Wrong Error Justification:</b> A reasoning error exists if: - The annotated error type is correct, - BUT the "reason" field contains incorrect analysis, such as: - Misidentifying the nature or location of the error; - Proposing a correction that is STILL factually wrong; - Explaining the error in a way that contradicts the original document. Note: - The problem must be</li> </ul>	<p>substantive and factual.</p> <p>As shown in Table 11, most errors originated in the initial claim generation or evidence extraction phases, underscoring the inherent difficulty of grounded generation and highlighting the robustness of our controlled hallucination pipeline.</p> <p><b>H Cost Efficiency Evaluation</b></p> <p>Table 12 shows the result for cost efficiency analysis. We compared the computational cost of InFi-Checker and baseline LLMs. The cost of InFi-Checker is converted using the prediction time on the GPUs, while the cost of LLMs is computed through tokens in API calls.</p>	<p>1404</p> <p>1405</p> <p>1406</p> <p>1407</p> <p>1408</p> <p>1409</p> <p>1410</p> <p>1411</p> <p>1412</p> <p>1413</p> <p>1414</p> <p>1415</p> <p>1416</p>
---	--	---	---

Categorization		Method	Example
Intrinsic Errors	Semantic Frame Errors	Predicate Error (PredE)	Swapping Relation Masking Its impact on theaters is now <b>emerging</b> . Its absence from theaters is now <b>receding</b> .
			Modifying Predictions Despite it all, 2023 <b>should</b> still <b>reach</b> the \$9 billion in domestic gross hoped for this year. Despite it all, 2023 <b>reached</b> the \$9 billion in domestic gross hoped for this year.
		Entity Error (EntE)	Swapping Entities Her last stage role was in <b>My Fair Lady</b> . Her last stage role was in <b>Bless This House</b> .
			Compressing Words In <b>France and Italy</b> , he wrote his last work. In <b>France</b> , he wrote his last work.
	Circumstance Error (CircE)	Swapping Circumstances The shooting left <b>10</b> students and <b>2</b> teachers dead. The shooting left <b>2</b> students and <b>10</b> teachers dead.	
	Discourse Errors	Co-reference Error (CorefE)	Swapping Pronouns <b>Gonzales</b> was also indicted. <b>She</b> was also indicted.
			Merging Sentences The charges were first reported by the San Antonio Express-News. <b>District Attorney Christina Mitchell</b> did not return requests for comment. The charges were first reported by the San Antonio Express-News, <b>who</b> did not return requests for comment.
Discourse Link Error (LinkE)		Reverse Logical Relationship The six-month Hollywood labor disruption has finally ended. <b>Immediately following the settlement</b> , Disney announced delays in its upcoming release schedule. <b>After</b> the announcement of the delays in Disney's upcoming release schedule, the six-month Hollywood labor disruption has finally ended.	
Extrinsic Errors		Introducing Extrinsic Information	Robert escaped to Visegrád disguised as a civilian, aided by Nicholas, son of Radoslav, who defended him against five attackers. Robert escaped to Visegrád disguised as a civilian, aided by Nicholas, son of Radoslav, <b>a renowned swordsman known for his exceptional skill in battle</b> , who defended him against five attackers.

Table 13: Categories of the error data. The blue part is the selected text for modification, and the red part is the modified text.

I'll provide you with a document. Your task is to write a short summary for this document according to the following requirements:

1. The length of the summary should be within <WORD\_CONSTRAIN> words.
2. Every sentence in the summary should be directly supported by the content of the document.
3. For each event, make sure every important entity such as person, location and time is kept in the summary, especially entities that occurs in parallel.
4. When doing simplification, make sure each complex event or idea remains true to the original meaning. Avoid over-simplification that leads to in-consistency with the origin document.

Document:<Document>

Directly output the summary without any extra words.

Figure 5: Prompt for writing summaries.

Here is a document with a corresponding summary. Your task is to analyze the summary sentence by sentence. For each sentence in the summary, provide the exact sentences from the document that supports the summary sentence. If the summary comes from multiple sentences, report all sentences.

You should ensure that sentences are considered as individual units based on punctuation. Specifically:

- Treat each sentence as ending at a period (.), question mark (?), or exclamation mark (!), even if multiple sentences are enclosed within quotation marks.
- Do not truncate any sentence. If a portion of a sentence is extracted (e.g., ending at a comma, semicolon, or any punctuation other than ., ?, !), you must include the rest of the sentence so that the entire sentence is fully reproduced.

You should only respond in format as described below. Do not return anything else. START YOUR RESPONSE WITH '['

Return the result as a Python list of dictionaries, where each dictionary has the following keys:

"summary sentence": The sentence from the summary.

"sentences from the document": A Python list of the exact supporting sentences from the document.

Ensure that the sentences are in the same order as they appear in the original document. Each sentence should be reported fully, without any omission.

Figure 6: Prompt for locating the grounding sentences.

Here are some pieces from the SOURCE\_DESCRIPTION, and a summary sentence of it. Your task is to:

1. Compare the summary with the document and determine whether if the summary is fully supported by the content in the document. Specifically, verify if all key points made in the summary are traceable back to the sentence in the document. State whether the summary is fully supported with 'YES' or 'NO'.
2. If the answer is 'NO', revise the summary to align it fully with the document. Make sure the fluency and grammar correctness of the revised sentence, and ensure it accurately reflects the information.

You should only respond in format as described below. Do not return anything else. START YOUR RESPONSE WITH '['

Return the result as a Python list of dictionaries, where each dictionary has the following keys:

"summary sentence": The sentence from the summary.

"sentences from the document": A Python list of the exact supporting sentences from the document.

Ensure that the sentences are in the same order as they appear in the original document. Each sentence should be reported fully, without any omission.

Figure 7: Prompt for determining whether a summary sentence is sufficiently supported by its grounding sentences. The purple part is only used in the LLM for re-writing.

Here is a document with a summary. Please create a fake summary based on the origin summary by the following steps:

Instruction:

1. Analyze the document and identify sentences that contain future predictions (e.g., those using modal verbs like 'will,' 'might,' or phrases like 'predict,' 'suppose').
2. Select a sentence where the prediction is the main clause, not just a subordinate clause, and modifying the prediction into a factual statement will result in the most significant change in meaning. You can modify this sentence so that the prediction is completely transformed into a factual statement about an event that has already occurred. Ensure the modified sentence is grammatically correct and fully removes any speculative language.
3. Based on the changed sentence, modify some part of the summary to include the fake information in the changed sentence, so the summary cannot be fully supported by the origin document.

Make sure the new summary should not be fully supported by the document, and not change any other part in the summary besides those associated with the modification.

You should only respond in format as described below. Do not return anything else. START YOUR RESPONSE WITH ‘{{‘.

Return the result as a Python dictionary with the following keys:

Format:

- "original text in summary": The original sentence containing the prediction from the document.
- "chosen element": The chosen prediction or future-oriented statement in the original text.
- "modification explanation": Description of the modification.
- "modified element": The new factual statement replacing the prediction.
- "modified text": The sentence after the modification, now a factual statement.
- "explanation": A clear explanation of how the meaning of the original text has been altered.
- "full text of modified summary": The full text of the modified summary.
- "wrong information": Point out the specific wrong information introduced in the summary after the modification.

Replace any line breaks in the values with '\n' so that the dictionary can be parsed using eval().

Figure 8: Prompt for generating InFi-Check-FG. The colored part is construction-method specific.

**\*\*Task:\*\***

Evaluate the given summary by comparing it with the original document and identify any errors. These errors may include incorrect information, over-simplifications, misrepresentations, or other discrepancies. The possible types of errors to consider are as follows: Predicate Error, Entity Error, Circumstance Error, Co-reference Error, Discourse Link Error, Extrinsic Error.

**\*\*Instructions:\*\***

You are provided with the full text of the original document and a summary that may contain errors. You should analyze the summary sentence by sentence and returning the results in the following Python list format:

Each item in the list should correspond to a summary sentence and be represented as a Python dictionary with the following keys:

- "summary sentence": The summary sentence.
- "related sentence(s) from the document": A list of sentences from the original document that support the summary sentence, if the summary sentence is fully supported. If the summary sentence contains an error, list the sentences needed to point out the error.
- "supported or not": "YES" if the summary sentence is fully supported by the related sentences, otherwise "NO".
- "reason": A brief analysis explaining whether the summary sentence is supported or not. If the summary sentence is not supported, specify where the error occurs, explain the incorrect information conveyed in the summary, and provide a corrected version of the sentence.
- "error type": The type of error found (choose from the types listed above), or "No Error" if the sentence is fully supported.

Document: <Document>

Summary: <Summary>

Figure 9: Prompt for evaluating LLM baseline.

You are an expert auditor for synthetic data quality. Your task is to strictly examine whether the given synthetic annotation data contains ANY clear construction errors, according to the error categories defined below. If a clear problem exists, you MUST identify it and mark the data as problematic. If no clear problems exist, mark the data as valid.

<POSSIBLE ERROR TYPES>

<INPUT DATA>

### Final Judgment Rules

- If ANY of the five error types is found in ANY summary sentence, the entire data sample is problematic.
- If no errors are found, the data sample is valid.
- When evidence is unclear or debatable, flag an error.

### Requirements

1. For the entire data sample, judge whether it contains ANY of the 5 types of construction errors;
2. Output ONLY JSON format with the following fields:
  - is\_problematic: Boolean (True = has any construction error, False = no errors);
  - error\_details: element is a dict with "sentence\_idx", "error\_type"
  - empty if no errors;
  - reason: Str (brief summary of all errors, or "No construction errors detected").

Figure 10: Prompt for the filtering process of InFi-Check-FG. Possible error types are the same as human evaluation instructions, and the input data are the document-claim pairs.